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

Online Appendix — not for publication

Jaroslav Borovička

New York University

jaroslav.borovicka@nyu.edu

October 17, 2016

Abstract

This Online Appendix provides further background and extends the results of the main paper. It also includes lengthier derivations from the proofs. It is not intended for publication.

OA.1 Introduction

This Online Appendix contains several specific results that illuminate and extend the analysis from the paper, and lengthier derivations from the proofs. 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 OA.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 OA.3 discusses details of the information structure in the economy and general modeling of subjective beliefs in the Brownian information environment. I also explain contractual details that lead to a complete-market decentralization with role for speculative trade even in economies with constant aggregate endowment. The section also provides a change of measure result that helps express survival outcomes under agents' subjective beliefs.

Section OA.4 summarizes additional survival results that are not included in the paper: the case of multiple, mutually correlated shocks, 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. I also include a discussion of the role of parameteric restrictions that guarantee existence of an equilibrium, and provide further details on the dynamics of the Pareto share. Section OA.5 compares the survival results with those derived in Kogan, Ross, Wang, and Westerfield (2006) for economies with no intermediate consumption and a terminal consumption payout. Section OA.6 discusses in more detail possible extensions of the framework introduced in the paper, including model uncertainty and learning, and robust utility.

Finally, Sections OA.7 and OA.8 contain proofs omitted from the main text of the paper.

OA.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 detail 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.

OA.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 $\widetilde{V}$ for an agent endowed with Epstein-Zin preferences is given by

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

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 $\mathcal{R}$ 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 $\mathcal{R}_t\left(\widetilde{V}_{t+1}\right) = h^{-1}\left(E_t\left[h\left(\widetilde{V}_{t+1}\right)\right]\right)$ is not linear in $\widetilde{V}$, 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

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

that implies the recursion

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

This transformation reduces the certainty equivalence $\mathcal{R}_t(V_{t+1}) = E_t^Q V_{t+1}$ 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 $E_t^Q[V_{t+\varepsilon}]$ from (OA.1) to obtain

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

Applying the L'Hospital rule leads to

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

The function $f(C, V)$ is called the aggregator function. Integrating this expression over time and taking expectations yields

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

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

$$V_t = E_t^Q\left[\int_t^{\infty} f(C_s, V_s)\,ds\right] \tag{OA.2}$$

[1]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.

with the aggregator defined as

$$f(C,V) = \frac{\beta}{\rho}\left[(C)^{\rho}(\gamma V)^{1-\frac{\rho}{\gamma}} - (\gamma V)\right]. \tag{OA.3}$$

The aggregator $f(C_s, V_s)$ 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 $V(C)$. 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 $V(C)$ 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 (OA.2) with subjective beliefs.

OA.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 (OA.2). 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

$$\lambda_t V_t = \sup_{\nu} E_t^Q\left[\int_t^{\infty} \lambda_s F(C_s, \nu_s)\, ds\right] \tag{OA.4}$$

subject to

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

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

$$f(C,V) = \sup_{\nu \in R}\left[F(C,\nu) - \nu V\right] \tag{OA.5}$$

$$F(C,\nu) = \inf_{V \in R}\left[f(C,V) + \nu V\right]. \tag{OA.6}$$

The transformation (OA.5)–(OA.6) 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 $F(C_t, \nu_t)\, dt$, 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 (OA.5) yields the maximized instantaneous discounted surplus $f(C_t, V_t)\,dt$ of the decision maker, and the recursive utility representation aggregates the maximized surplus.

For the case of the Duffie–Epstein–Zin preferences (OA.3), transformation (OA.6) implies

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

corresponding to the felicity function specification considered in the paper.

OA.3 Information structure and subjective beliefs

OA.3.1 Information structure

Uncertainty in the economy is modeled using a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ generated by a univariate Brownian motion W. Aggregate endowment Y follows a geometric Brownian motion

$$d\log Y_t = \mu_y dt + \sigma_y dW_t \tag{OA.7}$$

with constant parameters μ_y and σ_y.

Agents know the parameters μ_y and σ_y and observe the realizations of the Brownian motion W and hence also the realizations of Y. Observing W is equivalent to observing Y when $\sigma_y > 0$ but this distinction will become material when we consider the case without aggregate uncertainty ($\sigma_y = 0$) in Section OA.3.3.

OA.3.2 Modeling subjective beliefs

Subjective beliefs are modeled as disagreement about the distribution of W (and hence of Y). Here, we show that under this Brownian information structure and under mild square integrability conditions, absolutely continuous subjective beliefs can be generally expressed using local drifts distortions of the Brownian motion W.

In line with the literature (Sandroni (2000), Blume and Easley (2006), Kogan, Ross, Wang, and Westerfield (2006, 2011), Yan (2008) and others), I impose agents' heterogenous beliefs by specifying alternative subjective probability measures Q^n. Here, I show that constructing a particular Q^n is equivalent to appropriately specifying a stochastic process u^n that describes the local evolution of the belief distortion.

In order to prevent arbitrage opportunities and other pathologies, we require subjective probability measures Q^n to be equivalent to each other (and, for convenience, also to the data generating measure P). Hence, there exist martingales M^n adapted to $\{\mathcal{F}_t\}$ that are strictly positive P-a.s. such that

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

Assume that this martingale is square integrable, i.e., $E\left[(M_t^n)^2\right] < \infty$. By the Martingale Representation Theorem (see, e.g., Øksendal (2007), Theorem 4.3.4), there exists a unique square integrable process $\widetilde{u}^n$ such that

$$M_t^n = M_0^n + \int_0^t \widetilde{u}_s^n dW_s$$

and hence, defining $u^n \doteq \widetilde{u}^n / M^n$,

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

The Girsanov Theorem then implies that the process

$$W_t^n = W_t - \int_0^t u_s^n ds$$

is a Brownian motion under Q^n. Substituting this expression into (OA.7) implies that

$$d\log Y_t = (\mu_y + \sigma_y u_s^n)\, dt + \sigma_y dW_t^n.$$

Hence, under the subjective probability measure Q^n (which we have not restricted beyond technical conditions involving absolute continuity and square integrability), the agent perceives the original Brownian motion to have a local drift distortion u_t^n, and the logarithm of the aggregate endowment to have a local trend $\widetilde{\mu}_{y,t} \doteq \mu_y + \sigma_y u_s^n$.

Reverting the argument, subjective beliefs represented by Q^n in this Brownian information environment can be generally modeled by directly specifying the processes u^n. In the paper, u^n is taken to be constant, as a particular special case. These specific belief distortions have been studied by Yan (2008), Kogan, Ross, Wang, and Westerfield (2011) and others.

OA.3.3 Contractual structure in the economy without aggregate risk

In the model economy driven by a univariate Brownian motion, dynamic trade in two suitably chosen assets provide a dynamically complete market in the sense of Harrison and Kreps (1979). When $\sigma_y > 0$, it is intuitively convenient to specify the two assets as an infinitesimal risk-free asset (with a locally safe return $r_t dt$) in zero net supply and a claim on the aggregate endowment (with return $d\log R_t$) in unit supply.

Introducing additional redundant assets into this environment does not change allocations or asset prices. Consider therefore dynamic trade in an additional asset that pays off the amount W_t and is provided in zero net supply. When $\sigma_y > 0$, the redundancy of this asset implies that a feasible decentralization involves zero positions of both agents in this asset and all trade is conducted in the infinitesimal risk-free asset and the claim on the aggregate endowment.

Now consider the case without aggregate uncertainty, $\sigma_y = 0$. In this case, trade only in the infinitesimal risk-free asset and the (now safe) claim on aggregate endowment would not dynamically span the market, as neither of the asset payoffs is exposed to the shock W. However, a proper complete-market decentralization involves trade in the claim on W that is in zero net supply and the claim on the (deterministic) aggregate endowment in unit supply. This is one feasible decentralization that supports the results in Section 4.4.3 of the paper. Speculative trade in the claims on realizations of W is voluntary due to heterogeneity in beliefs between the two agents and, despite the lack of aggregate uncertainty, generates fluctuations in the wealth distribution.

OA.3.4 Change of measure and survival under subjective beliefs

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 OA.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 the 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.

OA.4 Specific survival results

OA.4.1 Role of restrictions on the time preference parameter

Assumption A.1 imposes parameteric restrictions (17)–(18) that are sufficient for the equilibrium in the economy to exist and the individual decision problems to be well-defined both for the large agent and for the infinitesimal agent. Beyond sufficient time discounting, these conditions are immaterial for the long-run results, as the survival conditions in Propositions 3.2 and 3.4 do not depend on β, beyond the requirement that an equilibrium exists.

To get a sense how tight these conditions are quantitatively, observe first that when IES $= 1$ ($\rho = 0$), they amount to $\beta > 0$, i.e., any positive degree of impatience is sufficient. The reason is

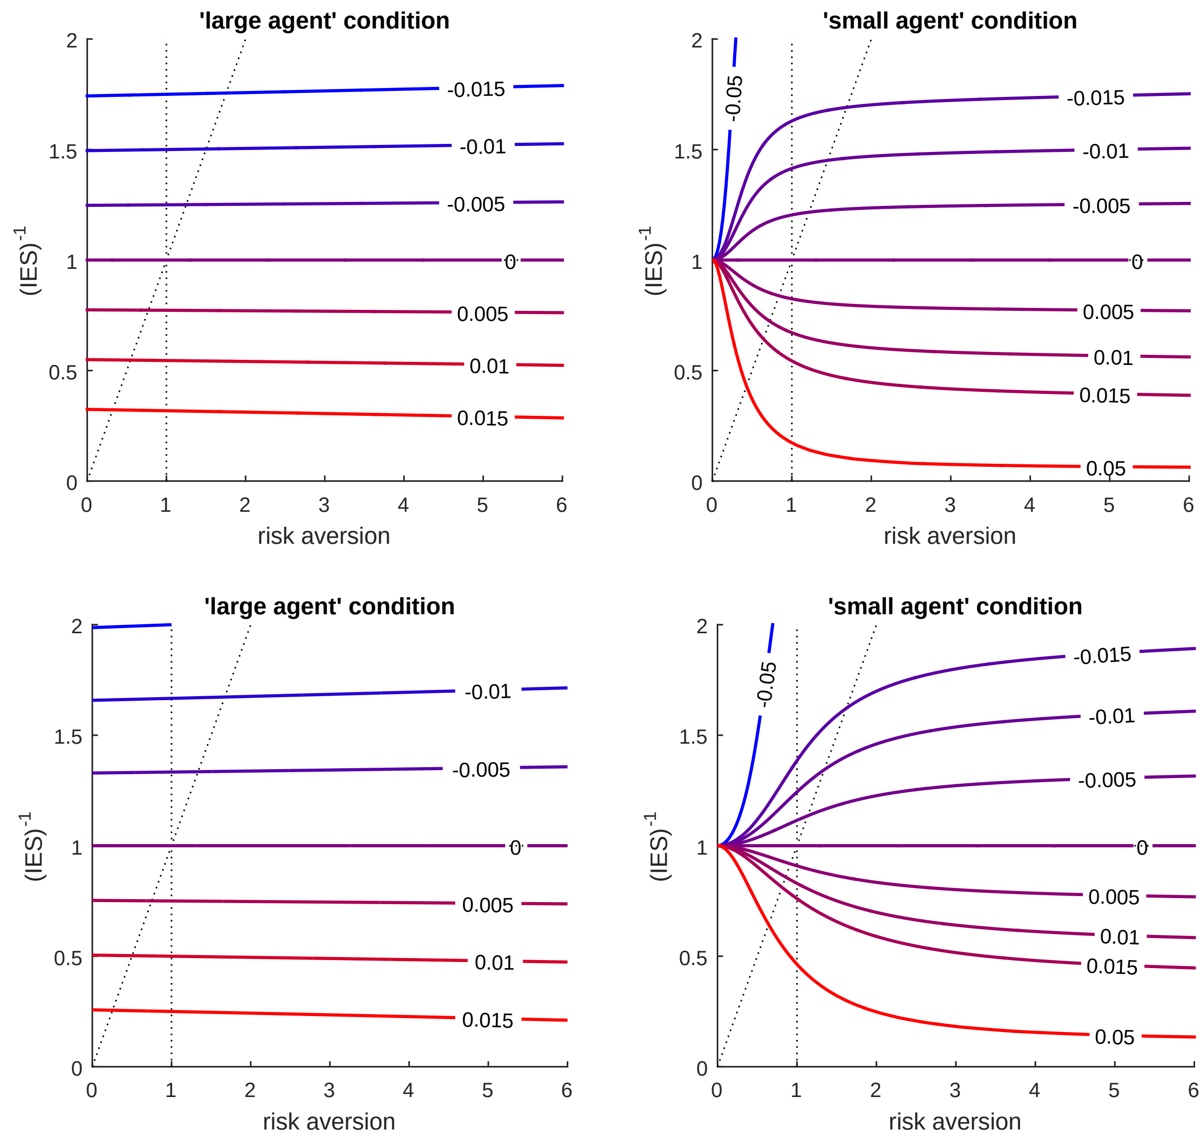

Figure OA.1: Contour plots for minimum sufficient values of β in Assumption A.1 for alternative preference parameters. The left panels ('large agent' conditions) plot the right-hand side of condition (17), while the right-panels ('small agent' conditions) plot the right-hand side of condition (18). The top row are economies with an optimistic agent 1, $u^1 = 0.1$, the bottom row economies with a pessimistic agent 1, $u^1 = -0.25$. The remaining parameters are $u^2 = 0$, $\mu_y = 0.02$, $\sigma_y = 0.02$.

that under unitary elasticity of substitution, the wealth-consumption ratios satisfy $\xi^n(\theta) = \beta^{-1}$ as under logarithmic preferences. Figure OA.1 then provides contour plots for the minimum values of β for selected economies from Figure 2 from the paper. The top row shows conditions for economies with an optimistic agent (corresponding to top left panel in Figure 2), while the bottom row shows conditions for economies with a pessimistic agent (corresponding to bottom right panel in Figure 2). Both the 'large agent' condition (17) in the left panels and the 'small agent' condition (18) in the right panels have to be satisfied.

As the graphs show, the (17)–(18) do not impose a severe restriction on the time preference parameter. The tightest restriction is the 'small agent' restriction (18) for high values of IES. This is not surprising—the negligible agent forms a speculative portfolio with a high subjective expected return, which, under a high IES, induces her to choose a high saving rate. Sufficient impatience is needed to make the consumption-wealth ratio of the negligible agent positive (compare the condition with formula (15)).

OA.4.2 Relative patience and the dynamics of the Pareto share

The survival conditions in Proposition 3.4 are stated in terms of the logarithmic growth rates of wealth. However, these conditions can also be restated in terms of the behavior of relative patience at the boundaries.

For instance, for the left boundary, the limiting discount rate of the large agent $\nu^2(\theta)$ converges to $\bar{\nu}^2$ from (OA.13) as $\theta \searrow 0$. Similarly, the limiting discount rate $\nu^1(\theta)$ for the infinitesimal agent 1 can be inferred from her portfolio problem outlined in the proof of Proposition 5.3 in equations (36)–(37), which leads to the following result for the limiting behavior of relative patience.

Proposition OA.1 *The expressions for the limiting behavior of the relative patience in Proposition 3.2 are*

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

Proof. See Section OA.8. ■

Proposition 3.2 then implies that both agent survive (conditions (i) and (ii) in that Proposition hold) when

$$\lim_{\theta^1 \searrow 0} \nu^2(\theta) - \nu^1(\theta) > \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right] \tag{OA.8}$$
$$\lim_{\theta^1 \nearrow 1} \nu^2(\theta) - \nu^1(\theta) < \frac{1}{2}\left[\left(u^1\right)^2 - \left(u^2\right)^2\right]. \tag{OA.9}$$

These conditions show that differences in patience must compensate for differences in belief distortions in order for the agents to survive. For instance, if agent 1's beliefs are less accurate than agent 2's beliefs, $\left|u^1\right| > \left|u^2\right|$, then at the left boundary, agent 2 has to be sufficiently more impatient than agent 1 to guarantee survival of agent 1.

The left panel of Figure OA.2 displays the behavior of relative patience $\nu^2(\theta) - \nu^1(\theta)$ in the interior of the state space for three different economies. Under CRRA preferences, the relative patience would be identically zero. The dash-dotted line represents a high risk aversion economy in which both survival conditions from Proposition 3.2 (equivalent to conditions (OA.8)–(OA.9)) hold and both agents survive. The dashed line corresponds to a parameterization that is close to the CRRA case when only the survival condition for the rational agent 2 is satisfied. At the left bound-

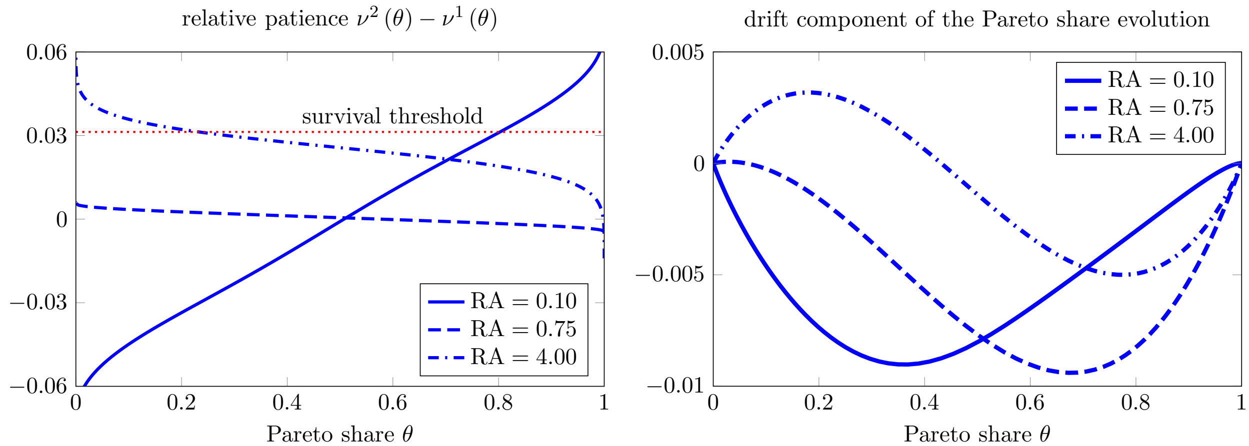

Figure OA.2: Relative patience $\nu^2(\theta) - \nu^1(\theta)$ (left panel) and the drift component of the Pareto share evolution $E_t[d\theta_t]/dt$ (right panel) as functions of the Pareto share θ. All models are parameterized by $u^1 = 0.25$, $u^2 = 0$, IES $= 1.5$, $\beta = 0.05$, $\mu_y = 0.02$, $\sigma_y = 0.02$, and differ in levels of risk aversion. The dotted horizontal line in the left panel represents the survival threshold $\frac{1}{2}(u^1)^2 - \frac{1}{2}(u^2)^2$.

ary, relative patience is not sufficiently high to exceed the 'survival threshold' $\frac{1}{2}\left[(u^1)^2 - (u^2)^2\right]$. Finally, the solid line captures a low risk aversion economy in which both attracting conditions from Proposition 3.2 hold and each of the agents dominates with a strictly positive probability.

The behavior of relative patience directly affects the dynamics of the state variable θ. An application of Itô's lemma yields

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

The right panel of Figure OA.2 depicts the impact of relative patience on the drift coefficient of the Pareto share process. The drift vanishes at the boundaries and the boundaries are unattainable (a reflection of the Inada conditions), but sufficiently large positive (negative) slopes at the left (right) boundaries assure the existence of a nondegenerate long-run distribution of the Pareto share.

OA.4.3 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, Corr $\left(dW_t^1, dW_t^2\right) = \varphi dt$. The aggregate endowment is driven by the first component of W,

$$d \log Y_t = \mu_y dt + \sigma_y dW_t^1, \qquad Y_0 > 0. \tag{OA.10}$$

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 process 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 exist in the real world, although, as discussed in Section 6.2 of the paper, it is harder to think about appropriate calibrations of the magnitude of these belief distortions.

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

$$d \log Y_t = \mu_y dt + \sigma_y \left(dW_t^1 - \varphi dW_t^2\right) + \varphi \sigma_y dW_t^2, \qquad Y_0 > 0,$$

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 symmetrically by both agents, and is thus cancelled out from the formula for the relative patience (this would not be the case if we considered 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

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

These formulas then enter the survival thresholds in Proposition OA.1. Recall that a sufficient condition for the existence of a nondegenerate long-run equilibrium is given by the pair of inequalities

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

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 and 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 $\theta \searrow 0$, 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.

OA.4.4 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 does not have a stationary distribution. What happens is that the conditional distribution of θ 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 $\theta_0 \in (0,1)$, the process θ visits every $\bar{\theta} \in (0,1)$, P-a.s. None of the agents vanishes with a strictly positive probability, yet a stationary distribution does not exist.

When preferences are not separable, this issue generally does not occur. The parameter space $(1-\gamma, 1-\rho)$ is divided into four regions, and one of the survival outcomes stated in the main survival proposition holds in each of these regions.

Figure 3 in the paper considers the case of an optimistic agent 1 and a pessimistic agent 2. The right panel of that figure plots the case of exactly mirror distortions, $u^1 = -u^2 = 0.025$. The division of the parameter space into the four survival regions occurs along the diagonal (the CRRA parameterization) and along a vertical line at risk aversion level $1-\gamma = |u^n|/\sigma_y$.

Figure OA.3 depicts perturbations of this belief parameterization (i.e., alternative perturbations of the belief parameters from the right panel of Figure 3). The thin dotted diagonal line represents the CRRA parameterizations.

The second and third panels of Figure OA.3 also reveal 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 (negative) sum of the two belief distortions is close to zero, 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

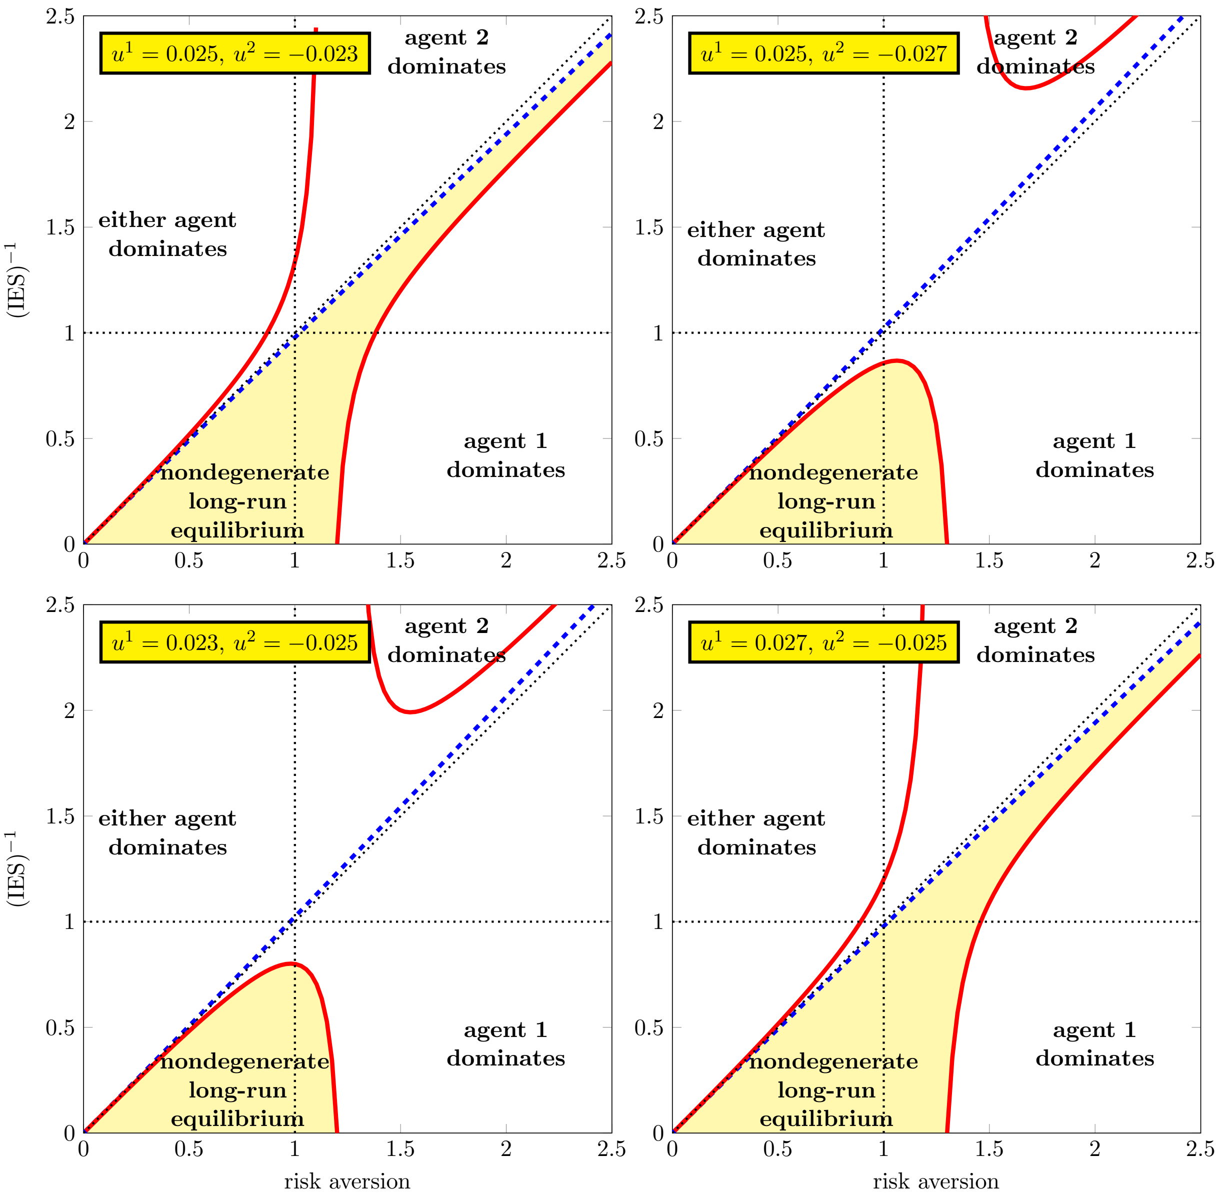

Figure OA.3: Survival regions for an optimistic agent 1 and a pessimistic agent 2 when their beliefs are close to symmetric (see legend of each plot). Volatility of aggregate endowment is $\sigma_y = 0.02$.

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.

OA.4.5 Exponential rate of convergence

When a stationary distribution for the Pareto share θ exists, convergence of the process to its stationary distribution occurs at an exponential rate, so that the process θ 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 OA.2 *Under the sufficient conditions for survival of both agents, the process θ 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

$$\sup_{\|\phi\|=1} \left\| E\left[\phi(\theta_t) \mid \theta_0 = \bar{\theta}_0\right] - \int_0^1 \phi(\theta)\, q(\theta)\, d\theta \right\| \doteq \rho_t \le B e^{-\delta t}.$$

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. ■

OA.5 Comparison to economies with only terminal consumption

In this paper, I analyze economies with intermediate consumption. Kogan, Ross, Wang, and Westerfield (2006) deal with a different framework with two agents endowed with CRRA preferences. In their economy, there is no intermediate consumption and the agents split and consume an aggregate dividend payoff at a terminal date T. The dividend evolves according to a geometric Brownian motion (OA.10) as in this paper, and agents can continuously retrade claims on the terminal payoff during the lifetime of the economy. The notion of survival in that framework is captured by analyzing the limit of the consumption share distribution in a sequence of economies as $T \nearrow \infty$.

Without intermediate consumption, the agent's intertemporal decision is reduced to the maximization of the (risk-adjusted) expected growth rate of the portfolio. In this respect, the framework is similar to a model in this paper under unitary IES when agents' consumption-wealth ratio is constant and equal to β, and intermediate consumption has no impact on the difference of the wealth growth rates.

The difference that prevents a direct comparison of the results lies in the valuation of wealth. In the absence of intermediate consumption, Kogan, Ross, Wang, and Westerfield (2006) use the price of a bond maturing at time T as numeraire and define the initial wealth in the economy with horizon T as the time-0 price of the terminal payoff. (In this paper, this quantity corresponds to the price of a single cash flow from the aggregate endowment paid out at time T, scaled by the price of a bond with corresponding maturity.) Then they consider two approaches to survival analysis.

In the ‘general equilibrium’ approach, they study the limiting properties of the terminal consumption allocation obtained as a solution of a sequence of planner’s problems as $T \nearrow \infty$. A critical assumption in this approach is the choice of the initial Pareto shares. These are chosen so that the initial wealth shares of the two agents are identical, which requires the initial Pareto share of the irrational agent to approach one as $T \nearrow \infty$. This mechanism reweighs the behavior of the tail and allows an optimistic agent to ‘survive’ in the sequence of planner’s problems. In economies with intermediate consumption, consumption at distant dates contributes only little to the wealth levels, and thus the reweighing of initial Pareto shares in order to achieve equal initial wealth levels would have no effect on the survival results. Under the ‘general equilibrium’ notion of survival in Kogan, Ross, Wang, and Westerfield (2006), optimistic agents can survive when risk aversion is larger than one but the survival regions differ from the results in this paper under unitary IES ($\rho = 0$).

Kogan, Ross, Wang, and Westerfield (2006) contrast their ‘general equilibrium’ to a simplified approach that is analogous to the boundary analysis in this paper and that they call the ‘partial equilibrium’ method. This method constructs a homogeneous economy injected with an infinitesimal agent with different beliefs under the assumption that she does not affect local price dynamics. It turns out that this method delivers exactly the same survival regions as those derived in this paper under unitary IES.

Propositions 5.1 and 5.2 from Section 3 of the paper show that in the model with intermediate consumption considered in this paper, the return on aggregate wealth and prices of individual finite-horizon cash flows from the aggregate endowment converge to their homogeneous economy counterparts and thus the ‘partial equilibrium’ approach is actually the correct method for this paper under general equilibrium. However, these results do not translate to the setup considered in Kogan, Ross, Wang, and Westerfield (2006). Although prices of individual cash flows from the aggregate endowment converge for every fixed $T \geq 0$, this convergence is not uniform on $T \in [0, \infty)$, which in general invalidates the result on converging returns and prices for the limit as $T \nearrow \infty$.

OA.6 Extensions

The analysis in this paper focuses on the case of exogenously specified time-invariant belief distortions. Agents are firm believers in their probability models, and do not use new data to update their beliefs. This can be interpreted as the strongest form of incorrect beliefs, and a bias against survival of agents whose beliefs are initially incorrect. However, the methodology can be applied to more complex belief distortions, including endogenously determined ones.

For instance, a natural question is to ask what happens when agents are allowed to learn. Learning can be incorporated into the current framework by introducing a law of motion that represents the Bayesian updating of the belief distortions u^n. These belief distortions become new state variables.

Blume and Easley (2006) provide a detailed analysis of the impact of Bayesian learning on survival under separable utility, and show that learning in general aids survival of agents who start with incorrect beliefs, by reducing their belief distortions. The message is much less clear in the nonseparable preference case. For instance, Figure 2 shows that the survival region of a pessimistic agent can shrink if her belief distortion diminishes. Whether the pessimist can learn quickly enough so that her beliefs converge to rational expectations at a rate that allows survival depends on the complexity of the learning problem, as shown by Blume and Easley (2006). The limiting distribution of θ as $t \nearrow \infty$ for the case of nonseparable preferences thus remains an open question.

Subjective beliefs can also arise from other decision-theoretical models. For instance, Bhandari (2015) uses the dynamics of Pareto weights to study a model where wealth dynamics interact with endogenous beliefs of agents concerned about model misspecification. Finally, formulas for survival regions can be extended by incorporating heterogeneity in preferences, as in Dumas, Uppal, and Wang (2000).

Here, I provide a more detailed outline of two specific 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.

OA.6.1 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 (OA.2) 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 (OA.2) 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

$$dp_t^n = \Delta\left(p_t^n\right)\left(d\log Y_t - \left(\mu_y^n\right)' p_t^n dt\right), \tag{OA.11}$$

where

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

is the regression 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 $\log Y_t$ to be $\left(\mu_y^n\right)' p_t^n$, and thus

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

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

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

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

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

is therefore the distorting martingale that replaces the original distorting martingale in the case of a learning agent. The agent acts as if there was a time-varying average distortion process $\bar{u}_t^n = \left(u^n\right)' p_t^n$. The optimization problem of the fictitious planner is extended by the filtering equation (OA.11) and the evolution of the modified discount factor becomes

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

Conjecturing a new Markov state $Z = \left(\bar{\lambda}', Y, \left(p^1\right)', \left(p^2\right)'\right)'$, 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.

OA.6.2 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

$$d\log Y_t = \mu_y dt + \sigma_y dW_t, \quad t \geq 0$$

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 (OA.4) 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

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

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

$$d\log Y_t = \mu_y dt + \sigma_y \left(u_t^n dt + dW_t^n\right),$$

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.

Bhandari (2015) uses the robust preference structure of Hansen and Sargent to study wealth dynamics in a two-agent economy in a discrete-time environment. Guerdjikova and Sciubba (2015) consider a heterogeneous-preference economy with one agent endowed with separable preferences and another endowed with smooth ambiguity averse preferences of Klibanoff, Marinacci, and Mukerji (2005, 2009). In both cases, endogenous belief distortions emerge as part of the solution of the problem.

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 model misspecification concerns emerging from robust preferences (or other forms of ambugity aversion) ex post as a specific endogenously generated belief distortion. Reverting the argument, the framework introduced in this paper can be used 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.

OA.7 Proofs omitted from Appendix A of the main text

Proof of Lemma A.3. Equation (8) can be written as

$$\begin{aligned} J\left(\bar{\lambda}_t, Y_t\right) &= \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^{\gamma} \sup_{a \in \mathcal{A}} E_t \left[\theta_t \int_t^{\infty} \frac{\bar{\lambda}_s^1}{\bar{\lambda}_t^1} \left(\frac{Y_s}{Y_t}\right)^{\gamma} F\left(\zeta_s^1, \nu_s^1\right) ds \right. \\ &\qquad\qquad \left. + \left(1-\theta_t\right) \int_t^{\infty} \frac{\bar{\lambda}_s^2}{\bar{\lambda}_t^2} \left(\frac{Y_s}{Y_t}\right)^{\gamma} F\left(\zeta_s^2, \nu_s^2\right) ds \right] \\ &\doteq \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^{\gamma} \widetilde{J}\left(\theta_t\right) \end{aligned}$$

because the ratios $\bar{\lambda}_s^n / \bar{\lambda}_t^n$ and Y_s/Y_t do not depend on $\bar{\lambda}_t^n$ and Y_t. Here, $a = \left(\zeta^1, \zeta^2, \nu^1, \nu^2\right) \in \mathcal{A}$ is a set of controls equivalent to Definition A.2 because $C^n = \zeta^n Y$. Further, since the individual value functions are increasing in consumption, we have

$$J\left(\bar{\lambda}_t, Y_t\right) = \sup_{C^1 + C^2 = Y} \bar{\lambda}_t^1 V_t^1\left(C^1\right) + \bar{\lambda}_t^2 V_t^2\left(C^2\right) \leq \bar{\lambda}_t^1 V_t^1\left(Y\right) + \bar{\lambda}_t^2 V_t^2\left(Y\right).$$

The value functions $V_t^n\left(Y\right)$ have a closed form solution for the iid growth process Y, given by $V_t^n\left(Y\right) = Y_t^{\gamma} \bar{V}^n$ where

$$\bar{V}^n = \frac{1}{\gamma}\left(\beta^{-1}\left[\beta - \rho\left(\mu_y + u^n \sigma_y + \frac{1}{2}\gamma \sigma_y^2\right)\right]\right)^{-\frac{\gamma}{\rho}} \tag{OA.12}$$

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

with the associated optimal discount rate

$$\bar{\nu}^n = \frac{\beta}{\rho}\left(\gamma + (\rho-\gamma)\left(\gamma\bar{V}^n\right)^{-\frac{\rho}{\gamma}}\right) = \beta + (\gamma-\rho)\left(\mu_y + u^n\sigma_y + \frac{1}{2}\gamma\sigma_y^2\right). \tag{OA.13}$$

$V_t^n(Y)$ and $\bar{\nu}^n$ are the value function and discount rate from a homogeneous economy populated only by agent n. These objects are well-defined when the first restriction in Assumption A.1 holds and satisfy the same homogeneity properties as the planner's value function. Therefore, $\widetilde{J}(\theta) \le \theta\bar{V}^1 + (1-\theta)\bar{V}^2$.

Finally, consider a suboptimal policy consisting of fixing, given an initial θ_t, the consumption shares $\bar{\zeta}^n$ for the two agents for the whole future. Since individual consumption processes now exhibit iid growth, the optimal choice of the discount rate will satisfy $\nu_t^n = \bar{\nu}^n$. We obtain

$$\begin{aligned} J\left(\bar{\lambda}_t, Y_t\right) &\ge \sup_{\bar{\zeta}^1+\bar{\zeta}^2=1}\left[\bar{\lambda}_t^1 V_t^1\left(\bar{\zeta}^1 Y_t\right) + \bar{\lambda}_t^2 V_t^2\left(\bar{\zeta}^2 Y_t\right)\right] = \\ &= \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^\gamma \sup_{\bar{\zeta}^1+\bar{\zeta}^2=1}\left[\theta_t\left(\bar{\zeta}^1\right)^\gamma \bar{V}^1 + (1-\theta_t)\left(\bar{\zeta}^2\right)^\gamma \bar{V}^2\right] \end{aligned}$$

and thus

$$\widetilde{J}(\theta) \ge \sup_{\bar{\zeta}^1+\bar{\zeta}^2=1} \theta\left(\bar{\zeta}^1\right)^\gamma \bar{V}^1 + (1-\theta)\left(\bar{\zeta}^2\right)^\gamma \bar{V}^2.$$

The first-order condition with respect to $\bar{\zeta}^1$ yields

$$\bar{\zeta}^1(\theta_t) = \frac{\left[\theta_t\gamma\bar{V}^1\right]^{\frac{1}{1-\gamma}}}{\left[\theta_t\gamma\bar{V}^1\right]^{\frac{1}{1-\gamma}} + \left[(1-\theta_t)\gamma\bar{V}^2\right]^{\frac{1}{1-\gamma}}} \tag{OA.14}$$

and $\bar{\zeta}^2(\theta_t) = 1 - \bar{\zeta}^1(\theta_t)$. Substituting this result back, we have

$$\begin{aligned} J\left(\bar{\lambda}_t, Y_t\right) &\ge \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^\gamma \frac{1}{\gamma}\left[\left[\theta_t\gamma\bar{V}^1\right]^{\frac{1}{1-\gamma}} + \left[(1-\theta_t)\gamma\bar{V}^2\right]^{\frac{1}{1-\gamma}}\right]^{1-\gamma} \\ &= \left(\bar{\lambda}_t^1 + \bar{\lambda}_t^2\right) Y_t^\gamma \bar{J}(\theta_t). \end{aligned} \tag{OA.15}$$

which establishes the lower bound on $\widetilde{J}(\theta_t)$. ∎

Proof of Lemma A.5. Introspection of the function $h^0(\nu,\theta)$ in (19) reveals that under Assumption A.4, this function is bounded away from zero, and thus there exists $M > 0$ such that $\left|h^0(\nu,\theta)\right| > M$. We have

$$\begin{aligned} J\left(\bar{\lambda}_t, Y_t\right) &= \sup_{a\in\mathcal{A}} E_t\left[\int_t^\infty \left[\bar{\lambda}_s^1 F\left(C_s^1,\nu_s^1\right) + \bar{\lambda}_s^2 F\left(C_s^2,\nu_s^2\right)\right] ds\right] = \\ &= Y_t^\gamma \sup_{(\nu^1,\nu^2)} E_t\left[\int_t^\infty \left(\bar{\lambda}_s^1 + \bar{\lambda}_s^2\right)\left(\frac{Y_s}{Y_t}\right)^\gamma h^0(\nu_s,\theta_s)\, ds\right]. \end{aligned}$$

Consider an arbitrary pair of processes $\left(\nu^1,\nu^2\right)$ and the associated optimal consumption shares ζ^n such that $a = \left(\zeta^1,\zeta^2,\nu^1,\nu^2\right)$ is admissible. Then

$$\begin{aligned} +\infty &> E_t\left[\int_t^\infty \left[\bar{\lambda}_s^1\left|F\left(C_s^1,\nu_s^1\right)\right| + \bar{\lambda}_s^2\left|F\left(C_s^2,\nu_s^2\right)\right|\right] ds\right] \ge \\ &\ge Y_t^\gamma E_t\left[\int_t^\infty \left(\bar{\lambda}_s^1 + \bar{\lambda}_s^2\right)\left(\frac{Y_s}{Y_t}\right)^\gamma \left|h^0(\nu_s,\theta_s)\right| ds\right] \ge M Y_t^\gamma E_t\left[\int_t^\infty \left(\bar{\lambda}_s^1 + \bar{\lambda}_s^2\right)\left(\frac{Y_s}{Y_t}\right)^\gamma ds\right] \end{aligned}$$

which proves (20). The limit in (21) is a direct consequence. ∎

Proof of Lemma A.6. Consider the case $\bar{\lambda}_t^1 \searrow 0$. Given optimal consumption streams $C^n\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right)$, we have

$$J\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right) = \bar{\lambda}_t^1 V_t^1\left(C^1\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right)\right) + \bar{\lambda}_t^2 V_t^2\left(C^2\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right)\right). \tag{OA.16}$$

Since $V_t^1\left(C^1\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right)\right)$ is bounded from above as a function of $\bar{\lambda}_t$ by $V_t^1(Y)$, it follows that

$$\lim_{\bar{\lambda}_t^1 \searrow 0} \bar{\lambda}_t^1 V_t^1\left(C^1\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right)\right) = \underline{V}^1 \leq 0$$

and thus

$$J\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right) \leq \lim_{\bar{\lambda}_t^1 \searrow 0} \bar{\lambda}_t^2 V_t^2\left(C^2\left(\bar{\lambda}_t^1, \bar{\lambda}_t^2, Y_t\right)\right) \leq \bar{\lambda}_t^2 V_t^2(Y).$$

Conversely, assume suboptimal policies $\zeta_u^n = \bar{\zeta}^n(\theta_t)$ for $u \geq t$ where $\bar{\zeta}^n(\theta_t)$ are given by (OA.14). Taking the limit in (OA.15) and noticing that $\bar{\lambda}_t^1 \searrow 0$ for a given $\bar{\lambda}_t^2 > 0$ implies $\theta_t \searrow 0$ yields

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

Combining the two inequalities yields (22). The limit in (23) is a direct consequence. ■

Remark OA.1 *The maximization over $\left(\nu^1, \nu^2\right)$ in the HJB equation (25) can be solved separately. Under the optimal discount rate process ν^n for agent n,*

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

The function f is the aggregator in the stochastic differential utility representation of recursive preferences postulated by Duffie and Epstein (1992b). Section OA.2 gives more detail on this relationship. Optimal consumption shares ζ^n are given by the first-order conditions in the consumption allocation

$$\zeta^1 = \frac{\left(\bar{\lambda}^1\right)^{\frac{1}{1-\rho}}\left(\gamma J_{\bar{\lambda}^1}\right)^{\frac{1-\rho/\gamma}{1-\rho}}}{\sum_{k=1}^2 \left(\bar{\lambda}^k\right)^{\frac{1}{1-\rho}}\left(\gamma J_{\bar{\lambda}^k}\right)^{\frac{1-\rho/\gamma}{1-\rho}}} = \frac{\theta^{\frac{1}{1-\rho}}\left(\gamma \widetilde{J}^1(\theta)\right)^{\frac{1-\rho/\gamma}{1-\rho}}}{\theta^{\frac{1}{1-\rho}}\left(\gamma \widetilde{J}^1(\theta)\right)^{\frac{1-\rho/\gamma}{1-\rho}} + (1-\theta)^{\frac{1}{1-\rho}}\left(\gamma \widetilde{J}^2(\theta)\right)^{\frac{1-\rho/\gamma}{1-\rho}}}, \tag{OA.17}$$

and $\zeta^2 = 1 - \zeta^1$, where $J_{\bar{\lambda}^n} = Y^\gamma \widetilde{J}^n(\theta)$ are the individual agents' continuation values under the optimal consumption allocation, with $\widetilde{J}^n(\theta)$ defined as

$$\begin{aligned}\widetilde{J}^1(\theta) &= \widetilde{J}(\theta) + (1-\theta)\widetilde{J}_\theta(\theta) \\ \widetilde{J}^2(\theta) &= \widetilde{J}(\theta) - \theta \widetilde{J}_\theta(\theta).\end{aligned} \tag{OA.18}$$

These are obtained from the envelope condition on the planner's value function (8).

Proof of Lemma A.8. The proof is a modification of the shooting algorithm argument from Strulovici and Szydlowski (2014). We first show the claim of the lemma for a fixed control $\nu(\theta) = \left(\nu^1(\theta), \nu^2(\theta)\right)$ that satisfies Assumption A.4. Then we extend the argument to the optimal control in the boundary value problem (28).

Consider an initial value problem (for a given $\nu(\theta)$) given by differential equation (28) together with initial conditions $\widetilde{J}^\varepsilon(\varepsilon) = \bar{J}(\varepsilon)$ and $\widetilde{J}_\theta^\varepsilon(\varepsilon) = y$. Since the functions $h^j(\nu(\theta), \theta)$ are bounded and satisfy Lipschitz continuity, it is well known (see, e.g., the references in Strulovici and Szydlowski (2014), Appendix B) that the initial value problem has a unique, twice continuously differentiable solution that is uniformly continuous in y. The goal is to show that we can find a unique value of y such that $\widetilde{J}^{\varepsilon}(1-\varepsilon) = \bar{J}(1-\varepsilon)$, so that the boundary value problem has a unique solution.

Define $K(\theta) = \widetilde{J}^{\varepsilon}_{\theta}(\theta)$ and $k^{j}(\theta) = -h^{j}(\nu(\theta), \theta)/h^{3}(\theta)$ for $j = 0, 1, 2$. Then (28) can be integrated to

$$
\begin{aligned}
K(s) &= y + \int_{\varepsilon}^{s} \left[k^{0}(r) + k^{1}(r)\widetilde{J}^{\varepsilon}(r) + k^{2}(r)K(r)\right] dr \\
\widetilde{J}^{\varepsilon}(\theta) &= \bar{J}(\varepsilon) + \int_{\varepsilon}^{\theta} K(s)\, ds.
\end{aligned} \tag{OA.19}
$$

We are interested in the sensitivity of $\widetilde{J}^{\varepsilon}(1-\varepsilon)$ to changes in the initial condition $\widetilde{J}^{\varepsilon}_{\theta}(\varepsilon) = y$. We have

$$
\begin{aligned}
\frac{d}{dy}K(s) &= 1 + \int_{\varepsilon}^{s} \left[k^{1}(r)\frac{d}{dy}\widetilde{J}^{\varepsilon}(r) + k^{2}(r)\frac{d}{dy}K(r)\right] dr = \\
&= 1 + \int_{\varepsilon}^{s} \left[k^{1}(r)\int_{\varepsilon}^{r}\frac{d}{dy}K(p)\,dp + k^{2}(r)\frac{d}{dy}K(r)\right] dr \\
&= 1 + \int_{\varepsilon}^{s} \left[\left(\int_{r}^{s} k^{1}(r')\,dr' + k^{2}(r)\right)\frac{d}{dy}K(r)\right] dr
\end{aligned}
$$

This is an integral form of a differential equation for $\frac{d}{dy}K(s)$ in s with $\frac{d}{dy}K(0) = 1$. Given the term $\int_{r}^{s} k^{1}(r')\,dr' + k^{2}(r)$ is bounded, take an $M > 0$ such that

$$
\left|\int_{r}^{s} k^{1}(r')\,dr' + k^{2}(r)\right| < M.
$$

Then

$$
e^{-M(s-\varepsilon)} \leq \frac{d}{dy}K(s) \leq e^{M(s-\varepsilon)}
$$

and therefore, using (OA.19),

$$
\frac{1}{M}\left[1 - e^{-M(1-2\varepsilon)}\right] \leq \frac{d}{dy}\widetilde{J}^{\varepsilon}(1-\varepsilon) \equiv \int_{\varepsilon}^{1-\varepsilon}\frac{d}{dy}K(s)\,ds \leq \frac{1}{M}\left[e^{M(1-2\varepsilon)} - 1\right].
$$

The sensitivity of the terminal value $\widetilde{J}^{\varepsilon}(1-\varepsilon)$ with respect to changes in the initial condition is therefore always positive, bounded and bounded away from zero. Moreover, the existence of the continuously differentiable solution for the initial value problem extends beyond $\theta = 1-\varepsilon$. Therefore, for an arbitrary choice of the initial slope y, the terminal value $\widetilde{J}^{\varepsilon}(1-\varepsilon)$ is finite. The lower bound on $\frac{d}{dy}\widetilde{J}^{\varepsilon}(1-\varepsilon)$ then implies that we can always sufficiently vary y to reach an arbitrary terminal value $\widetilde{J}^{\varepsilon}(1-\varepsilon)$. The fact that $\frac{d}{dy}\widetilde{J}^{\varepsilon}(1-\varepsilon)$ is always positive implies that the choice of y such that the terminal value yields the boundary condition $\widetilde{J}^{\varepsilon}(1-\varepsilon) = \bar{J}(1-\varepsilon)$ of the boundary value problem (28) is unique.

The extension of the proof to the optimal control $\nu(\theta)$ is a consequence of Berge's Maximum Theorem. The unique maximizers are given by

$$
\nu^{n}(\theta) = \frac{\beta}{\rho}\left(\gamma + (\rho-\gamma)\left(\frac{\zeta^{n}(\theta)^{\gamma}}{\gamma\widetilde{J}^{n}(\theta)}\right)^{\rho/\gamma}\right) \tag{OA.20}
$$

with $\widetilde{J}^{n}(\theta)$ defined as in (OA.18) except for $\widetilde{J}^{\varepsilon}(\theta)$ in place of $\widetilde{J}(\theta)$. These functions satisfy Assumption A.4 on every interval $[\varepsilon, 1-\varepsilon]$, $\varepsilon \in \left(0, \frac{1}{2}\right)$. The limits of these formulas at the boundaries as $\varepsilon \searrow 0$ are computed in the proof of Proposition OA.1 in Appendix B. This concludes the proof. ■

Proof of Lemma A.11. Consider time $\tau \geq t$. Then

$$J\left(\bar{\lambda}_\tau, Y_\tau\right) = J\left(\bar{\lambda}_t, Y_t\right) + \int_t^\tau \mu_{J,s} ds + \int_t^\tau \sigma_{J,s} dW_s$$

where

$$\mu_{J,s} = \left(\bar{\lambda}_s^1 + \bar{\lambda}_s^2\right) Y_s^\gamma \left\{ h^1\left(\nu_s, \theta_s\right) \widetilde{J}\left(\theta_s\right) + h^2\left(\nu_s, \theta_s\right) \widetilde{J}_\theta\left(\theta_s\right) + h^3\left(\theta_s\right) \widetilde{J}_{\theta\theta}\left(\theta_s\right) \right\}$$
$$\sigma_{J,s} = \left(\bar{\lambda}_s^1 + \bar{\lambda}_s^2\right) Y_s^\gamma \left\{ \left[\theta_s u^1 + \left(1-\theta_s\right) u^2 + \gamma \sigma_y\right] \widetilde{J}\left(\theta_s\right) + \theta_s\left(1-\theta_s\right)\left(u^1 - u^2\right) \widetilde{J}_\theta\left(\theta_s\right) \right\}.$$

It follows from the discussion in Section A.4 that the terms $\widetilde{J}(\theta)$ and $\theta(1-\theta)\widetilde{J}_\theta(\theta) = \widehat{J}_\vartheta(\vartheta(\theta))$ are bounded, and thus $\sigma_{J,s}$ is square integrable over $[t,\tau]$. As a consequence, the stochastic integral $\int_t^\tau \sigma_{J,s} dW_s$ is a martingale as a function of τ, and we have

$$E_t\left[J\left(\bar{\lambda}_\tau, Y_\tau\right)\right] = J\left(\bar{\lambda}_t, Y_t\right) + E_t\left[\int_t^\tau \mu_{J,s} ds\right].$$

The limiting version of the HJB equation (28) for $\varepsilon \searrow 0$ implies that for an arbitrary admissible control $\left(\zeta^1, \zeta^2, \nu^1, \nu^2\right)$ with optimal choice of the consumption shares $\zeta^n = \zeta^n(\nu)$ conditional on ν,

$$E_t\left[J\left(\bar{\lambda}_\tau, Y_\tau\right)\right] \leq J\left(\bar{\lambda}_t, Y_t\right) - E_t\left[\int_t^\tau \left(\bar{\lambda}_s^1 + \bar{\lambda}_s^2\right) Y_s^\gamma h^0\left(\nu_s, \theta_s\right) ds\right]$$

with equality for the optimal control $\nu = \left(\nu^1, \nu^2\right)$ given in (OA.20). Reorganizing and taking the limit $\tau \to \infty$, we obtain

$$E_t\left[\int_t^\infty \left(\bar{\lambda}_s^1 + \bar{\lambda}_s^2\right) Y_s^\gamma h^0\left(\nu_s, \theta_s\right) ds\right] \leq J\left(\bar{\lambda}_t, Y_t\right) \tag{OA.21}$$

where we utilized Lemma A.5 to show that

$$\lim_{\tau\to\infty} E_t\left[J\left(\bar{\lambda}_\tau, Y_\tau\right)\right] = \lim_{\tau\to\infty} E_t\left[\left(\bar{\lambda}_\tau^1 + \bar{\lambda}_\tau^2\right) Y_\tau^\gamma \widetilde{J}\left(\theta_\tau\right)\right] = 0$$

because $\widetilde{J}(\theta)$ is a bounded function. The left-hand side of (OA.21) evaluated for the maximizers (OA.20) is the value function, and since these maximizers are admissible, the inequality holds with equality for the value function. ■

OA.8 Proofs omitted from Appendix B of the main text

Optimal choice of ν^n in (9) implies that

$$\sup_{\nu \in R} F(C,\nu) - \nu V \doteq f(C,V) = \frac{\beta}{\rho}\left[C^\rho (\gamma V)^{1-\frac{\rho}{\gamma}} - \gamma V\right].$$

Substituting this expression into (9) for $F\left(\zeta, \widetilde{J}^1\right)$ and $F\left(1-\zeta, \widetilde{J}^2\right)$ leads to the ODE

$$\begin{aligned}0 = & \theta\frac{\beta}{\rho}\left(\zeta^1\right)^\rho\left(\gamma\widetilde{J}^1(\theta)\right)^{1-\frac{\rho}{\gamma}} + (1-\theta)\frac{\beta}{\rho}\left(\zeta^2\right)^\rho\left(\gamma\widetilde{J}^2(\theta)\right)^{1-\frac{\rho}{\gamma}} + \\ & + \gamma\left[-\frac{\beta}{\rho} + \left(\theta u^1 + (1-\theta)u^2\right)\sigma_y + \mu_y + \frac{1}{2}\gamma\sigma_y^2\right]\widetilde{J}(\theta) \\ & + \theta(1-\theta)\left(u^1-u^2\right)\gamma\sigma_y\widetilde{J}_\theta(\theta) + \frac{1}{2}\theta^2(1-\theta)^2\left(u^1-u^2\right)^2\widetilde{J}_{\theta\theta}(\theta)\end{aligned} \tag{OA.22}$$

where ζ^n are given by (OA.17). Appendix A in the main text proves the existence of a twice-continuously differentiable solution to this equation. The results that follow also utilize the third derivative of $\widetilde{J}$, which can be obtained by differentiating (OA.22).

Before proceeding with the proof of Proposition 5.1, we prove two lemmas that characterize the boundary behavior of $\widetilde{J}(\theta)$ and consumption shares of the two agents.

Lemma OA.1 *The solution of the planner's problem satisfies*

$$\lim_{\theta\searrow 0}\theta\widetilde{J}_\theta(\theta) = \lim_{\theta\searrow 0}(\theta)^2\widetilde{J}_{\theta\theta}(\theta) = \lim_{\theta\searrow 0}(\theta)^3\widetilde{J}_{\theta\theta\theta}(\theta) = 0.$$

Proof. Lemma A.6 implies that the planner's objective function can be continuously extended at $\theta = 0$ by the continuation value for agent 2 from a homogeneous economy. Expression (OA.16) scaled by $\left(\alpha^1+\alpha^2\right)\gamma^{-1}Y^\gamma$ leads to an equation in scaled continuation values

$$\widetilde{J}(\theta) = \theta\widetilde{J}^1(\theta) + (1-\theta)\widetilde{J}^2(\theta)$$

and the proof of Lemma A.6 yields

$$\lim_{\theta\searrow 0}\widetilde{J}(\theta) = \lim_{\theta\searrow 0}\widetilde{J}^2(\theta) = \bar{V}^2,$$

where $\bar{V}^2$ is defined in (OA.12). Since $\widetilde{J}^2(\theta) = \widetilde{J}(\theta) - \theta\widetilde{J}_\theta(\theta)$, then

$$\lim_{\theta\searrow 0}\theta\widetilde{J}_\theta(\theta) = 0. \tag{OA.23}$$

Further, consider the behavior of individual terms in ODE (OA.22) as $\theta\searrow 0$. Using expression (OA.17), the first term is proportional to

$$\begin{aligned}\theta(\zeta(\theta))^\rho\left(\widetilde{J}^1(\theta)\right)^{1-\frac{\rho}{\gamma}} &= (\theta)^{\frac{1}{1-\rho}}\left(\widetilde{J}^1(\theta)\right)^{\frac{1-\rho/\gamma}{1-\rho}}[K(\theta)]^{-\rho} = \\ &= \zeta(\theta)[K(\theta)]^{1-\rho},\end{aligned}$$

where $K(\theta)$ is the denominator in the formula for the consumption share (OA.17), and $\lim_{\theta\searrow 0}K(\theta) = \left(\bar{V}^2\right)^{\frac{1-\rho/\gamma}{1-\rho}} < \infty$. Since $\lim_{\theta\searrow 0}\zeta(\theta) = 0$, the first term in (OA.22) vanishes as $\theta\searrow 0$. The sum of the second and third term converges to

$$\frac{\beta}{\rho}\left(\bar{V}^2\right)^{1-\frac{\rho}{\gamma}} + \left(-\frac{\beta}{\rho} + \mu_y + u^2\sigma_y + \frac{1}{2}(\gamma-1)\sigma_y^2\right)\bar{V}^2$$

and formula (OA.12) implies that this expression is zero. Since the fourth term in (OA.22) also converges to zero due to (OA.23), the last term in (OA.22) must also converge to zero, or

$$\lim_{\theta\searrow 0}(\theta)^2\widetilde{J}_{\theta\theta}(\theta) = 0. \tag{OA.24}$$

Finally, differentiate the PDE (OA.22) with respect to θ and multiply the equation by θ. Using comparisons with results (OA.23)–(OA.24), the assumption that functions $\zeta^n(\theta) / \left(\gamma\widetilde{J}^n(\theta)\right)^{1/\gamma}$ are bounded and bounded away from zero and $\lim_{\theta\searrow 0}\zeta^1(\theta)=0$, we determine that all terms in the differentiated equation containing derivatives of $\widetilde{J}(\theta)$ up to second order vanish as $\theta\searrow 0$. The single remaining term that contains a third derivative of $\widetilde{J}(\theta)$ is multiplied by $(\theta)^3$ and must necessarily converge to zero as well, and thus

$$\lim_{\theta\searrow 0}(\theta)^3\,\widetilde{J}_{\theta\theta\theta}(\theta)=0.$$

■

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

$$\frac{d\widetilde{J}^n(\theta_t)}{\widetilde{J}^n(\theta_t)} \doteq \mu_{\widetilde{J}^n}(\theta_t)\,dt+\sigma_{\widetilde{J}^n}(\theta_t)\,dW_t \tag{OA.25}$$

$$\frac{d\zeta^n(\theta_t)}{\zeta^n(\theta_t)} \doteq \mu_{\zeta^n}(\theta_t)\,dt+\sigma_{\zeta^n}(\theta_t)\,dW_t. \tag{OA.26}$$

where the drift and volatility coefficients are functions of θ. The following lemma characterizes the boundary behavior of these coefficients for agent 2 as $\theta\searrow 0$.

Lemma OA.2 *The coefficients in equations (OA.25)–(OA.26) for agent 2 satisfy*

$$\lim_{\theta\searrow 0}\mu_{\widetilde{J}^2}(\theta)=\lim_{\theta\searrow 0}\sigma_{\widetilde{J}^2}(\theta)=\lim_{\theta^1\searrow 0}\mu_{\zeta^2}(\theta)=\lim_{\theta\searrow 0}\sigma_{\zeta^2}(\theta)=0.$$

Proof. The result follows from an application of Itô's lemma to $\widetilde{J}^2$ and ζ^2. Utilizing formulas (OA.17) and (OA.18), the coefficients contain expressions for the value function $\widetilde{J}(\theta)$ and its partial derivatives up to the third order, and all expressions can be shown to converge to zero using Lemma OA.1. Itô's lemma implies

$$\begin{aligned} d\widetilde{J}^2(\theta_t) &= d\left[\widetilde{J}(\theta_t)-\theta_t\widetilde{J}_\theta(\theta_t)\right]= \\ &= -(\theta_t)^2\,\widetilde{J}_{\theta\theta}(\theta_t)\,\frac{d\theta_t}{\theta_t}-\frac{1}{2}\left[(\theta_t)^2\,\widetilde{J}_{\theta\theta}(\theta_t)+(\theta_t)^3\,\widetilde{J}_{\theta\theta\theta}(\theta_t)\right]\left(\frac{d\theta_t}{\theta_t}\right)^2. \end{aligned}$$

Equation (26) implies that the drift and volatility coefficients of $d\theta_t/\theta_t$ are bounded by Assumption A.4. Applying results from Lemma OA.1 then proves the claim about the drift and volatility coefficients of $\widetilde{J}^2(\theta)$ ($\widetilde{J}^2$ itself converges to a nonzero limit so the scaling is innocuous). Further notice that

$$\begin{aligned} d\widetilde{J}^1(\theta_t) = d\left[\widetilde{J}(\theta_t)+(1-\theta_t)\,\widetilde{J}_\theta(\theta_t)\right] &= -(\theta_t)^2\,\widetilde{J}_{\theta\theta}(\theta_t)\,\frac{d\theta_t}{\theta_t}+ \\ &+\frac{1}{2}\left[(\theta_t)^2\,\widetilde{J}_{\theta\theta}(\theta_t)+(1-\theta_t)(\theta_t)^2\,\widetilde{J}_{\theta\theta\theta}(\theta_t)\right]\left(\frac{d\theta_t}{\theta_t}\right)^2 \end{aligned} \tag{OA.27}$$

and that

$$\frac{\zeta^1(\theta)}{\left(\gamma\widetilde{J}^1(\theta)\right)^{\frac{1}{\gamma}}}=(\theta)^{\frac{1}{1-\rho}}\left(\widetilde{J}^1(\theta)\right)^{\frac{1-1/\gamma}{1-\rho}}K(\theta)^{-1} \tag{OA.28}$$

is bounded and bounded away from zero by assumption. Denote the numerators of ζ^n in (OA.17) as

$$Z^1(\theta)=\theta^{\frac{1}{1-\rho}}\left(\gamma\widetilde{J}^1(\theta)\right)^{\frac{1-\rho/\gamma}{1-\rho}} \qquad Z^2(\theta)=(1-\theta)^{\frac{1}{1-\rho}}\left(\gamma\widetilde{J}^2(\theta)\right)^{\frac{1-\rho/\gamma}{1-\rho}}.$$

Then $\zeta^2 = Z^2 / \left(Z^1 + Z^2\right)$ and, omitting arguments,

$$\begin{aligned}
dZ^1 =& \frac{1}{1-\rho} Z^1 \frac{d\theta}{\theta} + \frac{1-\frac{\rho}{\gamma}}{1-\rho} Z^1 \frac{d\widetilde{J}^1}{\widetilde{J}^1} + \frac{1}{2} \frac{\rho}{(1-\rho)^2} Z^1 \left(\frac{d\theta}{\theta}\right)^2 + \\
&+ \frac{1}{2} \frac{\left(\rho - \frac{\rho}{\gamma}\right)\left(1-\frac{\rho}{\gamma}\right)}{(1-\rho)^2} Z^1 \left(\frac{d\widetilde{J}^1}{\widetilde{J}^1}\right)^2 + \frac{1-\frac{\rho}{\gamma}}{(1-\rho)^2} Z^1 \frac{d\theta}{\theta} \frac{d\widetilde{J}^1}{\widetilde{J}^1} \\
dZ^2 =& -\frac{1}{1-\rho} Z^2 \frac{\theta}{1-\theta} \frac{d\theta}{\theta} + \frac{1-\frac{\rho}{\gamma}}{1-\rho} Z^2 \frac{d\widetilde{J}^2}{\widetilde{J}^2} + \frac{1}{2} \frac{\rho}{(1-\rho)^2} Z^2 \left(\frac{\theta}{1-\theta}\right)^2 \left(\frac{d\theta}{\theta}\right)^2 + \\
&+ \frac{1}{2} \frac{\left(\rho - \frac{\rho}{\gamma}\right)\left(1-\frac{\rho}{\gamma}\right)}{(1-\rho)^2} Z^2 \left(\frac{d\widetilde{J}^2}{\widetilde{J}^2}\right)^2 - \frac{1-\frac{\rho}{\gamma}}{(1-\rho)^2} Z^2 \frac{\theta}{1-\theta} \frac{d\theta}{\theta} \frac{d\widetilde{J}^2}{\widetilde{J}^2}.
\end{aligned}$$

Since the drift and volatility coefficients of $d\widetilde{J}^2/\widetilde{J}^2$ vanish as $\theta \searrow 0$, and $\lim_{\theta \searrow 0} Z^2(\theta) = \left(\gamma \bar{V}^2\right)^{\frac{1-\rho/\gamma}{1-\rho}}$, the drift and volatility coefficients in the equation for dZ^2 vanish. In the equation for dZ^1, it remains to determine the behavior of terms containing $d\widetilde{J}^1$ (the remaining contributions to drift and volatility terms converge to zero due to $\lim_{\theta \searrow 0} Z^1(\theta) = 0$):

$$\frac{Z^1}{\widetilde{J}^1} = \theta \left[(\theta)^{\frac{1}{1-\rho}} \left(\gamma \widetilde{J}^1\right)^{\frac{1-1/\gamma}{1-\rho}} \right]^\rho,$$

where the term in brackets is bounded and bounded away from zero by utilizing (OA.28). Using the first θ to multiply the coefficients in $d\widetilde{J}^1$ in formula (OA.27), we conclude that the coefficients in $Z^1 d\widetilde{J}^1/\widetilde{J}^1$ vanish as $\theta \searrow 0$. Finally, the drift term arising from $\left(d\widetilde{J}^1\right)^2$ vanishes, and

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

Here, the last term has a bounded drift, the second last term is bounded, and the first term converges to zero as $\theta \searrow 0$, which can be shown using l'Hôpital's rule (the numerator converges to zero and the denominator to zero or $-\infty$, depending on the sign of γ):

$$\lim_{\theta \searrow 0} \frac{(\theta)^5 \left(\widetilde{J}_{\theta\theta}\right)^2}{\widetilde{J} + (1-\theta)\widetilde{J}_\theta} = \lim_{\theta \searrow 0} \frac{5(\theta)^4 \widetilde{J}_{\theta\theta} + 2(\theta)^5 \widetilde{J}_{\theta\theta\theta}}{1-\theta} = 0.$$

Thus all terms in the drift and volatility coefficients of dZ^1 vanish. Applying Itô's lemma to ζ^2 yields

$$\begin{aligned}
d\zeta^2 =& \frac{1}{Z^1 + Z^2} dZ^2 - \frac{Z^2}{\left(Z^1 + Z^2\right)^2} \left(dZ^1 + dZ^2\right) + \\
&+ \frac{Z^2}{\left(Z^1 + Z^2\right)^3} \left(dZ^1 + dZ^2\right)^2 - \frac{1}{\left(Z^1 + Z^2\right)^2} dZ^2 \left(dZ^1 + dZ^2\right)
\end{aligned}$$

and the results on the behavior of dZ^1 and dZ^2 as $\theta \searrow 0$ lead to the desired conclusion about the convergence of drift and volatility coefficients of $d\zeta^2$. ■

Proof of Proposition 5.1. I start by assuming that $\xi^n(\theta)$ in (34) are functions that are bounded and bounded away from zero. This implies that the discount rate functions $\nu^n(\theta)$ are bounded and that the drift and volatility coefficients in the stochastic differential equation for θ, (26), are bounded as well. The assumption will ultimately be verified by a direct calculation of the limits of $\xi^n(\theta)$ as $\theta \to \{0,1\}$. Without loss of generality, it is sufficient to focus on the case $\theta \searrow 0$.

Lemmas OA.1 and OA.2 characterize the convergence as $\theta \searrow 0$ of the local behavior of the stochastic discount factor $S_t^2(\theta)$ in (33) to $S_t^2(0)$, which is the limiting stochastic discount factor corresponding to the one prevailing in a homogeneous economy populated only by agent 2. Convergence of the risk-free interest rate follows from the direct calculation of

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

Similarly, convergence of the aggregate wealth-consumption ratio follows from

$$\xi(\theta) = \xi^1(\theta)\zeta^1(\theta) + \xi^2(\theta)\zeta^2(\theta).$$

Since $\xi^n(\theta)$ are bounded and $\zeta^1(\theta)$ converges to zero, we have

$$\lim_{\theta \searrow 0} \xi(\theta) = \lim_{\theta \searrow 0} \xi^2(\theta) = \frac{1}{\beta}\left(\gamma \bar{V}^2\right)^{\rho},$$

where $\bar{V}^2$ is given in Lemma A.3. In order to show convergence of the infinitesimal return, observe that

$$\xi^1(\theta)\zeta^1(\theta) = \beta^{-1}\theta\left(\gamma\widetilde{J}^1(\theta)\right)\left[Z^1(\theta) + Z^2(\theta)\right]^{\rho-1}$$

and

$$d\left[\theta\gamma\widetilde{J}^1(\theta)\right] = \theta\gamma\widetilde{J}^1(\theta)\frac{d\theta}{\theta} + \theta\gamma d\widetilde{J}^1(\theta) + \theta\gamma d\widetilde{J}^1(\theta)\frac{d\theta}{\theta}.$$

The drift and volatility coefficients of the first term on the right-hand side vanish as $\theta \searrow 0$ by the proof of Lemma OA.1, and the coefficients of the other two terms vanish by combining the results from the proofs of Lemma OA.1 and Lemma OA.2. Further,

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

and since dZ^1 and dZ^2 have vanishing coefficients by the proof of Lemma OA.2 and the remaining terms are bounded, we obtain that $d\left[\xi^1(\theta)\zeta^1(\theta)\right]$ has vanishing drift and volatility coefficients as $\theta \searrow 0$. The same argument holds for $d\left[\xi^2(\theta)\zeta^2(\theta)\right]$, and thus $d\xi(\theta)$ has vanishing coefficients as well. Therefore all but the first term in

$$dA_t = d\left[\xi(\theta_t)Y_t\right] = A_t\frac{dY_t}{Y_t} + Y_t d\xi(\theta_t) + d\xi(\theta_t)dY_t$$

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

Before we proceed with the proof of Proposition 5.3, we show a limiting result for the continuation value of the 'small' agent in the neighborhood of the boundary $\theta \searrow 0$. Homogeneity of the problem (36)–(37) motivates the guess

$$V_t^1 = \left(A_t^1\right)^{\gamma}\hat{V}^1(\theta_t). \tag{OA.29}$$

While a closed-form solution for $\hat{V}^1(\theta)$ is not available, it is again possible to characterize the asymptotic behavior as $\theta \searrow 0$. The next result will be useful and is stated without proof.

Lemma OA.3 *Let $f : \mathbb{R} \to \mathbb{R}$ be differentiable with a monotone first derivative in a neighborhood of $-\infty$*

and have a finite limit $\lim_{x\to-\infty} f(x)$*. Then* $\lim_{x\to-\infty} f'(x)=0$.

The following lemma establishes the boundary behavior of $\hat{V}^1(\theta)$, in an analogous way as Lemma OA.1 establishes the boundary behavior of the continuation value of the large agent.

Lemma OA.4 $\hat{V}^1(\theta)$ *satisfies*

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

Proof. Transformation (OA.29) together with the previously used $V^1_t=Y^\gamma\widetilde{J}^1(\theta_t)$ imply that

$$\hat{V}^1(\theta)=\beta^\gamma\left(\frac{\gamma\widetilde{J}^1(\theta^1)}{\zeta^1(\theta)^\gamma}\right)^{1-\rho}. \tag{OA.30}$$

Think of $\hat{V}^1$ as a function of $\log\theta$, where we are interested in the limiting behavior as $\log\theta\to-\infty$. We have

$$\theta\hat{V}^1_\theta=\hat{V}^1_{\log\theta}\text{ and }(\theta)^2\hat{V}^1_{\theta\theta}=\hat{V}^1_{(\log\theta)^2}-\hat{V}^1_{\log\theta}. \tag{OA.31}$$

Repeatedly differentiating (OA.30) and exploiting the local behavior of $\widetilde{J}(\theta)$ as $\theta\searrow 0$, we conclude that the assumptions of Lemma OA.3 hold, and thus both expressions in (OA.31) converge to zero as $\theta\searrow 0$. ■

Proof of Proposition 5.3. The drift and volatility coefficients in (37) depend explicitly on θ because A^1 and θ are linked through

$$A^1_t=Y_t\zeta^1(\theta_t)\beta^{-\frac{1}{1-\rho}}\left[\gamma\hat{V}^1(\theta_t)\right]^{\frac{\rho}{\gamma}\frac{1}{1-\rho}}. \tag{OA.32}$$

where we utilized the homogeneity property from (OA.29). Recall that we are interested in the characterization of the limiting solution as $\theta\searrow 0$. The associated HJB equation leads to a second-order ODE (omitting dependence on θ)

$$\begin{aligned}0=\max_{(C^1,\pi^1,\nu^1)}\frac{1}{\rho}\beta^{\frac{1}{1-\rho}}\left(\gamma\hat{V}^1\right)^{1-\frac{\rho}{\gamma}\frac{1}{1-\rho}}+\hat{V}^1\gamma\left(-\frac{\beta}{\rho}+\mu_{A^1}+u^1\sigma_{A^1}+\frac{1}{2}\gamma(\sigma_{A^1})^2\right)+\\+\hat{V}^1_\theta\theta\left(\mu_\theta+u^1\sigma_\theta+\gamma\sigma_\theta\sigma_{A^1}\right)+\hat{V}^1_{\theta\theta}(\theta)^2\frac{1}{2}(\sigma_\theta)^2,\end{aligned} \tag{OA.33}$$

which yields the first-order conditions on C^1_t and π^1_t:

$$\begin{aligned}\frac{C^1_t}{A^1_t}&=\beta^{\frac{1}{1-\rho}}\left(\gamma\hat{V}^1(\theta_t)\right)^{-\frac{\rho}{\gamma}\frac{1}{1-\rho}}\\ \pi^1_t&=\frac{[\xi(\theta_t)]^{-1}+\mu_A(\theta_t)+u_1\sigma_A(\theta_t)-r(\theta_t)+\frac{\theta\hat{V}^1_\theta(\theta_t)}{\hat{V}^1(\theta_t)}\sigma_\theta(\theta_t)\sigma_{A^1}(\theta_t)}{(1-\gamma)(\sigma_A(\theta_t))^2},\end{aligned} \tag{OA.34}$$

where μ_{A^1} and σ_{A^1} are the drift and volatility coefficients on the right-hand side of (37), and μ_θ and σ_θ are the coefficients associated with the evolution of $d\theta_t/\theta_t$ in (26). The portfolio choice π^1 almost corresponds to the standard Merton (1971) result, except the last term in the numerator which explicitly takes into account the covariance between agent's 1 wealth and the evolution in the state variable θ imposed by (OA.32). This term accounts for agent 1's knowledge about the impact of portfolio decisions of the 'small' class of agents on equilibrium prices.

Results from Lemma OA.4 imply that this term, represented by the derivatives of the agent's continuation value, vanishes as $\theta\searrow 0$, and we obtain the limit for $\hat{V}^1(\theta)$ and the evolution of A^1 in closed form. The agent understands that asymptotically as $\theta\searrow 0$ the portfolio decisions made by agents of her type do not have any impact on local equilibrium price dynamics, and thus behaves as if she resided in an economy populated only by agent 2. Utilizing these results from Lemma OA.4 to deduce which terms in ODE (OA.33) vanish and Proposition 5.1 to determine the limiting values of the remaining coefficients, we obtain

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

which is the limiting consumption-wealth ratio for agent 1. The formulas for the wealth share invested in the claim on aggregate consumption and the coefficients of the wealth process are obtained by plugging in the previous results into expressions (37) and (OA.34). ■

Proof of Proposition OA.1. Given convergence to the homogeneous economy counterpart, the expression for $\lim_{\theta\searrow 0}\nu^2(\theta)$ is given by equation (OA.13). Utilizing the formula for the wealth-consumption ratio (34) and the result from Proposition 5.3 then yields

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

The first two terms in the last expression are equal to the limit for $\nu^2(\theta)$, which yields the result for the difference of the discount rates. The expression for part (ii) is obtained by symmetry. ■

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.

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

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." *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.

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

Harrison, J. Michael, and David M. Kreps, 1979. "Martingales and Arbitrage in Multiperiod Securities Markets." *Journal of Economic Theory*. 20 (3): 381–408.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2005. "A Smooth Model of Decision Making under Ambiguity." *Econometrica*. 73 (6): 1849–1892.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2009. “Recursive Smooth Ambiguity Preferences.” *Journal of Economic Theory.* 144 (3): 930–976.

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

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

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.

Merton, Robert C., 1971. “Optimum Consumption and Portfolio Rules in a Continuous-time Model.” *Journal of Economic Theory.* 3 (4): 373–413.

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

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

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

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

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

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

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.