If You're NOT So Smart, Why Are You Rich?

Robust Market Selection with General Recursive Preferences*

Pablo F. Beker† Jaden Yang Chen‡

August 7, 2026

Abstract

This paper develops a unified framework for studying equilibrium consumption dynamics in complete-market stochastic exchange economies with a general class of recursive preferences, allowing for rich heterogeneity in both preferences and beliefs. We characterize long-run allocations under complete markets and reassess the *market selection hypothesis* that markets favor consumers with more accurate beliefs. We show that the failure of market selection is robust across a large class of recursive preferences: multiple preference types—including those with incorrect beliefs—survive regardless of the preferences and beliefs of other participants, so no single belief is selected in the long run. Our findings reveal a fundamental tension between two central economic principles: that markets allocate resources efficiently and that prices eventually reflect correct beliefs.

Keywords: Market selection hypothesis, rational expectations, belief heterogeneity, heterogeneous consumers, recursive preferences, non-expected utility, robustness

*The authors thank Jaroslav Borovička, Pietro Dindo, David Easley, Ani Guerdjikova, Marina Halac, Pellumb Reshidi, Philipp Sadowski, Alvaro Sandroni, Todd Sarver, Lorenzo Stanca and seminar participants at Duke, Sheffield, UNC-Chapel Hill, Virginia Tech, Warwick, Waterloo, SEA 2023, RUD 2024 (Northwestern), Midwest Theory Conference 2024 (Georgia Tech), Triangle Theory Workshop 2024 (NCSU) and the 12th Annual UTDT Conference. We also thank Genyu Li and Hanqing Ye for research assistance. This project was supported by the University of Warwick International Partnership Fund 2023/2024.

†University of Warwick, E-mail: pablo.beker@warwick.ac.uk

‡University of North Carolina at Chapel Hill, E-mail: yangch@unc.edu

1 Introduction

The *market selection hypothesis* (henceforth, MSH) postulates that agents who do not make correct predictions are driven out of financial markets, losing their wealth at the hands of those who make more correct predictions. The MSH has a long tradition in economics: it was first put forward by Alchian (1950) and Friedman (1953) and later applied to financial markets by Cootner (1964) and Fama (1965). It is often invoked to justify rational expectations equilibrium as a solution concept, as it implies that even if agents begin with heterogeneous beliefs, prices are eventually determined by the agents whose beliefs are most accurate—in particular, by agents with correct beliefs whenever they are present. Notably, as originally formulated, the MSH is a preference-neutral claim: it asserts that markets select for more accurate beliefs whatever the traders' attitudes towards risk, intertemporal substitution, or uncertainty, and however those attitudes are distributed in the population.

Sandroni (2000) and Blume and Easley (2006) were the first to study the MSH formally in a general equilibrium framework with intertemporal optimization.[1] They find that the MSH holds in a complete market with bounded endowments when consumers have *subjective expected utility* (SEU hereafter) preferences: survival depends only on the accuracy of beliefs and the discount rate. Among consumers who discount the future identically, attitudes toward risk do not matter. Within the SEU class, therefore, the MSH is indeed preference-neutral. Later work has explored deviations from SEU, but along a single dimension of heterogeneity at a time. One strand allows for heterogeneity in preferences while maintaining homogeneous beliefs (e.g., Condie, 2008; Guerdjikova and Sciubba, 2015). Another allows for heterogeneous beliefs but restricts consumers to identical homothetic Epstein–Zin preferences (e.g., Dindo, 2019; Borovička, 2020). These papers show that the MSH can fail outside the SEU class, and thus that whether markets select for accurate beliefs might depend on preferences. But this raises a question one-dimensional settings cannot answer. If survival depends jointly on preferences and beliefs, then any result obtained by fixing one dimension is, in effect, a statement about a particular correlation between preferences and beliefs—a correlation on which the hypothesis itself is silent—rather than a genuine consequence of equilibrium in complete markets. Therefore, a framework suited to assess the MSH must allow for rich heterogeneity in both dimensions simultaneously.

This paper develops a unified framework to study the MSH when consumers have heterogeneous beliefs and heterogeneous preferences belonging to a broad class of *recursive*

[1] Blume and Easley (1992) were the first to study the MSH in a market clearing model, but they do not allow for intertemporal optimization.

preferences, nesting expected utility and many common non-expected utility preferences as special cases. We allow preferences to take the recursive form

$$V_t = F\left(c_t, \mathcal{R}_t\left(V_{t+1}\right)\right),$$

where F is the time aggregator describing the attitude toward intertemporal substitution and the temporal resolution of uncertainty, and $\mathcal{R}$ is the certainty equivalent aggregator describing the attitude toward the uncertainty in continuation values. Through appropriate choices of F and $\mathcal{R}$, the framework nests almost all preferences studied in the MSH literature, as well as many deviations from SEU not explored before. The point of this generality is not breadth for its own sake: by embedding the existing analyses as special cases of a single selection criterion, the framework lets us pose—and answer—a question that cannot be formulated in those settings: which preferences and beliefs survive *robustly*, that is, regardless of the preferences and beliefs of every other market participant?

An illustrative example. The following economy previews both messages of the paper: (a) a homogeneous-preference model can deliver a clean market-selection prediction within a familiar class, yet that prediction may change once a single additional preference type enters, and (b) economies with more than two consumers cannot be understood through pairwise comparisons. Consider an endowment economy with aggregate risk and two agents with homogeneous Epstein–Zin (EZ) homothetic preferences over two equally likely i.i.d. states; one agent has correct beliefs, the other overestimates the probability of the high-endowment state. As in Borovička (2020) and Dindo (2019), the left panel of Figure 1 maps the regions of the agents' (common) risk aversion $\gamma \geq 0$ and inverse elasticity of intertemporal substitution (IES, hereafter) $\rho \geq 0$ in which the MSH holds or fails.

The left panel suggests a simple lesson—the MSH holds when preferences are close to expected utility (blue) and fails when they are far from it (orange, green, red). The right panel shows that this lesson is not robust: once a third consumer with expected-utility preferences and correct beliefs enters, the two-consumer predictions remain valid only in the dashed regions.[2] Predictions correct within the homogeneous EZ class can thus be overturned by the entry of a single new preference type; neither the failure nor the validity of the MSH in such an economy need carry over to markets with heterogeneity in both dimensions. The rest of the paper makes this precise and characterizes what *does* survive.

[2] In particular, only EZ agents who display preference for the late resolution of uncertainty ($\rho > \gamma$) may survive in the presence of an EU agent with correct beliefs when $\rho > 1$. See Section 7.5

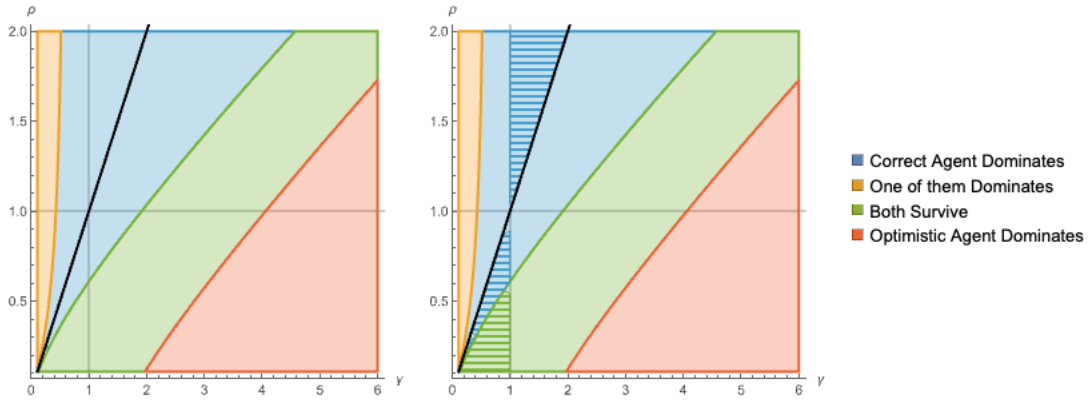

Figure 1: Survival regions in (γ, ρ) space. The left panel considers two homogeneous-preferences EZ agents, one with correct beliefs and one with incorrect beliefs. The right panel adds a third SEU agent with correct beliefs. The two-agent predictions continue to hold only in the dashed region.

Approach and results. We consider an infinite-horizon exchange economy similar to Sandroni (2000) and Blume and Easley (2006), extended to general recursive preferences.[3] Our first main result characterizes the long-run consumption dynamics. As in Sandroni (2000) and Blume and Easley (2006), we track the ratio of any two consumers' marginal utilities, which evolves as the product of two terms: the ratio of their *effective beliefs*, the gradients $\nabla \mathcal{R}_t^i$ of their certainty equivalents, and the ratio of their *effective discount rates*, the derivatives F_2^i. Under SEU these reduce to the subjective beliefs and discount factors and are exogenous; with general recursive preferences they are endogenous, depending on continuation utilities and hence on the equilibrium allocation. This endogeneity is the central difficulty, and we address it with a local approach, in the spirit of Dindo (2019) and Borovička (2020) but extended to our general settings: near an allocation in which one consumer consumes almost the entire endowment, continuation utilities—and thus effective beliefs and discount rates—are approximately exogenous. We summarize the resulting drift by a *local survival index* and say that *i locally dominates j* if *i*'s index exceeds *j*'s when *i* consumes the aggregate endowment. Local dominance is self-reinforcing—the relevant log marginal-utility ratio is a local supermartingale near *i*'s dominance region—so that a locally dominant consumer captures the whole market with positive probability. Under a recurrence condition, the local dominance relation determines long-run survival.

Theorem 1 characterizes the two-consumer case, and the outcome already departs sharply from the SEU benchmark: one consumer may dominate almost surely, the two may coexist

[3] Existence of a competitive equilibrium is established in the Supplementary Material, Section S.6, building on Bewley (1969) and Marinacci and Montrucchio (2010).

almost surely, or exactly one may survive almost surely with each having positive probability of being that survivor. The common mechanism is that markets now select on effective beliefs and effective discount rates rather than subjective beliefs or discount rates alone. The certainty equivalent can reweight states—so that less accurate subjective beliefs may yield more accurate effective beliefs—and alter effective patience; either distortion can let an incorrect-belief consumer survive, or even dominate, alongside a correct-belief one. Because the size of these distortions depends on who bears the aggregate risk, local dominance can differ across neighborhoods of the allocation space, giving rise to coexistence and path-dependent survival.

We then extend the characterization to an arbitrary number of consumers, and this is where the contrast with the existing non-expected utility literature is starkest: with more than two consumers, the analysis is not a pairwise extension of the two-consumer case. Theorems 2 and 3 characterize the survivors in terms of the consumers' *local upper contour sets*, the sets of competitors who locally dominate a given consumer when she consumes the aggregate endowment. If these sets can be completely ranked by inclusion, Theorem 2 shows that the survivors are exactly the locally undominated consumers—those with an empty local upper contour set—each dominating with strictly positive probability, while everyone else vanishes almost surely. When the ranking is incomplete, Theorem 3 still restricts the long-run outcomes, identifying groups that cannot jointly drive out the rest of the market. This many-consumer structure is not a technical refinement: it can overturn the survival verdict that pairwise comparison of two-consumer sub-economies would predict.

The reason is that the survival "tournament" need not be transitive. Under SEU, survival indices are allocation invariant, so the tournament is transitive: a consumer who would vanish against some competitor in isolation also vanishes in any larger economy containing that competitor. With general recursive preferences, by contrast, survival indices are allocation dependent, so the tournament need not be transitive: a verdict in a two-consumer economy need not carry over to a larger one. Example 9 makes the point in the simplest way: a consumer who vanishes almost surely in a two-consumer economy survives almost surely once a third is added, because the third consumer's presence changes who bears the aggregate risk in the relevant near-dominance regions, and with it the effective beliefs that govern selection. This is no curiosity of exotic preferences: Example 10 shows it arises within the homothetic Epstein–Zin class itself, where the MSH can fail for parameters at which the homogeneous-preference analysis predicts it holds, and hold where that analysis predicts it fails. The two-agent, homogeneous-preference characterizations of Dindo (2019) and

Borovička (2020) therefore do not, by themselves, provide a complete basis for general selection results via pairwise comparison. Assessing the MSH in heterogeneous economies requires a joint analysis of preference and belief heterogeneity, as provided by our framework.

Our second main result uses this characterization to examine the *robustness* of the MSH—a question that, to our knowledge, has not been posed in this literature. For this exercise we restrict attention to preferences with a common separable time aggregator that is linear in the certainty equivalent, so that heterogeneity is confined to $\mathcal{R}$. We further require $\mathcal{R}$ to have a well-defined gradient at extinction, which yields a property we call *local expected utility at extinction*: at extinction a consumer's continuation utility is deterministic, so her effective belief is a probability—she prices risk as an SEU agent would—and her effective discount rate is exactly β. The class nests discounted SEU and many common smooth recursive preferences. Together these conditions rule out survival advantages from the discounting technology rather than from attitudes toward uncertainty, which is what makes the robustness question well posed: otherwise survival could always be improved by making a consumer mechanically more patient.[4]

Theorem 4 establishes that the MSH is globally non-robust. No preference—in particular, not SEU with correct beliefs—*robustly dominates*, that is, drives out every configuration of competitors. Instead, there exist multiple *robust survivors*: preferences that survive with positive probability against *any* preferences in the class. We provide a tight characterization: a preference robustly survives if and only if a consumer with it locally dominates one with SEU preferences and correct beliefs.[5] The criterion is clean precisely because of local expected utility at extinction: with each consumer's effective belief pinned down at the boundary, robust survival reduces to a comparison with correct-belief SEU. The implications of this characterization are stark: consumers with incorrect beliefs can robustly survive and, with enough aggregate risk, *any* full-support belief is consistent with robust survival. Thus, whereas Sandroni (2000) and Blume and Easley (2006) show that within the SEU class a unique belief dominates the market, no such belief exists in the general class, and the market may perform very poorly in selecting beliefs.

Theorem 5 shows that the MSH is locally non-robust as well: no preference locally robustly dominates. In two-consumer economies, a preference locally robustly survives if and only if bearing aggregate risk changes the entropy of its effective belief—equivalently,

[4]Linearity and a common β do not close the discounting channel on their own: it is local expected utility at extinction that guarantees the effective belief's mass equals one at the boundary; were it to differ, the effective discount rate would become endogenous and the channel would reopen.

[5]See Section 7.3.1 for a discussion in terms of the consumers' attitudes towards uncertainty.

if and only if the accuracy of the consumer's effective forecasts depends on how much aggregate risk she bears. We call a preference without this property *effectively uncertainty neutral*; SEU, whose effective belief always coincides with the subjective one and is therefore unaffected by risk-bearing, is the leading example. The SEU selection result is therefore a knife-edge local prediction: in any neighborhood of a preference that is not effectively uncertainty neutral, small perturbations leave either both consumers dominating with positive probability or both coexisting almost surely. With more than two consumers, the analogous implication is robust *co-survival*: at least two consumers survive with strictly positive probability.

The analysis above covers the large class of certainty equivalent aggregators displaying local expected utility at extinction—a standard regularity property in the non-expected-utility literature, closely related to the local expected-utility approach of Machina (1982). There are two ways in which this property can fail. The first is that the certainty equivalent may have kinks at extinction, as in Choquet, max-min, or rank-dependent expected-utility aggregators. The second arises for aggregators that are smooth at deterministic continuation utilities in the interior of the utility domain but fail to be differentiable at extinction. An example of the latter is the CRRA certainty equivalent when the extinction utility is exactly zero, as occurs when it is coupled with the homothetic Epstein–Zin time aggregator. We show the robustness question can still be answered in this second case. Here risk aversion and the preference over the timing of the resolution of uncertainty jointly determine robust survival: agents who prefer late resolution (PLRU) cannot robustly survive when $\rho < 1$, while those who prefer early resolution (PERU) cannot survive when $\rho > 1$. That the survival verdict turns on the timing of resolution—a preference dimension the homogeneous benchmark holds fixed—reinforces the message of the introductory example: conclusions drawn from a homogeneous economy need not survive the introduction of preference heterogeneity.

Taken together, our local and global robustness results imply that, in sharp contrast to the SEU case, the robust long-run outcome in a sufficiently rich preference space is the survival of *multiple* consumers with heterogeneous beliefs. This exposes a tension between two principles economists often hold at once: that complete markets allocate resources efficiently, and that prices eventually reflect correct beliefs. In our economy the two cannot be reconciled. Markets still allocate efficiently—every equilibrium is Pareto optimal—but they need not select for correct beliefs: prices can permanently reflect the forecasts of consumers with incorrect beliefs. This happens not because of market frictions, but because recursive preferences transform beliefs and patience into endogenous effective beliefs and

discount rates that can favor incorrect forecasts. Our results thereby provide a foundation for the behavioral finance idea that incorrect beliefs can have a persistent influence on asset prices—one that, we show, requires no friction in a world with sufficiently rich preferences.

Our work is most closely related to Guerdjikova and Sciubba (2015), Dindo (2019) and Borovička (2020), who establish that the MSH can fail when heterogeneity is confined to a single dimension—preferences in the first case, beliefs in the latter two—and within a single class of preferences; we defer a detailed comparison to Section 8. Relative to this literature, our contribution is not another example showing that non-expected utility can matter. Rather, we provide a common local object—the survival index—for a broad recursive class, allow beliefs and preferences to vary jointly across an arbitrary finite set of consumers, and turn the question of whether the SEU selection result survives a richer preference space into global and local robustness theorems. Along the way, as a by-product of the Epstein–Zin analysis, we obtain a closed-form solution to the consumer's problem in complete Arrow–Debreu markets, which we believe to be new and of independent interest.[6]

This paper is organized as follows. Sections 2–4 introduce the economy, characterize equilibrium, and derive the dynamics of marginal-utility ratios; Sections 5 and 6 develop the local approach and the long-run characterization. Section 7 studies robustness and Section 8 reviews the literature. All proofs are in the Appendix and the Supplementary Material.

2 Model Setup

In this section, we introduce the basic notions for our infinite-horizon stochastic exchange economy and define a competitive equilibrium.

2.1 Uncertainty and Preferences

Time is discrete and denoted by $t \in \mathbb{T} = \{0, 1, 2, 3, \ldots\}$. Let S denote the finite set of states of nature. The initial state at time 0 is s_0. At each period $t \geq 1$, a state $s_t \in S$ is realized. Let Σ denote the set of all paths with a typical element $\sigma = (s_1, s_2, \ldots)$. Let Σ_t denote the set of histories up to time t with a typical element $\sigma_t = (s_1, \ldots, s_t)$. Further, $\Sigma(\sigma_t) = \{\tilde{\sigma} \in \Sigma : \tilde{\sigma}_t = \sigma_t\}$, denotes a cylinder set with base σ_t. Let $\mathcal{F}_t$ denote the σ-algebra that consists of all finite unions of the sets $\Sigma(\sigma_t)$, where $\mathcal{F}_0 = \{\emptyset, \Sigma\}$ is the trivial σ-algebra and $\mathcal{F}$ is the σ-algebra generated by $\cup_{t=0}^{\infty} \mathcal{F}_t$. The states of nature are drawn according to a distribution π_0 with full support on S. We also use $\mathbb{P}$ and $\mathbb{E}$ to denote the probability measure

[6] We characterize the optimal allocation and continuation utilities in a discrete-time stochastic economy with finitely many states, for arbitrary summable contingent-claim prices and any values of the IES and risk aversion; the closest results are more restrictive (e.g., unitary IES, continuous-time planner's allocations, or specific asset structures); see Section S.7 of the Supplementary Material for more details.

on $(\Sigma, \mathcal{F})$ and expectation induced by π_0.

The economy is populated by I infinitely-lived consumers and there is a single good every period. With some abuse of notation, we also use I to denote the set of consumers. A consumption plan is a stochastic process $c : \Sigma \times \mathbb{T} \to \mathbb{R}_+$, where $c(\sigma_t) \equiv c(\sigma, t)$, also denoted by c_t, is $\mathcal{F}_t$-measurable. The set of all consumption plans is $\mathbb{C} = L^+_\infty$, where L^+_∞ denotes the set of non-negative adapted stochastic processes on Σ that are essentially bounded. Each consumer is endowed with a consumption plan $e^i \in \mathbb{C} \setminus \{0\}$. The aggregate endowment is denoted by e. Consumers have complete preferences over all feasible consumption plans that are represented by some mapping $V : \mathbb{C} \to L^+_\infty$, where $V_{\sigma_t}(c) \equiv V(c)(\sigma, t)$ denotes the utility of consuming c at σ_t. We assume that consumers have ***recursive preferences*** over the consumption plans which take the following form

$$V_{\sigma_t}(c) = F\left(c_t, \mathcal{R}_{\sigma_t}\left(\boldsymbol{V}_{\boldsymbol{\sigma_{t+1}}}(c)\right)\right), \tag{1}$$

where $F : \mathbb{R}^2 \to \mathbb{R}$ is the time aggregator which describes the attitude toward temporal resolution of uncertainty, $\mathcal{R}_{\sigma_t} : \mathbb{R}^S \to \mathbb{R}$ is the certainty equivalent aggregator which describes the attitude toward the uncertainty in the next period's utility values, and $\boldsymbol{V}_{\boldsymbol{\sigma_{t+1}}}(c) = \left(V_{\sigma_t, s}(c)\right)_{s \in S}$ is the vector of utility values at date $t+1$. This paper focuses on preferences that are monotone in c and continuous in product topology on $\mathbb{C}$ and that can be uniquely pinned down by the recursion.[7] To simplify notation, we denote V_{σ_t} and $\mathcal{R}_{\sigma_t}$ as V_t and $\mathcal{R}_t$ henceforth. Below we introduce some examples of recursive preferences.[8]

Example 1 (Discounted Expected Utility)**.** Let $F(c, y) = u(c) + \beta y$, and $\mathcal{R}_t(V) = \mathbb{E}_\pi(V|\mathcal{F}_t)$. Then, the utility satisfies

$$V_t(c) = u(c_t) + \beta \times \mathbb{E}_\pi\left(V_{t+1}(c)\,|\mathcal{F}_t\right),$$

which represents discounted expected utility preferences. Survival in economies with these preferences is studied by Sandroni (2000) and Blume and Easley (2006).

Example 2 (Epstein–Zin Preferences)**.** Let $F(c, y) = \left((1-\beta)c^{1-\rho} + \beta y^{1-\rho}\right)^{\frac{1}{1-\rho}}$ and

[7] These properties can also be proved from more primitive assumptions such as monotonicity and continuity in both aggregators, as shown in Supplementary Material, Section S.6 (also see Marinacci and Montrucchio, 2010). Moreover, they are satisfied by most preferences of interest, so we assume them directly.

[8] In addition to previous examples, there are many other important recursive preferences that have not been explored in the market selection literature, including rank-dependent utility (Quiggin, 1982), uncertainty averse preferences (Cerreia-Vioglio et al., 2011; Strzalecki, 2013), dynamic preference for flexibility (Krishna and Sadowski, 2014), cautious expected utility (Cerreia-Vioglio et al., 2015) and dynamic mixture-averse preferences (Sarver, 2018) and many others.

$\mathcal{R}_t(V) = \left[\mathbb{E}_\pi\left(V^{1-\gamma}(c)\,|\mathcal{F}_t\right)\right]^{\frac{1}{1-\gamma}}$. Then, these preferences satisfy

$$V_t(c) = \left((1-\beta)\,c_t^{1-\rho} + \beta\left[\mathbb{E}_\pi\left(V_{t+1}^{1-\gamma}(c)\,|\mathcal{F}_t\right)\right]^{\frac{1-\rho}{1-\gamma}}\right)^{\frac{1}{1-\rho}},$$

which are Epstein–Zin preferences with a homothetic aggregator, see Epstein and Zin (1989). Here, $1/\rho$ stands for the elasticity of intertemporal substitution (EIS), and γ stands for the coefficient of relative risk aversion. The survival in economies with these preferences is studied by Borovička (2020) and Dindo (2019).

Example 3 (Recursive Smooth Ambiguity)**.** Suppose that consumers face model uncertainty and entertain a set of distributions $\mathbb{M} = \{\pi_1, \ldots, \pi_M\}$. Let $F(c, y) = u(c) + \beta y$, and $\mathcal{R}_t(V) = \phi^{-1}\left(\mathbb{E}_\mu \phi\left[\mathbb{E}_\pi\left(V|\mathcal{F}_t\right)\right]\right)$, where μ stands for the distribution over $\mathbb{M}$. So, we have

$$V_t(c) = u(c_t) + \beta \times \phi^{-1}\left(\mathbb{E}_\mu \phi\left[\mathbb{E}_\pi\left(V_{t+1}(c)\,|\mathcal{F}_t\right)\right]\right),$$

which are the recursive smooth ambiguity preferences (Klibanoff et al., 2009). The survival in economies where consumers have these preferences and homogeneous beliefs is studied by Guerdjikova and Sciubba (2015). Condie (2008) studies the extreme case of the max-min preferences with heterogeneous beliefs.

In what follows we allow consumers to have heterogeneous preferences and use superscript i to denote consumer i's preferences.

2.2 Competitive Equilibrium

A contingent contract is a promise to deliver one unit of the consumption good contingent on some realization of the observable states of the world and zero otherwise. We assume there is a complete set of competitive markets to trade contingent contracts that opens at date zero.

A price system is a stochastic process $p : \Sigma \times \mathbb{T} \to \mathbb{R}$ and we define $p(\sigma_t) \equiv p(\sigma, t)$ to be the price of a contingent contract for σ_t. Let $\mathbb{Q} = \left\{p \geq 0 : \sum_{t=0}^{\infty}\sum_{\sigma_t \in \Sigma_t} p(\sigma_t) < \infty\right\}$ denote the set of summable contingent commodity prices. Consumer i's problem given price system p is defined as follows:

$$\begin{aligned}
&\max_{c \in \mathbb{C}} \quad V_0^i(c) \\
&\text{s.t.} \sum_{t=0}^{\infty}\sum_{\sigma_t \in \Sigma_t} p(\sigma_t) \times c(\sigma_t) \leq \sum_{t=0}^{\infty}\sum_{\sigma_t \in \Sigma_t} p(\sigma_t) \times e^i(\sigma_t), \\
&\quad c(\sigma_t) \geq 0 \text{ for all } \sigma \in \Sigma \text{ and } t \geq 0
\end{aligned}$$

where σ_0 is given and V^i satisfies (1). A competitive equilibrium is an allocation and a price system such that consumers' decisions are optimal in their budget sets and markets clear.

Definition 1. A ***competitive equilibrium*** is $(c, p) \in \mathbb{C}^I \times \mathbb{Q}$ such that:

1. For every i, c^i solves consumer i's problem given price system p;
2. $\sum_i c^i(\sigma_t) = e(\sigma_t)$ for all $\sigma_t \in \Sigma_t$ and all $t \in \mathbb{T}$.

3 Characterization of the Equilibrium

To characterize the equilibrium via a system of equations, we impose Assumptions 1 and 2. These assumptions make the time aggregator smooth and well behaved at the boundary, and the certainty equivalent smooth and strictly increasing in continuation utilities.

Assumption 1. *$F^i(x, y)$ is continuously differentiable, $F^i \geq 0$, and $F_1^i, F_2^i > 0$ and $\lim_{x \to 0} F_1^i(x, y) = +\infty$ for all $y \geq 0$.*

Assumption 2. *$\mathcal{R}^i_{\sigma_t}(V)$ is continuously differentiable and strictly increasing for all $V \in \mathbb{R}^S_{++}$, $\sigma_t \in \Sigma_t$ and $t \geq 1$.*

The following proposition shows that a competitive equilibrium can be characterized by a system of first-order conditions. The result follows from the one-deviation property of optimal plans together with the monotonicity and boundary properties of F imposed in Assumptions 1 and 2.

Proposition 1. *Suppose Assumptions 1 and 2 hold. If (c, p) is a competitive equilibrium, then for every consumer i and every $\sigma_t \in \Sigma_t$,*

$$\frac{p(\sigma_t, s)}{p(\sigma_t)} = \frac{F_1^i\left(c^i(\sigma_t, s), \mathcal{R}^i_{t+1}\left(V^i_{t+2}\right)\right)}{F_1^i\left(c^i(\sigma_t), \mathcal{R}^i_t\left(V^i_{t+1}\right)\right)} \times \nabla_s \mathcal{R}^i_t\left(V^i_{t+1}\right) \times F_2^i\left(c^i(\sigma_t), \mathcal{R}^i_t\left(V^i_{t+1}\right)\right), \quad (2)$$

where $\nabla_s \mathcal{R}^i_t\left(V^i_{t+1}\right) \equiv \partial \mathcal{R}^i_t\left(V^i_{t+1}\right) / \partial V^i_{\sigma_t, s}$.

This condition is the familiar first-order condition equating the marginal rate of substitution between present and future consumption to their relative price at any interior optimum. Unlike the discounted expected-utility case, where the marginal rate of substitution depends only on consumption at σ_t and (σ_t, s_{t+1}), here it also depends on continuation utilities. This dependence arises because the ***generalized marginal utility*** of consumption at σ_t, $F_1^i\left(c^i(\sigma_t), \mathcal{R}^i_t\left(V^i_{t+1}\right)\right)$, depends not only on the consumption at σ_t but also on the certainty equivalent $\mathcal{R}^i_t$. Likewise, the marginal utility of consuming at (σ_t, s) depends not only on the consumption level at (σ_t, s) but also on the continuation utilities at dates $t+1$ and $t+2$. Specifically, it equals the marginal utility of the certainty equivalent at σ_t, F_2^i, multiplied by the marginal effect of consumption at (σ_t, s) on the certainty equivalent, $\nabla_s \mathcal{R}^i_t$, and by the marginal utility of consumption.

In what follows, we refer to $\nabla\mathcal{R}_t^i \equiv \left(\nabla_s \mathcal{R}_t^i\right)_{s\in S}$ as consumer i's ***effective belief*** and F_2^i as the ***effective discount rate***. In the special case of subjective expected utility, the effective belief coincides with the subjective belief π^i and the effective discount rate equals the discount factor β^i. More generally, $\nabla\mathcal{R}_t^i$ represents the gradient of the certainty equivalent and need not add up to one.

4 Consumption Dynamics

This section introduces the notions of survival and derives a recursive description of consumption dynamics, in the spirit of Sandroni (2000) and Blume and Easley (2006).

Definition 2. Consumer i ***survives*** on a path σ if $\limsup_{t\to\infty} c_t^i(\sigma) > 0$, ***vanishes*** on a path σ if $\lim_{t\to\infty} c_t^i(\sigma) = 0$ and ***dominates*** on a path σ if all other consumers vanish on that path.

We impose boundedness of endowments and normalize the certainty equivalent at constant continuation utilities.

Assumption 3. *There exists $0 < m < M < +\infty$ such that $m < e^i(\sigma_t) < M$ for all $\sigma_t \in \Sigma_t$ and $t \in \mathbb{T}$.*

Assumption 4. *$\mathcal{R}_t^i(k \cdot \mathbf{1}) = k$ for all $i \in I$, $t \in \mathbb{T}$ and $k \geq 0$.*

We focus throughout on equilibrium consumption allocations and analyze survival by studying the dynamics of marginal utility ratios between any two consumers. Lemma 1 establishes a key property used throughout the paper.

Lemma 1. *Suppose Assumptions 1-4 hold. Then we have*

$$L^{ij}(\sigma_t) \equiv \frac{F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right)}{F_1^j\left(c^j(\sigma_t), \mathcal{R}_t^j\left(V_{t+1}^j\right)\right)} \to \infty \text{ for some } j \Longrightarrow c^i(\sigma_t) \to 0 \quad \mathbb{P}\text{-a.s.},$$

*where $L^{ij}(\sigma_t)$ is the **generalized marginal utility ratio** between i and j.*

Lemma 1 implies that consumer i vanishes on any path along which her generalized marginal utility ratio relative to some other consumer diverges. To study these dynamics, we divide the first-order conditions (2) for any pair of consumers, yielding the following result.

Proposition 2. *Suppose Assumptions 1 and 2 hold. In any competitive equilibrium,*

$$L^{ij}(\sigma_t, s_{t+1}) = L^{ij}(\sigma_t) \times B^{ji}(\sigma_t, s_{t+1}) \times D^{ji}(\sigma_t), \tag{3}$$

for every i, j and $\sigma_t \in \Sigma_t$ and $s_{t+1} \in S$, where

$$B^{ji}(\sigma_t, s) \equiv \frac{\nabla_s \mathcal{R}_t^j\left(V_{t+1}^j\right)}{\nabla_s \mathcal{R}_t^i\left(V_{t+1}^i\right)}, \text{ and } D^{ji}(\sigma_t) \equiv \frac{F_2^j\left(c^j(\sigma_t), \mathcal{R}_t^j\left(V_{t+1}^j\right)\right)}{F_2^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right)},$$

*are the **effective belief ratio** and **discount rate ratio** respectively between j and i.*

Equation (3) shows that if consumer i has a higher effective belief on state s and a higher effective discount rate than consumer j—that is, $B^{ji}(\sigma_t, s) < 1$ and $D^{ji}(\sigma_t) < 1$—then i's generalized marginal utility decreases relative to j when state s is realized. This induces a tendency for consumer i's consumption to increase relative to consumer j's. Intuitively, a higher effective belief assigns greater marginal value to continuation utility in state s, while a higher effective discount rate reflects greater effective patience. Both shift consumption toward state s or the future. We illustrate these dynamics in several special cases.

Example 4 (Dynamics of discounted expected utility)**.** Suppose consumers have discounted expected utility preferences with i.i.d. beliefs. Then (3) reduces to

$$\frac{u_i'\left(c^i\left(\sigma_{t+1}\right)\right)}{u_j'\left(c^j\left(\sigma_{t+1}\right)\right)} = \frac{u_i'\left(c^i\left(\sigma_t\right)\right)}{u_j'\left(c^j\left(\sigma_t\right)\right)} \times \frac{\pi^j\left(s_{t+1}\right)}{\pi^i\left(s_{t+1}\right)} \times \frac{\beta^j}{\beta^i}, \tag{4}$$

as in Sandroni (2000), Blume and Easley (2006, 2009). If the state process is also i.i.d., then

$$\begin{aligned}\lim_{T\to+\infty} \frac{1}{T} \log \frac{u_i'\left(c^i\left(\sigma_T\right)\right)}{u_j'\left(c^j\left(\sigma_T\right)\right)} &= \lim_{T\to+\infty} \frac{1}{T} \sum_{t=1}^{T} \log \frac{\pi^j\left(s_t\right)}{\pi^i\left(s_t\right)} + \log \frac{\beta^j}{\beta^i} \\ &\to I\left(\pi^i\right) - I\left(\pi^j\right) + \log \frac{\beta^j}{\beta^i} \quad \mathbb{P}-a.s.,\end{aligned}$$

where $I\left(\pi^i\right) = -\mathbb{E}\log\left[\pi^i\left(s\right)\right]$ represents the relative entropy of π^i. Suppose that $\beta^i = \beta^j$, then the consumer with a larger entropy—thus less precise belief—will vanish in the limit, as predicted by the market selection hypothesis.

In the expected utility case, the marginal utility ratio follows an exogenous Markov process depending only on beliefs and discount factors, allowing the dynamics to be analyzed independently of the equilibrium allocation. This property relies on the linearity of the certainty equivalent and the separability of the time aggregator. With general recursive preferences, by contrast, effective beliefs and discount rates depend endogenously on continuation utilities. Below are two examples.

Example 5 (Dynamics of Epstein–Zin preferences)**.** Suppose consumers have Epstein–Zin preferences with i.i.d. beliefs. Then

$$F_2^i = \underbrace{\left(\frac{V^i_{\sigma_t}}{\mathcal{R}^i_t(V^i_{t+1})}\right)^{\rho^i}}_{\text{time adjustment}} \beta^i, \quad \nabla_s \mathcal{R}^i_t\left(V^i_{t+1}\right) = \underbrace{\left(\frac{V^i_{\sigma_t,s}}{\mathcal{R}^i_t(V^i_{t+1})}\right)^{-\gamma^i}}_{\text{risk adjustment}} \pi^i\left(s\right), \tag{5}$$

and the generalized marginal utility is $F_1^i = \left(1-\beta^i\right)\left(V^i_{\sigma_t}/c^i_t\right)^{\rho^i}$. The dynamics of the

generalized marginal utility ratios are described by

$$L^{ij}\left(\sigma_t, s_{t+1}\right) = L^{ij}\left(\sigma_t\right) \times \frac{\left(\frac{V^j_{\sigma_t, s_{t+1}}}{\mathcal{R}^j_t\left(V^j_{t+1}\right)}\right)^{-\gamma^j} \pi^j\left(s_{t+1}\right)}{\left(\frac{V^i_{\sigma_t, s_{t+1}}}{\mathcal{R}^i_t\left(V^i_{t+1}\right)}\right)^{-\gamma^i} \pi^i\left(s_{t+1}\right)} \times \frac{\left(\frac{V^j_{\sigma_t}}{\mathcal{R}^j_t\left(V^j_{t+1}\right)}\right)^{\rho^j} \beta^j}{\left(\frac{V^i_{\sigma_t}}{\mathcal{R}^i_t\left(V^i_{t+1}\right)}\right)^{\rho^i} \beta^i}. \tag{6}$$

Relative to the discounted expected utility case (4), the dynamics feature endogenous time and risk adjustments to standard beliefs and discount rates. The time adjustment depends on current utility $V^i_{\sigma_t}$ relative to the date-t certainty equivalent of future utility: when current utility is low (high), the consumer becomes more impatient (patient) and discounts the future more (less) heavily. The risk adjustment for state s_{t+1} depends on the continuation utility $V^i_{\sigma_t, s_{t+1}}$ relative to the same certainty equivalent: states with lower (higher) continuation utility receive greater (smaller) weight. Thus, better continuation utilities make the EZ agent more patient but also more pessimistic about that state, so her survival chances depend on which effect dominates, governed by ρ^i and γ^i.[9]

Example 6 (Dynamics of smooth ambiguity)**.** Suppose consumers have recursive smooth ambiguity preferences with i.i.d. beliefs. Then

$$F^i_2 = \beta^i, \quad \nabla_s \mathcal{R}^i_t\left(V^i_{t+1}\right) = \mathbb{E}_{\mu^i}\left(\underbrace{\frac{\phi_i'\left[\mathbb{E}_\pi\left(V^i_{t+1}\left(c\right)\right)\right]}{\phi_i'\left(\mathcal{R}^i_t(V^i_{t+1})\right)}}_{\text{ambiguity adjustment}} \pi\left(s\right)\right).$$

and the generalized marginal utility is $F^i_1 = u'(c_i)$. Note that effective discount rates and the generalized marginal utility are both standard because F is separable. The dynamics of generalized marginal utility ratios follow

$$L^{ij}\left(\sigma_t, s_{t+1}\right) = L^{ij}\left(\sigma_t\right) \times \frac{\mathbb{E}_{\mu^j}\left(\frac{\phi_j'\left[\mathbb{E}_\pi\left(V^j_{t+1}(c)\right)\right]}{\phi_j'\left(\mathcal{R}^j_t(V^j_{t+1})\right)} \times \pi\left(s_{t+1}\right)\right)}{\mathbb{E}_{\mu^i}\left(\frac{\phi_i'\left[\mathbb{E}_\pi\left(V^i_{t+1}(c)\right)\right]}{\phi_i'\left(\mathcal{R}^i_t(V^i_{t+1})\right)} \times \pi\left(s_{t+1}\right)\right)} \times \frac{\beta^j}{\beta^i}, \tag{7}$$

which corresponds to the dynamics in Guerdjikova and Sciubba (2015). When consumers are ambiguity neutral, ϕ' is constant and (7) becomes (4) with $\pi^i = \mathbb{E}_{\mu^i}(\pi)$. When consumers

[9] Under SEU ($\rho^i = \gamma^i$) the two cancel and continuation utilities do not matter. When $\rho^i > \gamma^i$ ($\rho^i < \gamma^i$), the effective-discount-rate (belief) effect dominates and the product of effective discount rates and beliefs increases (decreases) with continuation utilities. When $\rho^i = 1$, that product is a probability measure—a property that will be useful for studying survival in Example 10.

are ambiguity-averse, there is an ambiguity adjustment for each model π^i: models with lower (higher) expected future utility receive greater (smaller) weight.

These examples illustrate that, with general recursive preferences, the evolution of generalized marginal utilities is endogenous and potentially complex. To analyze long-run consumption dynamics, we therefore adopt a *local approach*, studying approximate dynamics when a given consumer nearly dominates the market and consumes almost all of the aggregate endowment. This approach allows endogenous continuation values to be approximated by exogenously determined objects and is developed in detail in the following sections.

5 Local Survival Index

In this section we introduce some concepts to implement the aforementioned local approach. For simplicity, we focus on the stationary case and impose the following assumption.

Assumption 5 (Stationarity)**.** *States and endowments are i.i.d., and* $\mathcal{R}_t^i = \mathcal{R}^i$ *for all* $t \in \mathbb{T}$.

Assumption 5 imposes a stationary structure on uncertainty. First, uncertainty manifests itself in a time-independent manner: states and endowments are i.i.d. distributed. Second, consumers have time-independent preferences toward uncertainty; for example, beliefs are i.i.d. and there is no learning. This assumption simplifies the analysis and allows us to focus on the pure effect of market selection.[10]

Definition 3 (Local utility)**.** Under Assumption 5, we define $\underline{V}^i = V^i_{\sigma_t}(0)$ and $\overline{V}^i_{s_t} = V^i_{\sigma_t}(e)$ for all $s_t \in S$ and $i \in I$.

In words, $\underline{V}^i$ denotes the continuation utility of consumer i when she consumes zero in every period, while $\overline{V}^i_s$ denotes the continuation utility when the current state is s and she consumes the aggregate endowment in every period. Under Assumption 5, $\underline{V}^i$ is constant and $\overline{V}^i_s$ depends only on the current state. Both objects can be obtained by solving a system of equations.[11]

We now define the local survival index.

Definition 4 (Local survival index)**.** For all i, j, define

$$\mathcal{S}_i^j \equiv \begin{cases} \mathbb{E}\log\left[\nabla_s \mathcal{R}^i\left(\overline{V}^i\right)\right] + \mathbb{E}\log F_2^i\left(e(s), \mathcal{R}^i\left(\overline{V}^i\right)\right), & \text{if } j = i \\ \mathbb{E}\log\left[\nabla_s \mathcal{R}^j\left(\underline{V}^j\right)\right] + \mathbb{E}\log F_2^j\left(0, \mathcal{R}^j\left(\underline{V}^j\right)\right). & \text{if } j \neq i. \end{cases}$$

[10] We believe that our approach can be extended to more general environments. For instance, if there is learning but beliefs converge, the same analysis applies. More generally, one may redefine the entropy using time averages to characterize non-i.i.d. dynamics.

[11] From (1), we know that $\underline{V}^i$ and $\overline{V}^i_s$ satisfy $\underline{V}^i = F^i\left(0, \mathcal{R}^i(\underline{V}^i)\right)$ and $\overline{V}^i_s = F^i\left(e(s), \mathcal{R}^i(\overline{V}^i)\right)$ for all $s \in S$.

We call $\mathcal{S}_i^j$ the ***local survival index*** of consumer j when consumer i dominates.

The local survival index $\mathcal{S}_i^j$ has two components. The first component is the expected log, or the negative of the relative entropy, of consumer j's effective belief when i dominates the market, i.e., consumes the total endowment.[12] The second component is the expected log of consumer j's effective discount rate when i dominates. A consumer has a higher survival index if her effective belief has a lower entropy, and if her effective discount rate is higher. A special case is the discounted expected utility where $\mathcal{S}_i^j = -I(\pi^j) + \log \beta^j$ for all i, j.

Remark 1 (Well-definedness). For some preferences, the formula for the local survival index may not be well defined at the boundary because an argument of a logarithm vanishes or because either $\nabla \mathcal{R}^j$ or F_2^j fails to exist. In such cases, the survival index is defined by taking the relevant limit along the equilibrium path, as discussed in Section 7.5.[13]

Now we introduce the following concept.

Definition 5 (Local Dominance)**.** For all i, j, we denote by

$$i \succ_i j \quad \text{if } \mathcal{S}_i^i > \mathcal{S}_i^j,$$

and say that i ***locally dominates*** j. Define $j \succ_i i$ if the inequality is reversed.

Definition 5 states that i locally dominates j if, when i consumes almost the entire endowment, her survival index exceeds that of j. As shown in the next section, local dominance serves as a key tool for analyzing equilibrium dynamics.

6 Characterizations of Long-Run Dynamics

6.1 Assumptions

To implement the local approach, we introduce an irreducibility assumption, that is every consumer has positive probability of almost dominating the market, or equivalently, the marginal utility ratio between any consumer i and all other consumers can be arbitrarily small with positive conditional probability at every date.

Assumption 6 (Irreducibility)**.** *For all* $0 < \underline{L} < \overline{L} < +\infty$ *and* $i \in I$*, there exists* $K < \infty$ *such that whenever* $\max_{j \neq i} L^{ij}(\sigma_t) < \overline{L}$*, we have*

$$\max_{j \neq i} L^{ij}(\sigma_t, s_1, \ldots, s_K) < \underline{L} \text{ for some } s_1, \ldots, s_K \in S,$$

[12] With some abuse of language, we refer to $-\mathbb{E} \log \left(\nabla_s \mathcal{R}^i \right)$ as the relative entropy of i's effective belief though $\nabla \mathcal{R}^i$ may not be a probability.

[13] Alternatively, the survival index can be defined in terms of equilibrium prices, which covers more general cases. The main text uses the current definition because it yields a simpler expression in terms of primitives and is more convenient to many of the preference classes studied in the paper.

with each signal realized with strictly positive probability, for all $t \in \mathbb{T}$ *and* $\sigma_t \in \Sigma_t$.

Notice this is an assumption on endogenous variables as L^{ij} depends on the equilibrium consumption allocation. In the Supplementary Material, Section S.2 we provide conditions on exogenous variables to ensure this assumption holds.[14] We retain Assumption 6 in the main text because it explicitly highlights the crucial property of the marginal utility ratio.

To simplify the discussion, we also impose the following assumption.

Assumption 7 (Bounded Patience)**.** *For every pair of consumers* i, j*, there exist constants* $0 < \underline{d}_{ij} \leq \overline{d}_{ij} < \infty$ *such that* $D^{ji}(\sigma_t) \in \left[\underline{d}_{ij}, \overline{d}_{ij}\right]$ *for all* t *and* σ_t.

Assumption 7 requires relative effective discount rates to remain uniformly bounded along equilibrium paths. The assumption rules out cases in which one consumer becomes infinitely more or less patient than another near the extinction boundary. The role of this assumption is explained below.

6.2 A Relevant Process

It turns out that a highly relevant process for the generalized marginal utility ratio, $\{L^{ij}\}$, is a modified process, $\{\hat{L}^{ij}\}$, where

$$\hat{L}^{ij}_{t+1}\left(\sigma_t, s_{t+1}\right) = \hat{L}^{ij}_t\left(\sigma_t\right) \times B^{ji}\left(\sigma_t, s_{t+1}\right) \times D^{ji}\left(\sigma_t, s_{t+1}\right). \tag{8}$$

The process $\hat{L}^{ij}$ differs from the original dynamics (3) only in that the relative effective discount rate is evaluated one period later. This shift pairs the effective belief term and the effective discount term with the same realized next state:

$$\mathbb{E}\left[\log \hat{L}^{ij}_{t+1} \mid \mathcal{F}_t\right] = \log \hat{L}^{ij}_t + \mathbb{E}\left[\log B^{ji}\left(\sigma_t, s_{t+1}\right) + \log D^{ji}\left(\sigma_t, s_{t+1}\right) \middle| \mathcal{F}_t\right].$$

Thus the drift comes cleanly from effective beliefs and effective discounting. Near the region where consumer i dominates, the drift becomes $\mathcal{S}^j_i - \mathcal{S}^i_i$. Hence the local survival index determines the local drift of the modified process.

Assumption 7 ensures that the modified process is uniformly comparable to the original generalized marginal utility ratio: the ratio $L^{ij}/\hat{L}^{ij}$ is uniformly bounded away from zero and infinity. Therefore, any local convergence or supermartingale argument for $\hat{L}^{ij}$ transfers back to L^{ij}, so the two processes imply the same survival conclusions.[15]

[14] For example, the assumption holds as long as consumers' disagreement, as captured by their belief ratio, about one of the states is large enough. Later in Section 7.2, we also show that this assumption holds when we allow for small perturbations of the economy.

[15] One could work directly with the original dynamics, but then the effective discount term is predetermined at date t, so the proof requires a state-dependent normalization. Conclusions from the normalized process transfer back to survival only if this normalization remains uniformly comparable to the original marginal-utility ratio, which again requires a boundedness condition.

More generally, the paper's characterization results extend beyond Assumption 7. They apply to *any* positive process admitting a decomposition of the form (8) and serving as the relevant variable for relative survival. For example, standard homothetic Epstein–Zin preferences may fail Assumption 7, but their relative consumption or wealth dynamics can often be written in an analogous form, thus the same local-index logic applies, see Example 10 below.[16]

6.3 A General Dominance Criterion

We begin with a general principle: a consumer can dominate the market with positive probability if and only if she is locally dominant within her own near-dominance neighborhood (up to irreducibility).

Proposition 3 (Conditions for Dominance)**.** *Under Assumptions 1–5 and 7, for every $i \in I$:*

(i) If there exists $j \in I$ with $j \succ_i i$, then consumer i dominates the market with $\mathbb{P}$-zero probability.

(ii) If Assumption 6 holds and $i \succ_i j$ for all $j \neq i$, then consumer i dominates the market with strictly positive $\mathbb{P}$-probability.

Proposition 3 provides a sharp characterization of positive-probability market dominance. Local dominance is necessary for a consumer to dominate, and—under irreducibility—also sufficient. Part (i) identifies consumers who can never dominate the market and does not rely on Assumption 6. Part (ii) uses irreducibility to ensure that the near-dominance neighborhood of i is reached with strictly positive probability.[17] Proposition 3 identifies which consumers can dominate with positive probability, but it does not characterize survival along other sample paths. In the next subsection, we show that under additional structure one can identify the entire set of surviving consumers.

6.4 Two-Consumer Economies

We first show that in a two-consumer economy the consumption dynamics can be fully characterized.

Theorem 1 (Two-Consumer Economy)**.** *Suppose $|I| = 2$ and Assumptions 1–7 hold.*

(i) If $i \succ_i j$ and $i \succ_j j$, then i survives and j vanishes $\mathbb{P}$-almost surely.

[16] In particular, one can show that the Epstein–Zin dynamics (6) can be equivalently written as

$$\frac{c^i(\sigma_t, s_{t+1})^{-\rho_i}}{c^j(\sigma_t, s_{t+1})^{-\rho_j}} = \frac{c^i(\sigma_t)^{-\rho_i}}{c^j(\sigma_t)^{-\rho_j}} \times B^{ji}(\sigma_t, s_{t+1}) \times \tilde{D}^{ji}(\sigma_t, s_{t+1}),$$

which is already the form of (8) with $\hat{L}_t^{ij} = \frac{c^i(\sigma_t)^{-\rho_i}}{c^j(\sigma_t)^{-\rho_j}}$. See Supplementary Material, Section S.7.

[17] The proof only requires that this neighborhood is visited with strictly positive probability. This is weaker than Assumption 6, which requires the probability to be bounded away from zero uniformly over t and σ_t.

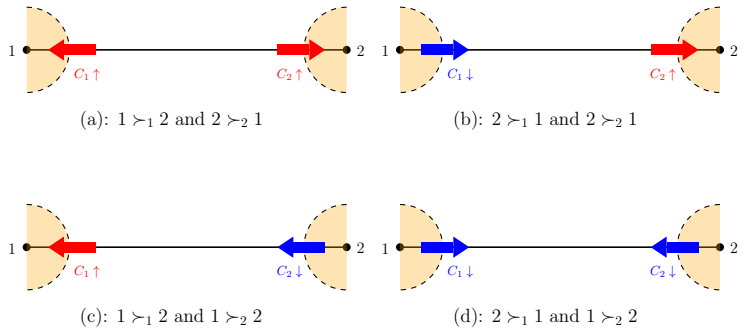

Figure 2: Two-Consumer Case

(ii) If $i \succ_j j$ and $j \succ_i i$, then i and j survive together $\mathbb{P}$-almost surely.

(iii) If $i \succ_i j$ and $j \succ_j i$, then exactly one consumer survives $\mathbb{P}$-almost surely, and each survives with $\mathbb{P}$-strictly positive probability.

Theorem 1 shows that, under non-expected utility, a rich variety of long-run outcomes can arise. This contrasts with the SEU benchmark, in which a unique survivor typically emerges except in knife-edge tie cases. Figure 2 depicts the four possible configurations of the local dominance relation. Panel (a) illustrates case (iii), where both consumers are locally dominant. To see the mechanism, consider consumer 1. The condition $1 \succ_1 2$ implies that when consumer 1 is close to dominating the market, her relative consumption share increases in expectation. Thus, near-dominance is self-reinforcing. The same reasoning applies symmetrically to consumer 2.

Assumption 6 ensures that the economy repeatedly visits neighborhoods in which either consumer is close to dominance with uniformly positive probability. This recurrence allows us to strengthen the positive-probability dominance result into an almost-sure statement.

6.4.1 Examples

We illustrate Theorem 1 with two preference specifications that generate survival patterns absent under expected utility.

Example 7 (Survival Patterns Under Risk-Sensitive Preferences)**.** Let $S = \{s_1, s_2\}$ and $I = \{1, 2\}$. Consumers have i.i.d. subjective beliefs $\pi^1 = (1/3, 2/3)$ and $\pi^2 = (3/4, 1/4)$, while the true distribution is $\pi_0 = (1/2, 1/2)$. Normalize the utility of the aggregate endowment by setting $u(s_1) = 2$ and $u(s_2) = 1$, where $u(s) \equiv u(e(s))$. Consumers have risk-sensitive preferences (Hansen and Sargent, 1995):

$$V^i_{\sigma_t} = u(c_t) - \beta \frac{1}{\gamma^i} \log\left(\sum_{s_{t+1} \in S} \exp\left(-\gamma^i V^i_{\sigma_t, s_{t+1}}\right) \pi^i(s_{t+1}) \right),$$

where $\gamma^i > 0$ measures risk sensitivity. When consumer i consumes the entire endowment,

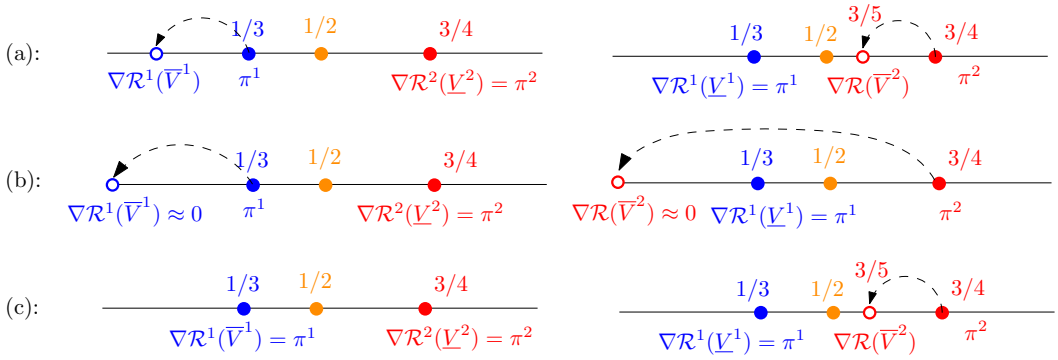

Figure 3: Survival with non-expected utility in Example 7

Note: All beliefs depicted in the figure, e.g., π^i and $\nabla\mathcal{R}^i$, denote the probability on the first state, s_1. Graphs on the left (right) depict effective beliefs when consumer 1 (2) consumes all the endowment.

effective beliefs satisfy

$$\nabla_s\mathcal{R}^j(\underline{V}^j)=\pi^j(s),\qquad \nabla_s\mathcal{R}^i(\overline{V}^i)=\frac{e^{-\gamma^i u(s)}\pi^i(s)}{\sum_{s'\in S}e^{-\gamma^i u(s')}\pi^i(s')}. \tag{9}$$

Since $u(s_2)<u(s_1)$, risk sensitivity tilts effective beliefs toward the low-endowment state s_2. This distortion can either move beliefs farther from the truth or closer to it, depending on whether the consumer is initially pessimistic or optimistic.

Under SEU ($\gamma_1=\gamma_2=0$), consumer 1 survives since π^1 is closer to π_0 (Sandroni, 2000; Blume and Easley, 2006). With risk sensitivity, we obtain the following patterns.

1. The less correct consumer dominates. Let $\gamma_1=\gamma_2=\ln 2$. Then

$$\left(\nabla\mathcal{R}^1(\overline{V}^1),\,\nabla\mathcal{R}^2(\underline{V}^2)\right)=\left((1/5,4/5),(3/4,1/4)\right)\;\Rightarrow\;2\succ_1 1,$$

and

$$\left(\nabla\mathcal{R}^1(\underline{V}^1),\,\nabla\mathcal{R}^2(\overline{V}^2)\right)=\left((1/3,2/3),(3/5,2/5)\right)\;\Rightarrow\;2\succ_2 1.$$

Mechanism. When a consumer bears aggregate risk, risk sensitivity shifts her effective belief toward the bad state s_2. This shift makes consumer 1, already pessimistic about s_1, even less accurate, while partially correcting consumer 2's optimism. Thus consumer 2 is locally more accurate in both neighborhoods and dominates the market $\mathbb{P}$-a.s. (Figure 3(a)).

2. Coexistence. Let $\gamma_1,\gamma_2\to+\infty$. Then $\nabla\mathcal{R}^i(\overline{V}^i)\to(0,1)$ for $i=1,2$, so $2\succ_1 1$ while $1\succ_2 2$. Thus neither consumer locally dominates in both neighborhoods and both survive $\mathbb{P}$-a.s. (Figure 3(b)).

Mechanism. Extreme risk sensitivity makes the consumer who bears aggregate risk behave as if the worst state is almost certain, which is very inaccurate given $\pi_0=(1/2,1/2)$. Since the identity of the risk-bearer switches across neighborhoods, local dominance reverses,

preventing either consumer from driving out the other.

3. Path-dependent dominance. Let $\gamma_1 = 0$ and $\gamma_2 = \ln 2$. Then consumer 1 behaves as SEU, so $\nabla\mathcal{R}^1(\overline{V}^1) = \pi^1$. We obtain $1 \succ_1 2$ and (as above) $2 \succ_2 1$, so each consumer dominates with strictly positive probability, and exactly one survives $\mathbb{P}$-a.s. (Figure 3(c)).
Mechanism. Consumer 1's effective belief is allocation-invariant, so she is locally dominant when she is the risk-bearer (by the standard belief-accuracy logic). Consumer 2's effective belief is distorted only when she bears risk, and this distortion can improve accuracy in her own near-dominance neighborhood. Mutual local dominance therefore leads to endogenous path dependence: either consumer can end up dominating, but one dominates almost surely.

The preferences we assumed in Example 7 imply that the effective beliefs when some consumer dominates the market are probability measures, so a consumer with correct beliefs cannot be driven out of the market by consumers with incorrect beliefs. The following example shows that for some other preferences a consumer i with incorrect beliefs can actually drive another with correct beliefs out of the market almost surely if the mass i allocates to future states becomes larger than one.

Example 8 (Correct Consumers Vanish)**.** Let $I = \{1, 2\}$, $S = \{s_1, s_2\}$, and $\pi_0 = (1/2, 1/2)$. Beliefs are $\pi^1 = (3/4, 1/4)$ and $\pi^2 = \pi_0$. Preferences have CRRA certainty equivalent:

$$V_t^i(c) = u(c_t) + \beta\left[\mathbb{E}_{\pi^i}\left((V_{t+1}^i(c))^{1-\gamma^i} \mid \mathcal{F}_t\right)\right]^{\frac{1}{1-\gamma^i}}, \qquad u(0) > 0.$$

When consumer i consumes the entire endowment, effective beliefs satisfy

$$\nabla_s\mathcal{R}^j(\underline{V}^j) = \pi^j(s), \qquad \nabla_s\mathcal{R}^i(\overline{V}^i) = \left(\frac{\overline{V}_s^i}{\mathcal{R}^i(\overline{V}^i)}\right)^{-\gamma^i}\pi^i(s). \tag{10}$$

Unlike Example 7, these effective beliefs need not sum to one. For suitable parameter values,[18] one obtains $1 \succ_1 2$ and $1 \succ_2 2$. Hence, by Theorem 1(i), consumer 1 dominates $\mathbb{P}$-a.s., even though her subjective belief is incorrect.

Mechanism. When the sum of effective beliefs is greater than 1, consumer 1 behaves as if she were endogenously more patient. This mass distortion can compensate for incorrect odds and allow an incorrect-belief consumer to dominate.

Remark 2 (Three Channels of Distortion). These examples isolate three channels through which recursive preferences alter selection. First, effective beliefs can reweight states, improving or worsening subjective forecasts. Second, effective beliefs need not sum to one,

[18] For instance, take $\gamma^1 = \gamma^2 = 3/2$ and choose endowments so that $\sum_s \nabla_s\mathcal{R}^1(\overline{V}^1) > 1$; a concrete calibration is available upon request.

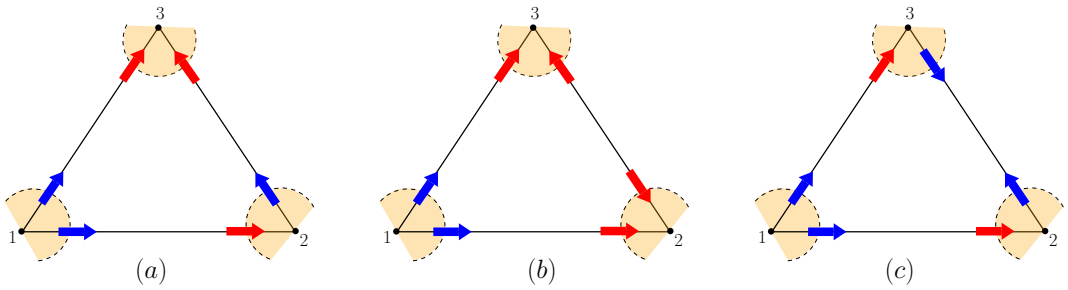

Figure 4: Multiple Consumers

which changes effective patience. Third, both distortions depend on who bears aggregate risk, so local dominance can reverse across near-dominance regions. These channels are absent under SEU, where effective beliefs and discount rates are allocation invariant.

6.5 Multiple-Consumer Economies

In this section, we extend the previous section characterization to general cases with an arbitrary number of consumers. We also argue that the survival patterns with multiple consumers cannot be considered as a simple extension of the two-consumer case.

Definition 6 (Local Contour Sets)**.** For each i, we define i's ***local upper contour set*** $U_i = \{j : j \succ_i i\}$ and i's ***local lower contour set*** $D_i = \{j : i \succ_i j\}$.

We impose the following assumption.

Assumption 8 (No-Indifference)**.** *For all i and $j \neq i$, either $j \in U_i$ or $j \in D_i$ holds.*

The local upper (lower) contour set of consumer i consists of the consumers who dominate i (are dominated by i) when i consumes the aggregate endowment. Assumption 8 stipulates that every consumer belongs to either the local upper contour set or the lower contour set of any other consumer, thereby ruling out indifference cases.

Theorem 2 (Complete Ranking of Upper Contour Sets)**.** *Suppose that Assumptions 1 to 8 hold, and that there exists a re-labeling of consumers $\{i_1, i_2, \ldots, i_n\}$ such that*

$$U_{i_1} \subset U_{i_2} \subset \ldots \subset U_{i_n}.$$

Let $I^* = \{i : U_i = \emptyset\}$. Then (i) all consumers not in I^* will vanish $\mathbb{P}$-almost surely, and (ii) every consumer in I^* will dominate the market with strictly positive $\mathbb{P}$-probability.

Theorem 2 states that if every consumer's local upper contour set can be ranked completely based on set inclusion, then survivors are the consumers whose local upper contour set is empty, indicating that they are locally undominated. Besides, each of these survivors will dominate the market with a strictly positive probability. Figure 4 provides some illustrative

examples. In Figure 4 (a), we have

$$U_1 = \{2, 3\}, U_2 = \{3\}, \text{and } U_3 = \emptyset,$$

resulting in a complete ranking: $U_3 \subset U_2 \subset U_1$. According to Theorem 2, the only survivor is $I^* = \{3\}$, meaning that consumer 3 almost surely dominates the market, while consumers 1 and 2 almost surely vanish. In Figure 4 (b), we have

$$U_1 = \{2, 3\}, U_2 = U_3 = \emptyset,$$

resulting in a set of survivors: $I^* = \{2, 3\}$. Hence, consumer 1 will vanish almost surely, and consumers 2 and 3 will dominate the market, each with strictly positive probability.

Intuition. To illustrate the argument, consider Figure 4(a), where $U_3 \subset U_2 \subset U_1$. Under irreducibility, L_t visits each one-consumes-all neighborhood infinitely often. When L_t enters consumer 1's neighborhood, $U_1 = \{2, 3\}$ implies that consumer 1 cannot dominate by Proposition 3, so the process must eventually exit this neighborhood. Once outside, either consumer 2 or 3 must hold a non-negligible share, and irreducibility implies that L_t subsequently enters the neighborhood of either consumer 2 or 3.

Applying the same reasoning to consumer 2, we conclude that L_t must eventually enter consumer 3's neighborhood with positive probability. Since $U_3 = \emptyset$, consumer 3 is locally undominated, and Proposition 3 implies that this neighborhood is absorbing with positive probability. Because it is visited infinitely often, eventual absorption occurs almost surely. Hence consumer 3 dominates.

There are situations in which Theorem 2 does not apply. Figure 4(c) provides such an example, where the upper contour sets cannot be completely ranked by set inclusion. In this case, Proposition 3 implies that no single consumer can dominate the market, so at least two consumers must survive. However, the proposition alone does not identify which consumers survive. The following theorem provides additional structure.

Theorem 3 (Incomplete Ranking of Upper Contour Sets)**.** *Suppose that Assumptions 1 to 8 hold. Let $G \subset I$ be any subset such that $\{U_i \cap G : i \in G\}$ can be completely ranked by set inclusion, and let $G^* = \{i \in G : U_i \cap G = \emptyset\}$. If $U_i \neq \emptyset$ for all $i \in G^*$, then $\mathbb{P}$-almost surely some consumer in $I \backslash G$ survives.*

Applying Theorem 3 to Figure 4(c), we can further conclude that both consumers 2 and 3 will almost surely survive. Taking $G = \{1, 3\}$, then

$$U_1 \cap G = \{3\}, \text{and } U_3 \cap G = \emptyset,$$

so they can be ranked by set inclusion, and we obtain $G^* = \{3\}$. Because $U_3 = \{2\} \neq \emptyset$,

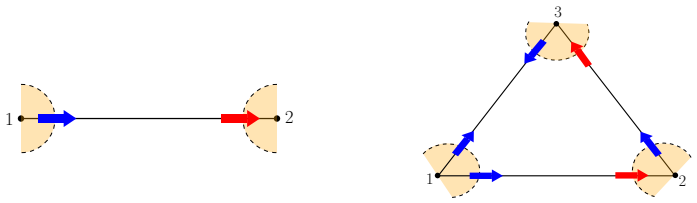

(a): consumer 1 vanishes a.s. (b): consumer 1 survives a.s.

Figure 5: Revival of an extinct consumer

Theorem 3 implies consumer 2 (i.e., someone in $I\backslash G$) survives almost surely. A symmetric argument with $G = \{1, 2\}$ implies that consumer 3 survives almost surely. Hence both consumers 2 and 3 survive.

Intuition. If all consumers in $I\backslash G$ were to vanish, the economy would reduce to one containing only G. By Theorem 2, some consumer in G^* would dominate. However, since each $i \in G^*$ is locally dominated by someone outside G, Proposition 3 rules out their dominance. Therefore, the event that all consumers in $I\backslash G$ vanish must have probability 0.

It is worth noting that the multiple-consumer case cannot be considered as a straightforward extension of the two-consumer case. The following example, using Theorem 3, illustrates that it is possible for a consumer who almost surely vanishes in a two-consumer economy to almost surely survive in a larger economy.

Example 9 (Revival of an Extinct Consumer)**.** Consider Example 7 with $\gamma_1 = \gamma_2 = \ln 2$.

- **Two-consumer economy.** With $I = \{1, 2\}$, we have $2 \succ_1 1$ and $2 \succ_2 1$, so consumer 1 *vanishes almost surely* (Figure 5(a)).
- **Three-consumer economy.** Add consumer 3 with $\pi^3 = (1/2, 1/2)$ and $\gamma_3 = \ln(7/3)$. The induced effective beliefs are

$$\nabla\mathcal{R}^i(\overline{V}^i) = \begin{cases} (0.2, 0.8) & i = 1, \\ (0.6, 0.4) & i = 2, \\ (0.3, 0.7) & i = 3, \end{cases} \qquad \nabla\mathcal{R}^i(\underline{V}^i) = \pi^i.$$

  The resulting dominance relations are depicted in Figure 5(b).

  Taking $G = \{2, 3\}$, we have $U_3 \cap G = \emptyset \subset U_2 \cap G = \{3\}$, so $G^* = \{3\}$. Since $U_3 \neq \emptyset$, consumer 1 *survives almost surely* by Theorem 3.

This example shows that the analysis of survival of a two-consumer economy with general preferences in the literature (e.g., Guerdjikova and Sciubba, 2015; Borovička, 2020) is not sufficient to understand the prospects for survival in general. With general preferences, the local survival indices can be allocation dependent, leading to a reversal of the indices' order. In our example above, the index of consumer 2 is larger than that of consumer 1

when either of them dominates the market (so consumer 1 would vanish in a two-consumer economy with $I = \{1, 2\}$) and the local survival index of consumer 3 is larger than that of consumer 2 (so consumer 2 would vanish in a two-consumer economy with $I = \{2, 3\}$). However, consumer 1's local survival index is the largest when consumer 3 dominates the market, which prevents consumer 1 from vanishing in the three-consumer economy.[19]

Example 10 (Heterogeneity of Preferences in EZ economies)**.** This example shows that preference heterogeneity overturns the selection patterns suggested by the homogeneous-preference benchmark. We consider the case of unit elasticity of intertemporal substitution. In the Supplementary Material, Section S.7, we derive the local survival index for $\gamma^j \neq 1$:[20]

$$\mathcal{S}_i^j = \begin{cases} \mathbb{E}_{\pi_0}\left(\log \dfrac{\pi^i(s)\left(\bar{V}_s^i\right)^{1-\gamma^i}}{\sum_{s'} \pi^i(s')\left(\bar{V}_{s'}^i\right)^{1-\gamma^i}}\right) + \log \beta^i & \text{if } j = i \\ \mathbb{E}_{\pi_0}\left(\log \dfrac{\left(\frac{\pi^j(s)}{\pi^i(s)}\right)^{\frac{1}{\gamma^j}} \left(\frac{\phi^j(s|\bar{p}^i)}{\phi^i(s|\bar{p}^i)}\right)^{\frac{1-\gamma^j}{\gamma^j}} \pi^i(s)\left(\bar{V}_s^i\right)^{\gamma^i \frac{1-\gamma^j}{\gamma^j}}}{\sum_{s'}\left(\frac{\pi^j(s')}{\pi^i(s')}\right)^{\frac{1}{\gamma^j}} \left(\frac{\phi^j(s'|\bar{p}^i)}{\phi^i(s'|\bar{p}^i)}\right)^{\frac{1-\gamma^j}{\gamma^j}} \pi^i(s')\left(\bar{V}_{s'}^i\right)^{\gamma^i \frac{1-\gamma^j}{\gamma^j}}}\right) + \log \beta^j & \text{if } j \neq i \end{cases}$$

The key property is that each agent j's index is the negative of the relative entropy of a probability distribution (see footnote 9); in particular, the index of both EZ agent is no larger than that of an EU agent with correct beliefs and identical discount rate. Figure 6 shows the indices of an EU agent with correct beliefs and agents 1 and 2 as functions of their common risk-aversion coefficient γ^j when agent 1 consumes all (left) and agent 2 consumes all (right).

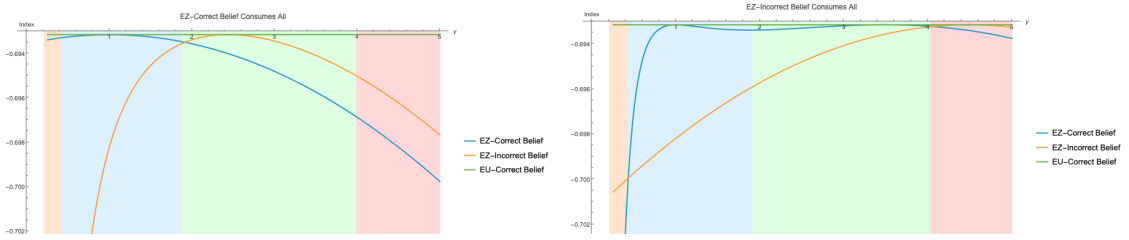

Figure 6: Local survival indices of agents 1 and 2.

- In the blue region the MSH holds under homogeneous preferences but it fails under heterogeneous ones: for any $\gamma^j \neq 1$, there exists an open set of incorrect beliefs for which the EU agent's index exceeds that of the correct-belief EZ agent when the latter consumes

[19] Note that this reversal of local survival indices cannot occur when all consumers have SEU preferences. Therefore, under SEU, a consumer who vanishes in a two-consumer economy also vanishes in a larger one.

[20] Where $\bar{p}^i(\sigma_t, s_{t+1}) = \beta^i \pi^i(s_{t+1}) \left(\frac{e(s_{t+1})}{e(s_t)}\right)^{-\rho^i} \left(\frac{\bar{V}_{s_{t+1}}^i}{\mathcal{R}^i(\bar{V}^i)}\right)^{\rho^i - \gamma^i} \bar{p}^i(\sigma_t)$ are the prices in a representative agent economy with only agent i and $\phi^j(\sigma_t, s|\bar{p}^i)$ is agent j's marginal utility of wealth at (σ_t, s) and prices $\bar{p}^i$.

everything, so the correct-belief EZ agent dominates with zero probability.

- In the red region the MSH fails under homogeneous EZ preferences but holds under heterogeneous ones: for any $\gamma^j \neq 1$—as Figure 6 shows—the correct-belief EU agent (agent 3) has a higher index than both EZ agents when they consume everything, and—as Figure 7 shows—the highest index when agent 3 consumes everything.

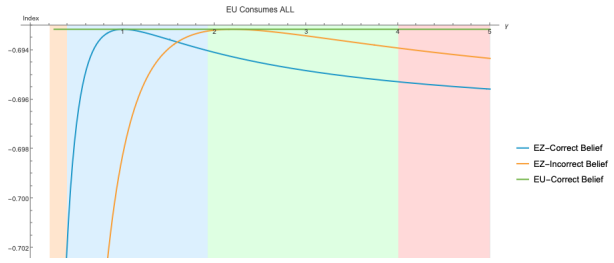

Figure 7: Local survival indices of agents 1, 2 and 3.

The conclusion then follows from Theorem 2, since $U_2 = \{1, 3\}$, $U_1 = \{3\}$ and $U_3 = \emptyset$.

7 Robustness of the Market Selection Hypothesis

This section studies whether the classical market selection hypothesis (MSH) is robust once we enlarge the preference space from discounted SEU to a more general preference space. We ask two robustness questions:

- **Global robustness:** Is there a preference such that a consumer endowed with it always survives with positive probability *regardless of the other consumers' preferences*?
- **Local robustness:** Is there a preference such that a consumer endowed with it always survives with positive probability against *arbitrarily small perturbations* of the other consumers' preferences?

To address these questions, we must specify the class of preferences over which robustness is evaluated. The class should be sufficiently general to encompass SEU preferences as a special case, yet structured enough to allow meaningful comparisons across agents.

7.1 Assumptions

We restrict attention to preferences with a common time aggregator that is separable and linear in the second argument. This setup allows heterogeneity only in the certainty equivalent operator. Formally, we impose the following assumptions.

Assumption 9 (Aggregate uncertainty)**.** *For all $t \in \mathbb{T}$ and $\sigma \in \Sigma$, there exist $s, s' \in S$ such that $e(\sigma_t, s) \neq e(\sigma_t, s')$.*

Assumption 10 (Common separable linear class)**.** *For all $i \in I$, (i) $F^i(c, y) = u(c) + \beta y$ for some $\beta \in (0, 1)$ and (ii) $\nabla\mathcal{R}^i$ is well-defined at local utilities.*

Assumption 9 rules out riskless economies. Assumption 10 imposes a common discount factor and a separable aggregator and also requires the certainty equivalent aggregator to be differentiable at the boundary. The first restriction is natural for two reasons.

First, to make the robustness discussion non-trivial, consumers must discount the future in the same manner (i.e., they have the same F^i) and the effective discount rate is independent of future utility (i.e. F^i is linear in certainty equivalent). Otherwise, we can always make a consumer more likely to survive by increasing the discount factor.[21] Second, the separable aggregator nests the standard discounted SEU framework, which allows us to assess directly whether the classical market selection results are robust to richer attitudes toward uncertainty.[22]

Under Assumption 10, differences in long-run outcomes arise solely from the certainty equivalent operator $\mathcal{R}$. Throughout this section, we thus refer to $\mathcal{R}$ as *preferences*. Let $\mathscr{R}$ denote the class of certainty equivalent operators satisfying Assumptions 2, 4, 5, and 10(ii).

7.2 A Technical Device: ε-Perfections

Many of our characterizations rely on Assumption 6, which may fail in some economies due to particular features of the state space or endowment process. To avoid excluding such cases, we introduce a perturbation device.

For any economy and any $\varepsilon > 0$, one can construct a nearby economy (possibly with a larger state-space) that is ε-close in primitives and satisfies Assumption 6. We refer to such economies as *ε-perfections*. The formalization is in Supplementary Material, Section S.4.

Fact 1. *For every economy satisfying Assumptions 1 to 5 and every $\varepsilon > 0$, there exists an ε-perfection of that economy.*

This fact ensures that irreducibility can be attained arbitrarily close to any economy.[23] This fact therefore enables us to apply this paper's characterizations in a universal manner.[24]

[21] The essential component of Assumption 10 is linearity and the common discount rate β. The common u is imposed mainly for notational clarity, so that heterogeneity across consumers is captured entirely by $\mathcal{R}$. Allowing heterogeneous u would not affect the qualitative results.

[22] The separable structure also nests recursive smooth ambiguity preferences, monotonic transformations of homothetic Epstein–Zin preferences and many others.

[23] Intuitively, failures of irreducibility typically arise because some transitions happen with zero probability, which partitions the state space into disconnected communicating classes. An ε-perfection introduces an arbitrarily small "tremble"—for example, by slightly perturbing the transition law (or by adding auxiliary states with total probability mass below ε) so that every state becomes reachable with positive probability.

[24] An alternative approach would be to impose Assumption 6 throughout and study the implied global dynamics via local dominance patterns. We instead use ε-perfections to avoid restricting attention to economies

It can be applied directly to any given economy or to a small perturbation of it. Moreover, the perturbed economy can approximate the original preferences arbitrarily well, enabling us to explore the robustness of survival for a wide range of preferences.

7.3 Global Robustness

We begin with a global notion of robustness: robustness against *arbitrary* profiles of other consumers' preferences in $\mathscr{R}$.

Definition 7 (Robust dominator)**.** Preference $\mathcal{R}^i \in \mathscr{R}$ ***robustly dominates*** if, for every profile $\{\mathcal{R}^j\}_{j\neq i} \in \mathscr{R}^{I-1}$ with $\mathcal{R}^j \neq \mathcal{R}^i$ for all $j \neq i$, consumer i dominates $\mathbb{P}$-a.s. in all ε-perfected economies with small ε.[25]

Definition 8 (Robust survivor)**.** Preference $\mathcal{R}^i \in \mathscr{R}$ ***robustly survives*** if, for every profile $\{\mathcal{R}^j\}_{j\neq i} \in \mathscr{R}^{I-1}$ with $\mathcal{R}^j \neq \mathcal{R}^i$ for all $j \neq i$, consumer i survives with $\mathbb{P}$-strictly positive probability in all ε-perfected economies with small ε.

Robust dominance represents the strongest possible form of selection: a consumer with a robustly dominating preference $\mathcal{R}^i$ dominates almost surely for every configuration of other consumers' preferences, and in every sufficiently small ε-perfection of the economy. Robust survival is weaker in the sense that it only requires strictly positive probability of survival, but it also requires survival under all possible configurations of preferences. If there exist multiple robustly-surviving preferences associated with heterogeneous beliefs, then no preference or belief always dominates the market, and hence the market selection hypothesis is not robust.

We then present the following result.

Theorem 4 (Global Robustness)**.** *Under Assumptions 1–5 and 8–10:*

(i) no preference $\mathcal{R}^i \in \mathscr{R}$ robustly dominates;

(ii) $\mathcal{R}^i$ robustly survives if and only if $i \succ_i r$, where r denotes a consumer with correct SEU preferences.

Theorem 4(i) implies that the MSH is globally non-robust in $\mathscr{R}$—no preference guarantees dominance against all others. Theorem 4(ii) provides a simple characterization of robust survivors: A preference relation robustly survives *if and only if* a consumer with such a preference relation locally dominates a consumer with correct SEU preferences.[26] Based

that satisfy irreducibility exactly.

[25] Formally, there exists $\bar{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon})$ and every ε-perfection of $\mathcal{E}_S$, consumer i dominates $\mathbb{P}$-a.s.

[26] The if and only if condition relies on Assumption 8, which precludes tie cases where $i \sim_i r$. Without this assumption, we will have an almost if and only if condition that $\mathcal{R}^i$ robustly survives if $i \succ_i r$ and only if $i \succeq_i r$.

on the characterization, we can find many preferences—with heterogeneous beliefs—that always survive with strictly positive probability. Below is an example.

Example 11 (Robust survivors)**.** Consider $S = \{s_1, s_2\}$ and $\pi_0 = (1/2, 1/2)$, and suppose consumers have the following class of preferences

$$V_{\sigma_t} = u(c_t) + \beta \cdot \phi^{-1}\left(\mathbb{E}_\pi \phi(V_{\sigma_t,s})\right), \qquad \phi(x) = \tfrac{1-\gamma}{\gamma}\left(\tfrac{x}{1-\gamma} + m\right)^\gamma.$$

For simplicity, take $m > 0$ and $\beta \approx 0$.[27] The effective beliefs satisfy approximately

$$\nabla_s \mathcal{R}(\overline{V}) \approx \frac{u_s^{\gamma-1}\pi(s)}{\left(\sum_{s\in S} u_s^\gamma \pi(s)\right)^{1-1/\gamma}}, \qquad \nabla_s \mathcal{R}(\underline{V}) = \pi(s).$$

By Theorem 4, robust survival holds if

$$\mathbb{E}\log\left[\frac{u_s^{\gamma-1}\pi(s)}{\left(\sum_{s\in S} u_s^\gamma \pi(s)\right)^{1-1/\gamma}}\right] > \mathbb{E}\log \pi_0(s). \tag{11}$$

Figure 8 plots the robust-survivor set when $u_1 = 1$, $u_2 = \exp(k) > 1$, and the vertical axis is $\pi(s_1)$. Robust survivors exist for $\gamma \in (0, 1)$. As aggregate uncertainty increases ($k \to +\infty$) the set of robustly surviving beliefs—i.e., beliefs associated with some robustly surviving preference—expands to the entire interval $[0, 1]$. In other words, for *any* nondegenerate belief, there exists some economy such that some preference with this belief robustly survives.

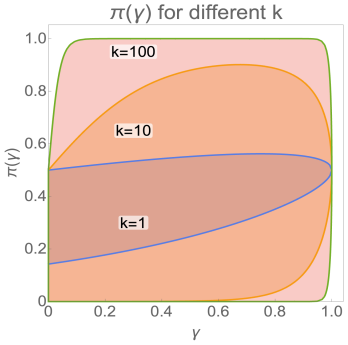

Figure 8: Robust Survivors

Intuition of Theorem 4. The key step is local expected utility at extinction. When an opponent j is almost extinct, her continuation utility is close to a constant vector $\underline{V}^j\mathbf{1}$. Since $\mathcal{R}^j(k\mathbf{1}) = k$, differentiability gives $\sum_{s\in S} \nabla_s \mathcal{R}^j(k\mathbf{1}) = 1$. Thus, $\nabla\mathcal{R}^j(\underline{V}^j\mathbf{1})$ is a probability vector. Near extinction, every admissible opponent is locally summarized by an SEU-type effective belief. Among SEU preferences, the correct-belief benchmark is locally the most difficult to beat. Therefore, if consumer i locally dominates the correct-belief SEU type, she locally dominates any preference in $\mathscr{R}$. This guarantees strictly positive probability of survival for any configuration of opponents, and hence $\mathcal{R}^i$ robustly survives. Conversely, if i

[27] The parameter $m > 0$ keeps the certainty equivalent differentiable at extinction; the approximation $\beta \approx 0$ makes $\overline{V}_s \approx u_s$ and yields a transparent closed-form condition.

is locally dominated by the correct-belief SEU consumer, one can construct a preference $\mathcal{R}^j \in \mathscr{R}$ such that i vanishes almost surely against j; that is, $\mathcal{R}^i$ does not robustly survive.

7.3.1 Robust Survival and Uncertainty Attitudes

Theorem 4(ii) shows that a preference robustly survives if and only if it locally dominates the SEU preference with correct belief. Although this condition tightly characterizes robust survivors, it is not totally obvious what it means in terms of properties of preferences and its implications for consumers' decisions. In Supplementary Material, Section S.5, we show that robust survival is highly related to the *decreasing absolute uncertainty aversion* (DAUA) property.[28] Specifically, we show that with the DAUA property, the total mass of i's effective belief is greater than one, making i effectively more patient than an SEU agent, thus increasing her survival prospects. On the other hand, if $\mathcal{R}^i$ exhibits *increasing absolute uncertainty aversion* (IAUA), the total mass is less than one, so i is effectively less patient. Formally, we prove one necessary condition for robust survival—if $\mathcal{R}^i$ is a robust survivor, then it must *not* exhibit IAUA. Hence, DAUA serves as a loose necessary condition for robust survival. In Example 11, robust survival occurs when $\gamma \in (0,1)$, which is the region where preferences exhibit DAUA.[29]

7.4 Local Robustness

We now examine the local robustness of the MSH. To formalize local perturbations of preferences, let $\mathcal{V}$ be a compact set containing all possible continuation-utility values.[30] Define

$$D_0(\mathcal{R},\mathcal{R}') = \max_{V\in\mathcal{V}} \left|\mathcal{R}(V) - \mathcal{R}'(V)\right|, \qquad D_1(\mathcal{R},\mathcal{R}') = \max_{V\in\mathcal{V},\ s\in S} \left|\nabla_s\mathcal{R}(V) - \nabla_s\mathcal{R}'(V)\right|,$$

and let $D(\mathcal{R},\mathcal{R}') = \max\{D_0, D_1\}$. We also define $B^\circ_\delta(\mathcal{R})$ as the set of preferences $\mathcal{R}' \in \mathscr{R}$ such that $0 < D(\mathcal{R}',\mathcal{R}) < \delta$.

Definition 9 (Locally robust dominator)**.** Preference $\mathcal{R}^i \in \mathscr{R}$ ***locally robustly dominates*** if there exists $\delta > 0$ such that for all $\{\mathcal{R}^j\}_{j\neq i} \in B^\circ_\delta\left(\mathcal{R}^i\right)^{I-1}$, consumer i dominates $\mathbb{P}$-a.s. in all ε-perfected economies with small ε.

Definition 10 (Locally robust survivor)**.** Preference $\mathcal{R}^i \in \mathscr{R}$ ***locally robustly survives*** if

[28] Formally, $\mathcal{R}^i$ displays DAUA if $k - \mathcal{R}^i(V + k\cdot\mathbf{1})$ is decreasing in k. Equivalently, $\mathcal{R}^i$ is constant superadditive. For risk certainty equivalents, this property specializes to decreasing absolute risk aversion; for ambiguity certainty equivalents, it specializes to decreasing absolute ambiguity aversion (Fabbri et al., 2024).

[29] Although DAUA is necessary in this sense, it alone is not sufficient to guarantee robust survival. In general, robust survival can rely intricately on other properties of the economy. As in Example 11, robust survivors only exist for particular beliefs and endowments when $\gamma \in (0,1)$.

[30] For instance one may take $\mathcal{V} = [0,\bar{v}]$ with $\bar{v} = u(\bar{e})/(1-\beta)$, where $\bar{e}$ is the maximal aggregate endowment.

there exists $\delta > 0$ such that for all $\{\mathcal{R}^j\}_{j\neq i} \in B^\circ_\delta\left(\mathcal{R}^i\right)^{I-1}$, consumer i survives with $\mathbb{P}$-strictly positive probability in all ε-perfected economies with small ε.

Local robust dominance and survival are much weaker than the global versions. They require dominance and survival only in the presence of sufficiently close preferences. To characterize local robust survival, we introduce the following concepts.

Definition 11 (Effective uncertainty neutrality)**.** Preference $\mathcal{R}^i \in \mathscr{R}$ is ***effectively uncertainty neutral*** if

$$\mathbb{E}\log\left[\nabla_s \mathcal{R}^i(\underline{V}^i)\right] = \mathbb{E}\log\left[\nabla_s \mathcal{R}^i(\overline{V}^i)\right], \tag{12}$$

and is ***effectively uncertainty favored*** (resp. ***unfavored***) if the left-hand side of (12) is smaller (resp. larger) than the right-hand side.

In words, $\mathcal{R}^i$ is effectively uncertainty neutral if the entropy of its effective belief, captured by the negative of $\mathbb{E}\log\nabla_s\mathcal{R}^i$, remains unchanged when i consumes all endowment (hence bears the aggregate uncertainty) and when i consumes nothing (hence faces no uncertainty). Correspondingly, it is called effectively uncertainty favored or unfavored if bearing all aggregate uncertainty decreases or increases the entropy, respectively.

Theorem 5 (Local Robustness)**.** *Under Assumptions 1–5 and 8–10:*

(i) no preference $\mathcal{R}^i \in \mathscr{R}$ locally robustly dominates;

(ii) if, in addition, $|I| = 2$, then $\mathcal{R}^i \in \mathscr{R}$ locally robustly survives if and only if it is not effectively uncertainty neutral.[31]

Recall that Theorem 4 shows that we do not have a global robust dominator, but one would conjecture that it is still likely that some preferences can locally dominate. Theorem 5(i) further establishes that such locally dominating preferences do not even exist. Correspondingly, there exist multiple preferences that can locally survive. Theorem 5(ii) characterizes locally robust survivors in a two-consumer economy. It shows that the market selection result with SEU preferences even represents a *knife-edge case* in a local sense—any preference that is not effectively SEU (i.e., not effectively uncertainty neutral) is a locally robust survivor.

Intuition. Suppose $I = \{i, j\}$. If $\mathcal{R}^i$ is effectively uncertainty favored, then the survival index of i when she consumes everything, $\mathcal{S}^i_i$, is higher than that when she consumes nothing,

[31] The "only if" part does not require $|I| = 2$ and holds for an arbitrary number of consumers. The "if" part is stated for $|I| = 2$ because our multi-consumer characterizations are less informative in the effectively-unfavored case; see Corollary 1.

$\mathcal{S}_j^i$. When the preferences of i and j are very similar, we have $\mathcal{S}_i^i \approx \mathcal{S}_j^j$ and $\mathcal{S}_j^i \approx \mathcal{S}_i^j$. Hence,

$$\mathcal{S}_i^i > \mathcal{S}_j^i \Longrightarrow \mathcal{S}_j^j > \mathcal{S}_j^i \text{ and } \mathcal{S}_i^i > \mathcal{S}_i^j.$$

That is, each consumer has a higher survival index when she almost dominates the market. This corresponds to the pattern in Figure 2(a), so both consumers dominate the market with positive probability. Similarly, when $\mathcal{R}^i$ is effectively uncertainty unfavored, both consumers coexist with probability 1 as in Figure 2(d).

Theorem 5(ii) characterizes only the case when $|I| = 2$. For $|I| > 2$, our theorems cannot pin down which specific agents survive without more structure, but they do imply that at least two agents coexist, leading to the following weaker notion.

Definition 12 (Locally robust co-survivor)**.** Preference $\mathcal{R}^i \in \mathscr{R}$ ***locally robustly co-survives*** if there exists $\delta > 0$ such that for all $\{\mathcal{R}^j\}_{j\neq i} \in B_\delta^\circ\left(\mathcal{R}^i\right)^{I-1}$, there exist *at least two* consumers who survive with strictly positive $\mathbb{P}$-probability in all ε-perfected economies with small ε.

Corollary 1. *Under Assumptions 1–5 and 8–10, $\mathcal{R}^i \in \mathscr{R}$ locally robustly co-survives if it is not effectively uncertainty neutral.*

Corollary 1 implies that local coexistence is dense in $\mathscr{R}$: in *any* neighborhood around any preference that is not effectively uncertainty neutral, at least two consumers survive with strictly positive probability. The coexistence pattern further depends on whether the preference is effectively uncertainty favored or unfavored: in the former multiple consumers can dominate the market, in the latter they survive on almost all paths.

7.5 Beyond Local Expected Utility at Extinction

The robustness analysis above relies on the local expected utility property: as a consumer vanishes, her continuation utility becomes nearly constant, and a consumer behaves as if she were an SEU agent. This is a standard regularity property in the non-expected-utility literature, closely related to the local expected-utility approach of Machina (1982) and to the distinction between smooth and first-order responses to small risks in Segal and Spivak (1990). It is satisfied by many smooth recursive specifications.

There are two natural ways in which this property may fail. First, the certainty equivalent may be non-smooth at deterministic continuations. Examples include preferences with kinks, such as max-min expected utility or rank-dependent and Choquet expected utility. For such preferences, we expect that one could either work with smooth approximations to which our theorem applies, or extend the local-index approach by replacing the gradient of the certainty equivalent with a set of local supporting beliefs. In Section S.9 of Beker and Chen

(2026), we provide survival analysis for two preferences with kinks.

Second, preferences may be smooth at deterministic continuations in the interior of the utility domain but fail the relevant boundary condition at extinction. One important example is homothetic Epstein–Zin preferences with a CRRA certainty equivalent. These preferences are smooth at interior deterministic continuations, but the continuation value at extinction is zero, where the CRRA certainty equivalent is not differentiable. For this class of preferences, the local survival index can still be computed, but it must be defined as the limit along the equilibrium-implied direction. Indeed, in the Supplementary Material, Section S.7, we show that for any $\rho^i > 0$ and $\gamma^i \neq 1$, the survival index is given by

$$\mathcal{S}_i^i = \left(1 - \frac{1-\rho^i}{1-\gamma^i}\right) \times \mathbb{E} \log \left(\frac{\pi^i(s) \left(\bar{V}_s^i\right)^{1-\gamma^i}}{\mathbb{E}_{\pi^i}\left[\left(\bar{V}_{s'}^i\right)^{1-\gamma^i}\right]} \right) + \frac{1-\rho^i}{1-\gamma^i} \times \mathbb{E} \log \pi^i(s) + \log \beta$$

and $0 < \frac{1-\rho^i}{1-\gamma^i} < 1$ implies that $r \succ_i i$. Additionally, $\mathcal{S}_r^i < \mathbb{E} \log \pi_0(s) + \log \beta$ and so $r \succ_r i$ if $0 < \frac{1-\rho^i}{1-\gamma^i} < 1$. It follows from Proposition 3 that EZ agents who display PLRU cannot robustly survive when $\rho^i < 1$, while those who display PERU cannot robustly survive when $\rho^i > 1$.

8 Related Literature

The MSH has been extensively analyzed for the case of consumers with SEU preferences. Sandroni (2000) and Blume and Easley (2006) were the first to analyze formally the MSH in a general equilibrium model with complete markets where agents make optimal intertemporal plans. They consider the class of SEU preferences and show that among consumers with the same discount rate and bounded endowments, survival depends only on the accuracy of beliefs as measured by their relative entropy. Yan (2008), Cvitanić and Malamud (2011) and Kogan et al. (2017) extend the analysis to the case of unbounded endowments.[32] These papers, however, left open the question of whether the MSH is an intrinsic property of optimal plans made by consumers with correct beliefs or whether it is due to the ancillary restriction to SEU preferences.

Some recent works on the MSH show that consumption dynamics can actually be quite different with non-expected utility preferences.[33] Those papers mostly focus on market

[32] Some other setups have been studied, for example incomplete markets (Blume and Easley, 2006; Beker and Chattopadhyay, 2010; Coury and Sciubba, 2012; Cao, 2017; Cogley et al., 2014), differential asset constraints (Guerdjikova and Quiggin, 2019), continuum of agents (Massari, 2019) and asymmetric information (Mailath and Sandroni, 2003). Cogley and Sargent (2009) study the quantitative effect of the MSH on asset prices.

[33] Condie (2008), Da Silva (2011) and Easley and Yang (2015) consider max-min, variational and loss

selection across only one dimension of heterogeneity: (i) preference heterogeneity while holding beliefs fixed, or (ii) belief heterogeneity while holding preferences fixed. Guerdjikova and Sciubba (2015) focus on *market selection across preferences* in the presence of recursive smooth ambiguity preferences. Their setup is closest to this paper in that they also consider a discrete-time, stochastic exchange economy with complete markets. They also study the consumption dynamics implied by the consumers' first-order conditions using effective beliefs and discount rates.[34] They analyze the case of two consumers with correct beliefs and heterogeneous preferences: one with SEU, the other with smooth ambiguity preferences. They show that if absolute ambiguity aversion decreases fast enough, then the ambiguity averse consumer can drive the SEU consumer out of the market. Their finding leads us to conjecture that the selection results within the SEU class of preferences may not be robust to the introduction of other preferences. Our analysis confirms this conjecture.

Dindo (2019) and Borovička (2020), instead, mainly focus on *market selection across beliefs* as their consumers have homothetic Epstein–Zin (EZ) preferences but heterogeneous beliefs.[35] Technically, both papers employ a local approach to study consumption dynamics, which is in a similar spirit to this paper's approach. Both papers find that consumption dynamics can exhibit a variety of patterns absent in the discounted expected-utility framework, and specifically, consumers with incorrect beliefs can survive in the long run. In particular, Dindo (2019) considers the case of multiple consumers but within the special class of preferences that contain either homothetic EZ preferences with unit elasticity of intertemporal substitution or SEU with constant relative risk aversion; Borovička (2020) considers the case of two consumers with identical arbitrary homothetic EZ preferences.[36] Although a framework to analyze the MSH must allow for rich belief heterogeneity, it must simultaneously allow for rich preference heterogeneity, as the MSH is a preference-neutral statement. The previous literature predominantly focuses on market selection along a single dimension of heterogeneity and within a single class of preferences. The general recursive preferences considered here, by contrast, nest expected utility and the common time-consistent non-expected utility models and allow heterogeneity to vary on both dimensions at once, enabling us to examine the validity of the market selection hypothesis when preferences and beliefs are jointly heterogeneous.

aversion preferences, respectively, but they fail to identify a persistent impact on market outcomes.

[34] Their definitions of effective beliefs and discount rates are different from ours. They normalize effective beliefs to ensure that they add up to 1 and include the normalizing term in the effective discount rates.

[35] Dindo (2019) provides an example where one consumer has SEU with CRRA (although CRRA preferences can also be seen as a particular case of the homothetic EZ class).

[36] Dindo (2019) studies competitive-equilibrium dynamics in a discrete-time model, whereas Borovička (2020) studies Pareto optimal allocations in a continuous-time model.

9 Conclusion

We provide a unified framework for the dynamics of competitive-equilibrium consumption when markets are complete and consumers' preferences belong to a general class of recursive preferences, nesting both expected utility and many common non-expected-utility preferences, and we use it to study the robustness of the market selection hypothesis.

Sandroni (2000) and Blume and Easley (2006) show that among SEU consumers those with correct beliefs dominate the market, but leave open whether this reflects optimal plans under correct beliefs or the restriction to SEU. Enlarging the preference space to a rich recursive class, we find it is the latter. Globally, no preference robustly dominates, and a preference robustly survives if and only if—tie cases aside—it locally dominates a correct-belief SEU consumer, so that many preferences, associated with heterogeneous and even incorrect beliefs, robustly survive. Locally, no preference robustly dominates either, and any preference that is not effectively uncertainty neutral is a locally robust survivor, making the SEU selection result a knife-edge case. In sharp contrast to SEU, the robust long-run outcome in a sufficiently rich preference space is the survival of multiple consumers with heterogeneous beliefs.

These results expose an inherent tension between the efficiency of competitive-equilibrium allocations and rational expectations, and give a friction-free foundation for the behavioral-finance idea that incorrect beliefs can have persistent effects on asset prices.

We restrict attention to deviations from expected utility that preserve time consistency; the time-inconsistent case is left for future work. Likewise, the local approach developed here could be used to study consumption dynamics in incomplete-market economies.

Bibliography

Alchian, A. (1950) "Uncertainty, Evolution and Economic Theory," *Journal of Political Economy*, 58 (1), 211–221. 2

Aliprantis, C.D. and K.C. Border (1999) *Infinite Dimensional Analysis: A Hitchhiker's Guide*: Springer. 38

Beker, P. and S. K. Chattopadhyay (2010) "Consumption Dynamics in General Equilibrium: A Characterization When Markets Are Incomplete," *Journal of Economic Theory*, 145, 2133–2185. 33

Beker, Pablo F. and Jaden Yang Chen (2026) "Supplementary Material for "If You're NOT So Smart, Why Are You Rich? Robust Market Selection with General Recursive Preferences"," Available at SSRN, June, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7005738. 32

Bewley, T. (1969) "A Theorem on the Existence of Competitive Equilibria in a Market With a Finite Number of Agents and Whose Commodity Space is L_∞," in Khan, M. A. and N.C.Yannelis eds. *Equilibrium Theory in Infinite Dimensional Spaces*: Springer. 4

Blume, L. and D. Easley (1992) "Evolution and Market Behavior," *Journal of Economic Theory*, 58, 9–40. 2

——— (2006) "If You're So Smart, Why Aren't You Rich? Belief Selection in Complete and Incomplete Markets," *Econometrica*, 74 (4), 929–966. 2, 4, 6, 9, 12, 13, 20, 33, 35

——— (2009) "The Market Organism: Long Run Survival in Markets with Heterogeneous Traders," *Journal of Economic Dynamics and Control*, 33 (5), 1023–1035. 13

Borovička, J. (2020) "Survival and Long-run Dynamics With Heterogeneous Beliefs Under Recursive Preferences," *Journal of Political Economy*, 128 (1), 206–251. 2, 3, 4, 6, 8, 10, 24, 34

Cao, D. (2017) "Speculation and Financial Wealth Distribution Under Belief Heterogeneity," *Economic Journal*, 128, 2258–81. 33

Cerreia-Vioglio, S., D. Dillenberger, and P. Ortoleva (2015) "Cautious Expected Utility and the Certainty Effect," *Econometrica*, 83, 693–728. 9

Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci, and Luigi Montrucchio (2011) "Uncertainty averse preferences," *Journal of Economic Theory*, 146 (4), 1275–1330. 9

Cogley, T. and T. J. Sargent (2009) "Diverse Beliefs, Survival and the Market Price of Risk," *The Economic Journal*, 119 (March), 354–376. 33

Cogley, T., T. J. Sargent, and V. Tsyrennikov (2014) "Wealth Dynamics in a Bond Economy With Heterogeneous Beliefs," *The Economic Journal*, 124 (575), 1–30. 33

Condie, S. (2008) "Living with Ambiguity: Prices and Survival When Investors Have Heterogeneous Preferences for Ambiguity," *Economic Theory*, 36, 81–108. 2, 10, 33

Cootner, P. (1964) *The Random Character of Stock Market Prices*: MIT Press. 2

Coury, T. and E. Sciubba (2012) "Belief Heterogeneity and Survival in Incomplete Markets," *Economic Theory*, 49 (1), 37–58. 33

Cvitanić, J. and S. Malamud (2011) "Price Impact and Portfolio Impact," *Journal of Financial Economics*, 100 (1), 201–25. 33

Da Silva, C. (2011) "On Asymptotic Behavior of Economies With Complete Markets: The Role of Ambiguity Aversion," *Instituto Nacional de Matematica Pura e Aplicada, Rio de Janeiro, Brasil, mimeo*. 33

Dindo, P. (2019) "Survival in Speculative Markets," *Journal of Economic Theory*, 181, 1–43. 2, 3, 4, 5, 8, 10, 34

Easley, D. and L. Yang (2015) "Loss Aversion, Survival and Asset Prices," *Journal of Economic Theory*, 160, 494–516. 33

Epstein, L. G. and S. E. Zin (1989) "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57 (4), 937–969. 10

Fabbri, Francesco, Giulio Principi, and Lorenzo Stanca (2024) "Absolute and relative ambiguity attitudes," *arXiv preprint arXiv:2406.01343*. 30

Fama, E. (1965) "The Behavior of Stock Market Prices," *Journal of Business*, 38, 34–105. 2

Friedman, M. (1953) *Essays in Positive Economics*: University of Chicago Press. 2

Guerdjikova, A. and J. Quiggin (2019) "Market Selection with Differential Financial Constraints," *Econometrica*, 87 (5), 1693–1762. 33

Guerdjikova, A. and E. Sciubba (2015) "Survival With Ambiguity," *Journal of Economic Theory*, 155, 50–94. 2, 8, 10, 14, 24, 34, 43, 45

Hansen, Lars Peter and Thomas J Sargent (1995) "Discounted linear exponential quadratic gaussian control," *IEEE Transactions on Automatic control*, 40 (5), 968–971. 19

Klibanoff, P., M. Marinacci, and S. Mukerji (2009) "Recursive Smooth Ambiguity Preferences," *Journal of Economic Theory*, 144 (3), 930–976. 10

Kogan, L., S. A. Ross, J. Wang, and M. M. Westerfield (2017) "Market Selection," *Journal of Economic Theory*, 168, 209–236. 33

Krishna, V. and P. Sadowski (2014) "Dynamic Preferences for Flexibility," *Econometrica*, 82 (2), 655–703. 9

Machina, Mark J. (1982) ""Expected Utility" Analysis without the Independence Axiom," *Econometrica*, 50 (2), 277–323. 7, 32

Mailath, G. and A. Sandroni (2003) "Market Selection and Asymmetric Information," *Review Economic Studies*, 70 (2), 343–68. 33

Marinacci, M. and L. Montrucchio (2010) "Unique Solutions for Stochastic Recursive Utilities," *Journal of Economic Theory*, 145, 1776–1804. 4, 9

Massari, F. (2019) "Market Selection in Large Economies: A Matter of Luck," *Theoretical Economics*, 14 (2), 437–73. 33

Quiggin, J. (1982) "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3 (4), 323–343. 9

Sandroni, A. (2000) "Do Markets Favour Agents Able to Make Accurate Predictions?" *Econometrica*, 68 (6), 1303–1341. 2, 4, 6, 9, 12, 13, 20, 33, 35

Sarver, T. (2018) "Dynamic Mixture-Averse Preferences," *Econometrica*, 86 (4), 1347–1382. 9

Segal, Uzi and Avia Spivak (1990) "First Order versus Second Order Risk Aversion," *Journal of Economic Theory*, 51 (1), 111–125. 32

Strzalecki, T. (2013) "Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion," *Econometrica*, 81 (3), 1039–1074. 9

Yan, H. (2008) "Natural Selection in Financial Markets: Does It Work?" *Management Science*, 54, 1935–1950. 33

A Omitted Proofs

The main appendix contains the first-order condition proof, the two-consumer and multi-consumer characterization proofs, and the global and local robustness proofs. Proofs of some auxiliary lemmas are delegated to the Supplementary Material.

A.1 Proof of Proposition 1

Proof. Since (c, p) is a competitive equilibrium, $c^i \in \mathbb{C}$ must solve consumer i's maximization problem and, therefore, it must satisfy the one-deviation property. That is, for all $\sigma_t \in \Sigma_t$ and all t, $c(\sigma_t)$ and $\{c(\sigma_t, s) : s \in S\}$ must solve the following problem

$$\begin{aligned} \max \quad & F^i\left(x(\sigma_t), \mathcal{R}_t\left(F^i\left(x(\sigma_t, s), \mathcal{R}^i_{t+1}\left(V^i_{t+2}\right)\right)\right)\right) \\ \text{s.t.} \quad & p(\sigma_t)\, x(\sigma_t) + \sum_{s\in S} p(\sigma_t, s) \times x(\sigma_t, s) \leq p(\sigma_t)\, c(\sigma_t) + \sum_{s\in S} p(\sigma_t, s) \times c(\sigma_t, s), \\ & x(\sigma_t), x(\sigma_t, s) \geq 0. \end{aligned}$$

Since Assumptions 1 and 2 hold, $p(\sigma_t) > 0$ and $c^i(\sigma_t) > 0$ for all σ_t and all i. So the constraint qualification holds and there exists a Lagrange multiplier $\lambda^i(\sigma_t)$ such that

$\left(c^i(\sigma_t), \lambda^i(\sigma_t)\right)$ must satisfy the following FOC:

$$\begin{aligned}\lambda^i(\sigma_t)\, p(\sigma_t) =& F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right)\\ \lambda^i(\sigma_t)\, p(\sigma_t, s) =& F_2^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right) \times \frac{\partial \mathcal{R}_t^i}{\partial V_{\sigma_t, s}^i} \times F_1^i\left(c^i(\sigma_t, s), \mathcal{R}_{t+1}^i\left(V_{t+2}^i\right)\right).\end{aligned}$$

Therefore, we must have

$$\frac{p(\sigma_t, s)}{p(\sigma_t)} = \frac{F_1^i\left(c^i(\sigma_t, s), \mathcal{R}_{t+1}^i\left(V_{t+2}^i\right)\right)}{F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right)} \times \frac{\partial \mathcal{R}_t^i}{\partial V_{\sigma_t, s}^i} \times F_2^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right), \qquad \square$$

A.2 Proof of Lemma 1

Proof. The proof consists of the following four steps.

Step 1: We show that $\mathcal{R}_t^i\left(V_{t+1}^i\right)$ belongs to a compact set almost surely. Let $\overline{e} \equiv |I|\, M$ denote an upper bound of the aggregate endowment. Denote by $\overline{v}^i \equiv \operatorname{ess\,sup}_{\sigma,t} \left|V_{\sigma_t}^i(\overline{e})\right|$ and by $\underline{v}^i \equiv \operatorname{ess\,inf}_{\sigma,t} \left|V_{\sigma_t}^i(0)\right|$. By assumption, $V_{\sigma_t}^i(c)$ is monotonic in c, so $V_{\sigma_t}^i(c) \in \left[V_{\sigma_t}^i(0), V_{\sigma_t}^i(\overline{e})\right]$, and hence $V_t^i \in \left[\underline{v}^i, \overline{v}^i\right] \equiv \mathcal{V}^i$ $\mathbb{P}$-almost surely. As a consequence,

$$\mathcal{R}_t^i\left(V_{t+1}^i\right) \in \left[\mathcal{R}_t^i\left(\underline{v}^i \cdot \mathbf{1}\right), \mathcal{R}_t^i\left(\overline{v}^i \cdot \mathbf{1}\right)\right] = \left[\underline{v}^i, \overline{v}^i\right] = \mathcal{V}^i \quad \mathbb{P} - -a.s.,$$

which comes from the monotonicity of $\mathcal{R}_t^i$ and the fact that $\mathcal{R}_t^i(k \cdot \mathbf{1}) = k$. By assumption, $V^i(c) \in L_\infty^+$ for all $c \in \mathbb{C}$, so $0 \leq \underline{v}^i \leq \overline{v}^i < \infty$, and hence $\mathcal{V}^i$ is compact.

Step 2: We show $F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right)$ is almost surely bounded away from 0. Let $C \equiv [0, \overline{e}]$, which is the feasible set of consumption. From Assumption 1, $F_1^i : C \times \mathcal{V}^i \to (0, +\infty]$ is continuous and hence lower semicontinuous. Using the generalization of the Weierstrass' theorem to lower semicontinuous functions where $+\infty$ is allowed as a value (see Lemma 2.40 in Aliprantis and Border (1999)), we know that F_1^i obtains its minimum on $C \times \mathcal{V}^i$, so there exists $f_1 > 0$ such that $F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right) \geq f_1$ almost surely.

Step 3: We have $L^{ij}(\sigma_t) \to +\infty$ almost surely implies $F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right) \to +\infty$. This is a direct implication from Step 2.

Step 4: We show that $F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right) \to +\infty$ almost surely implies $c^i(\sigma_t) \to 0$. Suppose not. Then on a positive-probability set of paths, we can find an $\varepsilon > 0$ and an infinite sequence $\left\{c^i\left(\sigma_{t_k}\right)\right\}_{k=1}^\infty$ such that $c^i\left(\sigma_{t_k}\right) \geq \varepsilon$ for each k. Let $C_\varepsilon \equiv [\varepsilon, \overline{e}]$, and we have

$$F_1^i\left(c^i\left(\sigma_{t_k}\right), \mathcal{R}_{t_k}^i\left(V_{t_k+1}^i\right)\right) \leq \max_{(c,\mathcal{R}) \in C_\varepsilon \times \mathcal{V}^i} F_1^i(c, \mathcal{R}) < +\infty \quad \forall k, \quad \mathbb{P} - a.s.,$$

where the maximum can be obtained since $F_1^i : C_\varepsilon \times \mathcal{V}^i \to (0, +\infty)$ is a continuous function on a compact domain, and hence the standard Weierstrass' theorem applies. However, this fact contradicts $F_1^i\left(c^i(\sigma_t), \mathcal{R}_t^i\left(V_{t+1}^i\right)\right) \to +\infty$, so we must have $c^i(\sigma_t) \to 0$. $\square$

A.3 Auxiliary Results

Throughout Appendices A.3–A.6, Assumptions 1–7 are maintained, together with the standing assumption that recursive utility is continuous in the product topology.

Recall from (3) that consumption dynamics are determined by

$$L_{t+1}^{ij} = L_t^{ij} \times B^{ji}(\sigma_t, s_{t+1}) \times D^{ji}(\sigma_t).$$

To characterize the dynamics, it is convenient to consider a modified process $\left\{\hat{L}_t^{ij}\right\}_{t=1}^{\infty}$ where

$$\hat{L}_{t+1}^{ij} = \hat{L}_t^{ij} \times B^{ji}(\sigma_t, s_{t+1}) \times D^{ji}(\sigma_t, s_{t+1}).$$

These two processes are *comparable*: their ratios are uniformly bounded.

Lemma 2 (Comparability)**.** *There exist $0 < \alpha < \beta < +\infty$ such that $\alpha < L_t^{ij}/\hat{L}_t^{ij} < \beta$ for all $t \geq 1$.*

Proof. See Supplementary Material, Section S.1. □

Lemma 3 (Continuity)**.** *For all $i, j \in I$ and $t \in \mathbb{T}$, and $\varepsilon > 0$, there exists a uniform $\overline{L} \in \mathbb{R}_+$ such that*

$$\hat{L}_t^{ij} > \overline{L} \Longrightarrow |V_t^i - \underline{V}^i| < \varepsilon.$$

Proof. See Supplementary Material, Section S.1. □

A.4 Proof of Proposition 3

A.4.1 Proof of Proposition 3 (i)

Lemma 4. *When $j \succ_i i$, there exists $\overline{L} < \infty$ such that for all $t \in \mathbb{T}$, we have*

$$\mathbb{P}\left(\hat{L}_\tau^{-i} \leq \overline{L} \text{ for some } \tau \geq t \,\middle|\, \hat{L}_t^{-i} > \overline{L}\right) = 1, \tag{13}$$

where $\hat{L}_t^{-i} \equiv \min_{k \neq i} \hat{L}_t^{ki}$.

Remark 3. $\hat{L}_t^{-i} \in \left(\overline{L}, +\infty\right)$ means i almost dominates the market, so Lemma 4 says if i is locally dominated by another consumer, i can't dominate the market with positive probability.

Proof. See Supplementary Material, Section S.1. □

Proposition 3 (i) then follows from Lemma 4, because if i dominated the market, we would have $\hat{L}_t^{-i} \to +\infty$ with positive probability, which contradicts Lemma 4.

A.4.2 Proof of Proposition 3(ii)

For some $\underline{L} > 0$, define the stopping time $T = \inf\left\{\tau : \hat{L}_\tau^{i-} > \underline{L}\right\}$, where $\hat{L}_\tau^{i-} = \max_{j \neq i} \hat{L}_\tau^{ij}$. We construct a local martingale $m_t^{ij} = \left(\hat{L}_{t\wedge T}^{ij}\right)^\rho$. That is,

$$m_{t+1}^{ij} = \begin{cases} m_t^{ij} \times \left[B^{ji}(\sigma_t, s_{t+1}) \times D^{ji}(\sigma_t, s_{t+1})\right]^\rho & t < T \\ m_t^{ij} & t \geq T \end{cases},$$

We present the following facts.

Fact 2. *If $i \succ_i j$, there exists $\rho > 0$ such that $\mathbb{E}\left[\left(B_i^{ji}(s) \times D_i^{ji}(s)\right)^\rho\right] < 1$.*

Proof. See Supplementary Material, Section S.1. □

Fact 3. *If $i \succ_i j$, there exists $\underline{L}, \rho > 0$ such that $\left\{m_t^{ij}\right\}$ is a supermartingale.*

Proof. See Supplementary Material, Section S.1. □

Lemma 5. *If $i \succ_i j$ for all $j \neq i$, there exists $p > 0$ and $\varepsilon' > \varepsilon > 0$ such that for all $t \in \mathbb{T}$, we have*

$$\mathbb{P}\left(\hat{L}_\tau^{i-} \leq \varepsilon', \text{for all } \tau \geq t \,\middle|\, \hat{L}_t^{i-} \leq \varepsilon\right) > p. \tag{14}$$

Remark 4. Lemma 5 says that if i locally dominates everyone else, then conditional on i almost dominating the market, she will maintain the dominating position with a strictly positive probability. As shown later, this further implies that consumer i will (unconditionally) dominate the market with a strictly positive probability.

Proof. From Fact 3, $\left\{m_t^{ij}\right\}$ is a non-negative supermartingale, therefore it almost surely converges to some finite random variable m_∞^{ij} for all $j \neq i$. We have

$$\mathbb{P}(T < \infty) \leq \mathbb{P}\left(\exists j \neq i \colon m_\infty^{ij} > \underline{L}^\rho\right) \leq \sum_{j\neq i} \mathbb{P}\left(m_\infty^{ij} > \underline{L}^\rho\right) \leq \sum_{j\neq i} \frac{\mathbb{E}m_\infty^{ij}}{\underline{L}^\rho} \leq (I-1) \times \left(\frac{\hat{L}_0^{i-}}{\underline{L}}\right)^\rho \tag{15}$$

where the third inequality comes from Markov inequality and the last one comes from the supermartingale property. Let $\epsilon \in (0,1)$, (15) implies that

$$\mathbb{P}\left(T = \infty \,\middle|\, \hat{L}_0^{i-} \leq \left(\tfrac{\epsilon}{I-1}\right)^{1/\rho} \underline{L}\right) > 1 - \epsilon > 0. \tag{16}$$

For all $t \in \mathbb{T}$ and $\hat{L}_t^{i-}$, we construct a stopped process $\left\{m_{t+k}^{ij}\right\}_{k=0}^{\infty}$ that starts from time t and applying the same arguments.[37] Then, we have

$$\mathbb{P}\left(\hat{L}_\tau^{i-} \leq \underline{L} \text{ for all } \tau \geq t \,\middle|\, \hat{L}_t^{i-} \leq \left(\tfrac{\epsilon}{I-1}\right)^{1/\rho} \underline{L}\right) > 1 - \epsilon, \tag{17}$$

so the lemma is proved by setting $\varepsilon = \left(\frac{\epsilon}{I-1}\right)^{1/\rho} \underline{L}$, $\varepsilon' = \underline{L}$ and $p = 1 - \epsilon$. □

Next we show that Lemma 5 implies that consumer i will dominate the market with a strictly positive probability. First, Assumption 6 implies that for all $\hat{L}_0^{i-} < \infty$, there exists some $t \in \mathbb{T}$ such that $\hat{L}_t^{i-} \leq \left(\frac{\epsilon}{I-1}\right)^{1/\rho} \underline{L}$ happens with a strictly positive probability. Therefore, (17) implies that there exists some $\delta > 0$ such that

$$\mathbb{P}\left(\hat{L}_\tau^{i-} \leq \underline{L} \text{ for all } \tau \geq t\right) > \delta. \tag{18}$$

The martingale convergence theorem and (18) imply that $\hat{L}_t^{ij}$ converges to some limit $\hat{L}_\infty^{ij} \in \left[0, \underline{L}\right]$ with a strictly positive probability for all $j \neq i$. Assumptions 5 and 6 imply that the limit can only be $\hat{L}_\infty^{ij} = 0$ (otherwise $\hat{L}_t^{ij}$ can always drift away from $\hat{L}_\infty^{ij}$ with

[37] Formally, define $T = \inf\left\{\tau : \hat{L}_{\tau+t}^{i-} > \underline{L}\right\}$ and $m_{t+k}^{ij} = \left(\hat{L}_{t+\min\{k,T\}}^{ij}\right)^\rho$.

a probability bounded away from 0, and hence unstable). So, (18) is equivalent to that consumer i dominates the market with a strictly positive probability.

A.5 Proof of Theorem 1

Proof. (i) Suppose that $i \succ_i j$ and $i \succ_j j$. From $i \succ_j j$ and Lemma 4, for all $t \in \mathbb{T}$,

$$\mathbb{P}\left(\hat{L}_\tau^{ij} \leq \overline{L} \text{ for some } \tau \geq t \middle| \hat{L}_t^{ij} > \overline{L}\right) = 1. \tag{19}$$

Assumption 6 implies that there exists $\delta > 0$ such that for all $t \in \mathbb{T}$,

$$\mathbb{P}\left(\hat{L}_{t+K}^{ij} \leq \left(\frac{\epsilon}{I-1}\right)^{1/\rho} \underline{L} \middle| \hat{L}_t^{ij} \leq \overline{L}\right) > \delta. \tag{20}$$

From (16), $i \succ_i j$ implies that for all $t \in \mathbb{T}$, we have

$$\mathbb{P}\left(\hat{L}_\tau^{ij} \leq \underline{L} \text{ for all } \tau \geq t \middle| \hat{L}_t^{ij} \leq \left(\frac{\epsilon}{I-1}\right)^{1/\rho} \underline{L}\right) > 1 - \epsilon > 0. \tag{21}$$

Denote by $E = \cup_{k=1}^{\infty} \left\{\hat{L}_t^{ij} \leq \underline{L} \text{ for all } t \geq k\right\}$. Combining (19), (20) and (21), we know that for all $t \in \mathbb{T}$, we have

$$\mathbb{P}\left(E|\mathcal{F}_t\right) \geq (1-\epsilon)\,\delta > 0,$$

which further implies that $\mathbb{P}(E) = 1$ by Levy's 0-1 law, i.e., $\hat{L}_t^{ij}$ will be trapped below $\underline{L}$ eventually with probability 1. We can choose $\underline{L}$ to be arbitrarily small, so we must have $\hat{L}_t^{ij} \to 0$ almost surely, which implies that $c_t^j \to 0$, i.e., consumer j vanishes and consumer i is the only survivor.

(ii) Suppose that $i \succ_j j$ and $j \succ_i i$. Proposition 3 implies that both i and j dominate with zero probability. That is, they survive simultaneously with probability 1.

(iii) Suppose that $i \succ_i j$ and $j \succ_j i$. Proposition 3 implies that both i and j dominate with a strictly positive probability. We define $E^i = \cup_{k=1}^{\infty} \left\{\hat{L}_t^{ij} \leq \underline{L} \text{ for all } t \geq k\right\}$ and $E^j = \cup_{k=1}^{\infty} \left\{\hat{L}_t^{ij} \geq 1/\underline{L} \text{ for all } t \geq k\right\}$. Following the proof in case (i), there exists some $p > 0$ such that for all $t \in \mathbb{T}$, we have

$$\mathbb{P}\left(E^i \cup E^j | \mathcal{F}_t\right) > p \Longrightarrow \mathbb{P}\left(E^i \cup E^j\right) = 1.$$

Similarly, because we can choose $\underline{L}$ to be arbitrarily small, it implies that

$$\mathbb{P}\left(\left\{\hat{L}_t^{ij} \to 0\right\} \cup \left\{\hat{L}_t^{ij} \to +\infty\right\}\right) = 1 \Rightarrow \mathbb{P}\left(\left\{c_t^j \to 0\right\} \cup \left\{c_t^i \to 0\right\}\right) = 1.$$

That is, with probability 1, one of the two consumers will dominate the market. Besides, both E^i and E^j occur with a strictly positive probability by Proposition 3, so both $\mathbb{P}\left(c_t^j \to 0\right)$ and $\mathbb{P}\left(c_t^i \to 0\right)$ are strictly positive. □

A.6 Proof of Theorem 2

Lemma 6. *There exists some $p > 0$ and $\varepsilon' > \varepsilon > 0$ such that for all i with $U_i \neq \emptyset$ and all $t \in \mathbb{T}$, we have*

$$\mathbb{P}\left(\max_{j\in U_i}\hat{L}_T^{ij}\geq\varepsilon'\middle|\hat{L}_t^{i-}<\varepsilon\right)\geq p,$$

where $T=\inf\left\{\tau>t:\hat{L}_\tau^{i-}\geq\varepsilon'\right\}$.

Remark 5. Lemma 6 improves Lemma 4. Recall that Lemma 4 says that if there is a consumer j such that $j\succ_i i$, then i cannot dominate the market. Lemma 6 further says that when consumer i almost dominates the market, her dominating position will be destroyed by some other consumer j such that $j\succ_i i$ (i.e., the marginal utility of consumption of i increases relative to some $j\in U_i$ and falls outside of the dominating neighborhood).

Proof. See Supplementary Material, Section S.1. □

Lemma 7. *For every* $\varepsilon>0$*, there is some* $q>0$ *such that for all* $i\in I$ *with* $U_i\neq\emptyset$ *and all* $t\in\mathbb{T}$*, we have*

$$\mathbb{P}\left(\cup_{j\in U_i}\left\{\hat{L}_\tau^{j-}<\varepsilon\text{ for some }\tau\geq t\right\}|\hat{L}_t^{i-}<\varepsilon\right)\geq q. \tag{22}$$

Remark 6. Lemma 7 says that conditional on i almost dominating the market, some other consumer who dominates i (if any) will eventually almost dominate it as well.

Proof. See Supplementary Material, Section S.1. □

Next, we prove the following claim, which implies Theorem 2 directly.

Lemma 8. $\mathbb{P}\left(\cup_{i\in I^*}\left\{c_t^j\to 0\text{ for all }j\neq i\right\}\right)=1$, *with* $\mathbb{P}\left(c_t^j\to 0\text{ for all }j\neq i\right)>0$ *for all* $i\in I^*$.

Proof. See Supplementary Material, Section S.1. □

A.7 Proof of Theorem 4

We first prove a lemma showing that each $\mathcal{R}\in\mathscr{R}$ exhibits a local-expected-utility property.

Lemma 9 (Local-Expected-Utility Property)**.** *Suppose Assumptions 2 and 4 hold. Then, for all* $\mathcal{R}\in\mathscr{R}$ *and* $k\in\mathbb{R}_{++}$*, we have* $\nabla\mathcal{R}\left(k\cdot\mathbf{1}\right)\in\Delta\left(S\right)$.

Proof. See Supplementary Material, Section S.1. □

Proof of Theorem 4

Proof. The key part is to prove Theorem 4(ii). Theorem 4(i) then follows directly since we can find at least two preferences with different beliefs that satisfy $i\succ_i r$.

The proof of Theorem 4(ii) consists of two parts.

The "If" direction: $i\succ_i r\Rightarrow$ **robust survival.** Suppose that $i\succ_i r$, then we have

$$\mathbb{E}\log\left[\nabla_s\mathcal{R}^i\left(\overline{V}^i\right)\right]>\mathbb{E}\log\left(\pi_0\left(s\right)\right)\geq\mathbb{E}\log\left[\nabla_s\mathcal{R}^j\left(\underline{V}^j\right)\right]\quad\text{for all }\mathcal{R}^j\in\mathscr{R}, \tag{23}$$

where the first inequality comes from $i \succ_i r$, and the second inequality employs the facts that: (i) $\nabla\mathcal{R}^j\left(\underline{V}^j\right)$ is a probability distribution, which comes from Lemma 9 and that $\underline{V}^j$ is constant across states, and (ii) the entropy achieves its minimum at the true distribution π_0. (23) implies that $i \succ_i j$ for all $\mathcal{R}^j \in \mathscr{R}$. As a consequence, in all ε-perfect economies with sufficiently small ε, we also have $i \succ_i j$ for all $\mathcal{R}^j \in \mathscr{R}$ by continuity arguments. The previous argument applies to all $j \neq i$, so we have $i \succ_i j$ for all $j \neq i$ for all $\left\{\mathcal{R}^j\right\}_{j\neq i} \in \mathscr{R}^{I-1}$. Therefore, consumer i dominates the market with positive probability in ε-perfect economies with small ε for all $\left\{\mathcal{R}^j\right\}_{j\neq i} \in \mathscr{R}^{I-1}$, which establishes that $\mathcal{R}^i$ robustly survives.

The "only if" direction: robust survival $\Rightarrow i \succ_i r$. Suppose instead that $i \not\succ_i r$. By Assumption 8, we have $r \succ_i i$. We now prove that $\mathcal{R}^i$ cannot robustly survive. Let $\mathcal{R}^+ \in \mathscr{R}$ be any preference such that: (a) a consumer with $\mathcal{R}^+$ locally dominates r, and (b) the effective belief at certainty is correct, that is, $\nabla\mathcal{R}^+(k \cdot \mathbf{1}) = \pi_0$ for all $k \in \mathbb{R}_{++}$.[38] Consider the preference profile $\left\{\mathcal{R}^j\right\}_{j\neq i}$ where $\mathcal{R}^j = \mathcal{R}^+$ for all $j \neq i$. In this construction, j's effective belief is correct at certainty, so when i dominates the market, we must have $j \succ_i i$ because $r \succ_i i$ by assumption. Therefore, the local upper contour set of i is $U_i = I \setminus \{i\}$. On the other hand, each j locally dominates r, so we have $j \succ_j r \succeq_j k$ for all $k \neq j$. Thus, the local upper contour set of j is $U_j = \emptyset$ for all $j \neq i$. Theorem 2 implies that i vanishes almost surely. We found a preference profile $\left\{\mathcal{R}^j\right\}_{j\neq i}$ such that i vanishes, so $\mathcal{R}^i$ cannot robustly survive. □

A.8 Proof of Theorem 5(i)

Case 1: $\mathbb{E}\log\left[\nabla_s\mathcal{R}^i\left(\overline{V}^i\right)\right] < \mathbb{E}\log\left[\nabla_s\mathcal{R}^i\left(\underline{V}^i\right)\right]$.

From the definition of D, we know that for all $\delta > 0$, there exists some $\epsilon > 0$ such that

$$D\left(\mathcal{R}^j, \mathcal{R}^i\right) < \epsilon \Rightarrow \left|\mathbb{E}\log\left[\frac{\nabla_s\mathcal{R}^i\left(\underline{V}^i\right)}{\nabla_s\mathcal{R}^j\left(\underline{V}^j\right)}\right]\right| < \delta, \tag{24}$$

that is, the entropy of $\nabla\mathcal{R}^i\left(\underline{V}^i\right)$ and $\nabla\mathcal{R}^j\left(\underline{V}^j\right)$ can be arbitrarily close when $\mathcal{R}^i$ and $\mathcal{R}^j$ are sufficiently similar. Therefore, there is $\epsilon > 0$ such that whenever $D\left(\mathcal{R}^j, \mathcal{R}^i\right) < \epsilon$, we have

$$\mathbb{E}\log\left[\nabla_s\mathcal{R}^i\left(\overline{V}^i\right)\right] < \mathbb{E}\log\left[\nabla_s\mathcal{R}^j\left(\underline{V}^j\right)\right], \tag{25}$$

which implies that $j \succ_i i$, so consumer i dominates the market with 0 probability from Proposition 3, and hence $\mathcal{R}^i$ cannot locally robustly dominate.

Case 2: $\mathbb{E}\log\left[\nabla_s\mathcal{R}^i\left(\overline{V}^i\right)\right] > \mathbb{E}\log\left[\nabla_s\mathcal{R}^i\left(\underline{V}^i\right)\right]$.

Similarly, for all $\delta > 0$, there exists some $\epsilon > 0$ such that

$$D\left(\mathcal{R}^j, \mathcal{R}^i\right) < \epsilon \Rightarrow \left|\mathbb{E}\log\left[\frac{\nabla_s\mathcal{R}^i\left(\underline{V}^i\right)}{\nabla_s\mathcal{R}^j\left(\underline{V}^j\right)}\right]\right| < \delta \text{ and } \left|\mathbb{E}\log\left[\frac{\nabla_s\mathcal{R}^i\left(\overline{V}^i\right)}{\nabla_s\mathcal{R}^j\left(\overline{V}^j\right)}\right]\right| < \delta. \tag{26}$$

[38] Such a preference exists, e.g., utility with correct belief as in Example 11, which can be viewed as a special case of the hyperbolic absolute ambiguity aversion (HAAA) preferences in Guerdjikova and Sciubba (2015).

So, there is $\epsilon > 0$ such that for all $\mathcal{R}^j, \mathcal{R}^k \in B_\epsilon\left(\mathcal{R}^i\right) \equiv \left\{\mathcal{R} \in \mathscr{R} : D\left(\mathcal{R}, \mathcal{R}^i\right) < \epsilon\right\}$, we have

$$\mathbb{E} \log \left[\nabla_s \mathcal{R}^j\left(\overline{V}^j\right)\right] > \mathbb{E} \log \left[\nabla_s \mathcal{R}^k\left(\underline{V}^k\right)\right]. \tag{27}$$

That is, for all $\mathcal{R}^j \in B_\epsilon\left(\mathcal{R}^i\right)$, we have $j \succ_j k$ for all $\mathcal{R}^k \in B_\epsilon\left(\mathcal{R}^i\right)$, so j dominates the market with a strictly positive probability, which implies that consumer i cannot dominate the market with probability 1 from Proposition 3, and hence $\mathcal{R}^i$ cannot locally robustly dominate.

Case 3: $\mathbb{E} \log \left[\nabla_s \mathcal{R}^i\left(\overline{V}^i\right)\right] = \mathbb{E} \log \left[\nabla_s \mathcal{R}^i\left(\underline{V}^i\right)\right]$**.**

To prove this case, we construct the following preference

$$\mathcal{R}^j(V) = \left(1 - \varepsilon_j\right) \mathcal{R}^i(V) + \varepsilon_j \times \mathcal{R}^*(V), \tag{28}$$

where $\varepsilon_j \in [0, 1]$ and $\mathcal{R}^*$ is a preference that satisfies $\mathbb{E} \log \left[\frac{\nabla_s \mathcal{R}^*\left(\overline{V}^i\right)}{\nabla_s \mathcal{R}^i\left(\underline{V}^i\right)}\right] > 0$, e.g., as in Example 11. Next, we want to prove that $j \succ_j i$ in some ε-perfected economy where ε can be arbitrarily small. We first state the following lemma.

Lemma 10. *Suppose $\mathcal{R}^i$ is effectively uncertainty neutral. For $\mathcal{R}^j$ defined by (28), we have $\mathbb{E} \log \left[\frac{\nabla_s \mathcal{R}^j\left(\overline{V}^i\right)}{\nabla_s \mathcal{R}^i\left(\underline{V}^i\right)}\right] > 0$ for every $\varepsilon_j > 0$.*

Proof. See Supplementary Material, Section S.1. □

Note that Lemma 10 does not directly imply $j \succ_j i$ because we only have $\nabla_s \mathcal{R}^j\left(\overline{V}^i\right)$ instead of $\nabla_s \mathcal{R}^j\left(\overline{V}^j\right)$, and it is possible that $\overline{V}^j \neq \overline{V}^i$. However, it turns out that there always exists some ε-perfection in which $\overline{V}^j = \overline{V}^i$, which is stated as a lemma below. Denote by $\mathscr{U}$ the set of feasible utility functions. We have the following lemma.

Lemma 11. *For every $\delta > 0$, there exists $\overline{\varepsilon}$ such that for every $\varepsilon_j \in (0, \overline{\varepsilon})$, there exists $u_j \in \mathscr{U}$ satisfying $||u_j - u_i|| < \delta$ and $\overline{V}^j = \overline{V}^i$, where $||\cdot||$ denotes the sup-norm over some compact set C.*

Proof. See Supplementary Material, Section S.1. □

Combining Lemmas 10 and 11, we can construct some $\mathcal{R}^j$ such that $j \succ_j i$ in some ε-perfected economy where ε can be arbitrarily small. Suppose we choose other consumers' preferences sufficiently close to i, we will have $j \succ_j k$ for all $k \neq j$, so consumer j dominates the market with strictly positive probability, thus $\mathcal{R}^i$ cannot locally robustly dominate.

A.9 Proof of Theorem 5(ii)

The "if" direction. Suppose $\mathcal{R}^i$ is not effectively uncertainty neutral. We have two cases.

Case 1: $\mathbb{E} \log \left[\nabla_s \mathcal{R}^i\left(\overline{V}^i\right)\right] < \mathbb{E} \log \left[\nabla_s \mathcal{R}^i\left(\underline{V}^i\right)\right]$.

As in (27), there is $\epsilon > 0$ so that for all $\mathcal{R}^j, \mathcal{R}^k \in B_\epsilon\left(\mathcal{R}^i\right) \equiv \left\{\mathcal{R} \in \mathscr{R} : D\left(\mathcal{R}, \mathcal{R}^i\right) < \epsilon\right\}$,

$$\mathbb{E} \log \left[\nabla_s \mathcal{R}^j \left(\overline{V}^j\right)\right] < \mathbb{E} \log \left[\nabla_s \mathcal{R}^k \left(\underline{V}^k\right)\right]. \tag{29}$$

That is, for all $\mathcal{R}^j \in B_\epsilon\left(\mathcal{R}^i\right)$, we have $k \succ_j j$ for all $\mathcal{R}^k \in B_\epsilon\left(\mathcal{R}^i\right)$, so every consumer j is locally dominated by all other consumers. Suppose $|I| = 2$, then both consumers coexist almost surely from Theorem 1, so $\mathcal{R}^i$ locally robustly survives.

Case 2: $\mathbb{E} \log \left[\nabla_s \mathcal{R}^i \left(\overline{V}^i\right)\right] > \mathbb{E} \log \left[\nabla_s \mathcal{R}^i \left(\underline{V}^i\right)\right]$**.**

By Case 2 in the proof of Theorem 5(i), every consumer j locally dominates the others, so i survives with positive probability and $\mathcal{R}^i$ locally robustly survives.

The "only if" direction. Suppose instead that $\mathcal{R}^i$ is effectively uncertainty neutral, then we want to show that $\mathcal{R}^i$ cannot locally robustly survive, i.e., it can almost surely vanish in the presence of some very close preferences and in small ε-perfected economies. Since we are interested in ε-perfection, it is without loss (see footnote 40) to focus on the following case

$$\mathbb{E} \log \left[\nabla_s \mathcal{R}^i \left(\overline{V}^i\right)\right] = \mathbb{E} \log \left[\nabla_s \mathcal{R}^i \left(\underline{V}^i\right)\right] \neq \mathbb{E} \log \left[\pi_0(s)\right]. \tag{30}$$

In this case, the entropy of local effective beliefs of $\mathcal{R}^i$ is not equal to that of the correct belief entropy.[39] We construct $\mathcal{R}^j$ with form (28), where $\mathcal{R}^*$ satisfies both: (i) $\mathbb{E} \log \frac{\nabla_s \mathcal{R}^*\left(\overline{V}^i\right)}{\nabla_s \mathcal{R}^i\left(\underline{V}^i\right)} > 0$, and (ii) $\mathbb{E} \log \frac{\nabla_s \mathcal{R}^*\left(\underline{V}^i\right)}{\nabla_s \mathcal{R}^i\left(\overline{V}^i\right)} > 0$, e.g., the HAAA preferences with correct beliefs as in Guerdjikova and Sciubba (2015). From the proof of Theorem 5, we already know that (i) implies that $j \succ_j i$ in some ε-perfected economies with sufficiently small ε. Similarly, we also have (ii) implies that $j \succ_i i$ in some ε-perfected economies with sufficiently small ε.[40]

In the case where $|I| = 2$, this directly implies that consumer j dominates almost surely, and hence i vanishes almost surely, so $\mathcal{R}^i$ cannot locally robustly survive. Note that the "only if" direction applies to an arbitrary number of consumers. In this case, the continuity further implies that for all $j \neq i$, we have: (i) $j \succ_j j'$ for all $j' \neq j$, and (ii) $j \succ_i i$. Therefore, we have $U_i = I \setminus \{i\}$ and $U_j = \emptyset$ for all $j \neq i$. Theorem 2 implies that consumer i vanishes almost surely, so $\mathcal{R}^i$ cannot locally robustly survive.

[39] Since $\nabla_s \mathcal{R}^i\left(\underline{V}^i\right)$ is a probability distribution, $\mathbb{E} \log \left[\nabla_s \mathcal{R}^i \left(\overline{V}^i\right)\right] = \mathbb{E} \log \left[\nabla_s \mathcal{R}^i \left(\underline{V}^i\right)\right] < \mathbb{E} \log \left[\pi_0(s)\right]$.

[40] Suppose instead (30) does not hold, then we still construct the same $\mathcal{R}^j$ with the same choice of $\mathcal{R}^*$. But the only difference is that $\mathcal{R}^*$ cannot satisfy (ii) as the local effective beliefs of $\mathcal{R}^i$ is correct. However, we can find some ε-perfected economies in which the perturbed $\mathcal{R}^i$ has an incorrect belief. Therefore, in these perturbed economies, we still have $j \succ_j i$ and $j \succ_i i$, and hence results won't change.