NBER WORKING PAPER SERIES

MARKET SELECTION

Leonid Kogan
Stephen Ross
Jiang Wang
Mark M. Westerfield

Working Paper 15189
http://www.nber.org/papers/w15189

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2009

Market Selection
Leonid Kogan, Stephen Ross, Jiang Wang, and Mark M. Westerfield
NBER Working Paper No. 15189
July 2009
JEL No. D51,D53,G1,G11,G12,G14

ABSTRACT

The hypothesis that financial markets punish traders who make relatively inaccurate forecasts and eventually eliminate the effect of their beliefs on prices is of fundamental importance to the standard modeling paradigm in asset pricing. We establish necessary and sufficient conditions for agents making inferior forecasts to survive and to affect prices in the long run in a general setting with minimal restrictions on endowments, beliefs, or utility functions. We show that the market selection hypothesis is valid for economies with bounded endowments or bounded relative risk aversion, but it cannot be substantially generalized to a broader class of models. Instead, survival is determined by a comparison of the forecast errors to risk attitudes. The price impact of inaccurate forecasts is distinct from survival because price impact is determined by the volatility of traders' consumption shares rather than by their level. Our results also apply to economies with state-dependent preferences, such as habit formation.

Leonid Kogan
MIT Sloan School of Management
50 Memorial Drive, E52-434
Cambridge, MA 02142
and NBER
lkogan@mit.edu

Stephen Ross
Sloan School of Management
Massachusetts Institute of Technology
50 Memorial Drive, E52-450
Cambridge, MA 02142-1347
and NBER
sross@mit.edu

Jiang Wang
E52-456, MIT
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
wangj@mit.edu

Mark M. Westerfield
Department of Finance & Business Economics
Marshall School of Business
Hoffman Hall 602D, Mailcode 1427
701 Exposition Blvd
Los Angeles, CA 90089-1427
mwesterf@usc.edu

1 Introduction

It has long been suggested that evolutionary forces work in financial markets: agents who are inferior at forecasting the future will either improve through learning or perish as their wealth diminishes relative to those superior in forecasting. This argument was first made in Friedman (1953), although much of the recent work stems from De Long, Shleifer, Summers, and Waldman (1991) and Blume and Easley (1992). If such an evolutionary mechanism works effectively, then in the long run only those agents with the best forecasts will survive the market selection process and determine asset prices. This "market selection hypothesis" (MSH) is one of the major arguments behind the assumption of rational expectations in neoclassical asset pricing theory. After all, if agents with more accurate knowledge of fundamentals do not determine the price behavior in the market, there is little reason to assume that prices are driven by fundamentals and not by behavioral biases. More generally, it may be comforting that markets select for those agents with more accurate forecasts, even if agents with less accurate forecasts are replenished over time (e.g. in overlapping generations economies). We show that in frictionless, complete-market exchange economies, both parts of the MSH – that traders with inferior forecasts do not survive and that extinction destroys their price impact – are false in general. With minimal restrictions on endowments, preferences, and beliefs, we develop necessary and sufficient conditions for the validity of the market selection hypothesis.

Despite the appeal and importance of the market selection hypothesis, its validity has remained ambiguous. Existing literature provides a number of examples in which agents with biased beliefs may or may not survive and/or influence prices. Relying on partial equilibrium analysis, De Long, Shleifer, Summers, and Waldman (1991) argue that agents making inferior forecasts can survive in wealth terms despite market forces exerted by agents with objective beliefs. Using a general equilibrium setting, Sandroni (2000) and Blume and Easley (2006) show that only agents with beliefs closest to the objective probabilities will survive and have price impact. Their results are obtained in economies with a limited range of primitives, specifically in those with bounded aggregate consumption.[1] Kogan, Ross, Wang, and Westerfield (2006) demonstrate in a setting without intermediate consumption that if aggregate endowment is unbounded, agents with incorrect beliefs can survive. Moreover, they show that even when such agents do not survive in the long-run, their impact on prices can persist. In other words, survival and price impact are two independent concepts and need to be considered separately. However, the absence of intermediate consumption leaves it less clear how important this distinction may be in more general models.

Existing analysis relies on specialized models, mostly for tractability and convenience, making it difficult to understand the economic mechanism behind the MSH and the scope of its validity. In this paper, we perform a comprehensive analysis of the MSH and its pricing implications in a general complete-market setting with time-separable preferences (including state-dependent preferences, e.g., catching up with the Joneses), not limiting ourselves to commonly used parametric specifications. We thus sharpen our understanding of the degree of robustness of the survival and price impact results and their connection to the economic fundamentals. For instance, in models with constant relative risk aversion (CRRA) preferences (e.g., Dumas, Kurshev, and Uppal (2008), Yan (2008)), the market selection hypothesis holds, and agents with inferior forecasts fail to survive or affect prices in the long run. However, as we discuss below, models with CRRA preferences are effectively a knife-edge case for the validity of the MSH. The MSH may be violated for utility functions arbitrarily close to CRRA, but with risk aversion depending on the consumption level.

We examine the MSH in frictionless and complete-market economies because common arguments in favor of its validity rely on unrestricted competition, the lack of limits to arbitrage, etc. To isolate the impact of disagreement, we populate our economies with competitive agents who only differ in their beliefs. We then analyze how survival and price impact properties of the economy depend on the primitives, such as errors in forecasts, endowment growth, and risk aversion.

[1] A significant body of work exists examining pricing implications of heterogeneous beliefs in specific parameterized models, including Dumas, Kurshev, and Uppal (2008), Fedyk and Walden (2007), Xiong and Yan (2008), and Yan (2008).

We find that if agents have bounded relative risk aversion or if the aggregate endowment is bounded, then agents with more accurate forecasts eventually dominate the economy and determine price behavior. Otherwise, if relative risk aversion is unbounded, the survival of agents with less accurate forecasts and their impact on state prices are effectively determined by the asymptotic growth rates of risk aversion and forecast errors: if forecast errors do not disappear fast enough compared to the growth rate of risk aversion, agents with less accurate forecasts can maintain a nontrivial consumption share and affect prices. Our results and counter-examples suggest it is not possible to substantially generalize the market selection hypothesis beyond the class of models with bounded relative risk aversion or aggregate endowment.

Intuitively, survival depends on the tradeoff between the forecast errors and the growth rate of risk aversion. Agents with heterogeneous beliefs trade with each other to share consumption across states. When two agents disagree in their probability assessment of a particular state, the more optimistic agent buys a disproportionate share of the state-contingent consumption. If two agents have diverging beliefs, they end up with extreme disagreement asymptotically over most states. Whether this extreme disagreement leaves one of the agents with a vanishing consumption share depends on agents' preferences. Pareto optimality implies that the ratio of agents' marginal utilities in each state must be inversely proportional to the ratio of their belief densities, and therefore, asymptotically, divergence in beliefs leads to divergence in marginal utilities. Whether or not large differences in marginal utilities correspond to small differences in consumption depends on the sensitivity of marginal utility to consumption, which is the same as the coefficient of relative risk aversion: $d \ln U'(C)/d \ln C = CU''(C)/U'(C) = -\gamma(C)$. If risk aversion of the two agents grows fast enough compared to their belief differences, their marginal utility differences may not translate into large consumption differences. In fact, as we show below (in Example 3.3), the two agents may consume equal consumption shares asymptotically despite their growing disagreement.

We show that in models with bounded relative risk aversion, the agent with inferior forecasts has no long-run impact on prices of Arrow-Debreu securities. This is not the case for models with unbounded relative risk aversion. In addition to the possibility that the agent making inferior forecasts maintains a nontrivial consumption share and thus affects prices in the long run, the precise conditions for survival and price impact are in general different. Therefore, certain economies may exhibit one without the other. This phenomenon was first observed by Kogan, Ross, Wang, and Westerfield (2006), and here we re-establish their result under a stricter definition of price impact and in much more general settings.

To gain intuition for why lack of survival does not always imply lack of price impact, consider the following exchange economy with two agents. Let D_t be the endowment, and let the agents have preferences given by $e^{-\rho t}U(C_t)$. Assume that belief differences imply that the second agent consumes a share w_t of the aggregate endowment. Then, the first agent consumes $C_t = D_t(1 - w_t)$. Assume that w_t vanishes asymptotically, and thus the second agent does not survive in the long run. Next, compare the stochastic discount factor in this economy to the one in an identical economy without the second agent, i.e., with $w_t = 0$. Assuming that all quantities are driven by Ito processes, in the second economy the volatility of the stochastic discount factor equals $\gamma(D_t)\mathrm{vol}(dD_t/D_t)$, where $\gamma(D) = -DU''(D)/U'(D)$. This compares to $\gamma(C_t)\mathrm{vol}(dC_t/C_t)$ in the first economy, where $\mathrm{vol}(dC_t/C_t) = \mathrm{vol}(dD_t/D_t - dw_t/(1 - w_t))$. Therefore, if the volatility of the consumption share w_t does not vanish relative to the volatility of endowment growth, the two discount factors may exhibit different volatilities.[2] Thus, in contrast to the concept of survival, which is defined by the magnitude of the consumption share, price impact depends on the volatility of the consumption share,

[2]As an example, we specify exogenously a pair of processes (D_t, w_t) and a utility function $U(C)$, such that the second agent has a vanishing consumption share but maintains price impact.

Let Z_t and Z'_t be two independent Brownian motions, let $f(t) = (1+t)^{-1}$, and define D_t and ω_t as

$$\begin{aligned} D_t &= \exp(t + Z_t), \\ w_t &= f(t)X_t, \\ dX_t &= f(t)^{-1}X_t(1 - X_t)\, dZ'_t. \end{aligned}$$

Assume that the utility function $U(C) = \ln(C)$. It is easy to check that in the economy without the second agent, the volatility of the stochastic discount factor equals one, while in the economy with the second agent, it is $\sqrt{1 + \frac{X_t^2(1-X_t)^2}{(1-f(t)X_t)^2}}$, which does not converge to one. Thus, as time goes to infinity, w_t vanishes, but the ratio of volatilities of the two pricing kernels does not converge to one.

In the above example, demand of the second agent is specified in reduced form, instead of being derived from his beliefs and preferences. As we show below, to obtain price impact without survival in an economy in which agents' demands differ only due to their disagreement, one must consider utility functions with unbounded relative risk aversion. We describe one such economy in Example 3.5.

and therefore lack of survival does not imply lack of price impact.

Our results cover general state-dependent preferences, such as external habit formation and catching-up-with-the-Joneses. State-dependent preferences change the risk attitudes of the agents in the economy, but they do not change how those risk attitudes affect survival or price impact. We are therefore able to apply the necessary and sufficient conditions for the validity of the MSH to models with state-dependent preferences that are commonly used in the literature. This aspect of our analysis is new to our knowledge, and provides a first exploration of the market selection hypothesis in the context of "behavioral" preferences.

The paper is organized as follows. Section 2 sets up the model and defines survival and price impact. Section 3 presents several examples of how survival and price impact results depend on the primitives of the economy. Sections 4 and 5 present our main results on survival and price impact. Section 6 covers economies with state-dependent preferences. Section 7 concludes. Proofs and derivations are in the Appendix.

2 The Model

We consider an infinite-horizon exchange (endowment) economy. Time is indexed by t, which takes values in $t \in [0, \infty)$. Time can either be continuous or discrete. While all of our general results can be stated either in discrete or continuous time, some of the examples are simpler in continuous time. We will use integrals to denote aggregation over time. When time is taken as discrete, time-integration will be interpreted as summation. We further assume that there is a single, perishable consumption good, which is also used as the numeraire.

Uncertainty and the Securities Market

The environment of the economy is described by a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Each element $\omega \in \Omega$ denotes a state of the economy. The information structure of the economy is given by a filtration on $\mathcal{F}$, $\{\mathcal{F}_t\}$, with $\mathcal{F}_s \subset \mathcal{F}_t$ for $s \leq t$. The probability measure $\mathbb{P}$ is referred to as the objective probability measure. The endowment flow is given by an adapted process D_t. We assume that the aggregate endowment is strictly positive: $D_t > 0$, $\mathbb{P} - a.s.$

In addition to the objective probability measure $\mathbb{P}$, we also consider other probability measures, referred to as subjective probability measures. Let $\mathbb{A}$ and $\mathbb{B}$ denote such measures. We assume that $\mathbb{A}$ and $\mathbb{B}$ share zero-probability events with $\mathbb{P}$. Denote the Radon-Nikodym derivative of the probability measure $\mathbb{A}$ with respect to $\mathbb{P}$ by $\xi_t^{\mathbb{A}}$. Then

$$\mathrm{E}_t^{\mathbb{A}}\left[Z_s\right] = \mathrm{E}_t^{\mathbb{P}}\left[\frac{\xi_s^{\mathbb{A}}}{\xi_t^{\mathbb{A}}} Z_s\right] \tag{1}$$

for any $\mathcal{F}_s$-measurable random variable Z_s and $s \geq t$, where $\mathrm{E}_t\left[Z\right]$ denotes $\mathrm{E}\left[Z|\mathcal{F}_t\right]$. In addition, $\xi_0^{\mathbb{A}} \equiv 1$ The probability measure $\mathbb{B}$ has a similar Radon-Nikodym derivative $\xi_t^{\mathbb{B}}$. The random variable $\xi_t^{\mathbb{A}}$ can be informally interpreted as the density of the probability measure $\mathbb{A}$ with respect to the probability measure $\mathbb{P}$ conditional on the time-t information set.

We use $\mathbb{A}$ and $\mathbb{B}$ to model heterogeneous beliefs. We define the ratio of subjective belief densities

$$\xi_t = \frac{\xi_t^{\mathbb{B}}}{\xi_t^{\mathbb{A}}}. \tag{2}$$

Since both $\xi^{\mathbb{A}}$ and $\xi^{\mathbb{B}}$ are nonnegative martingales, they converge almost surely as time tends to infinity (e.g., Shiryaev (1996, §7.4, Th. 1)), and therefore the process ξ_t also converges. Our results are most relevant for models in which the limit of ξ_t is either zero or infinity, implying that the agents' beliefs, described by subjective measures $\mathbb{A}$ and $\mathbb{B}$, are meaningfully different in the long run. To see that convergence to a finite limit implies that beliefs are not meaningfully different in the long run, consider the subset of the probability measure where ξ_t converges to a finite limit. Then, the ratio $\frac{\xi_{t+T}^{\mathbb{A}}/\xi_t^{\mathbb{A}}}{\xi_{t+T}^{\mathbb{B}}/\xi_t^{\mathbb{B}}}$, $T > 0$, converges to one, so the finite-period forecasts implied by the two subjective measures converge asymptotically.[3] We examine the asymptotic condition on subjective beliefs in more detail in Section 3.

We assume that there exists a complete set of Arrow-Debreu securities in the economy, so that the securities market is complete.

[3] See Blume and Easley (2006) for further discussion.

Agents

There are two competitive agents in the economy. They have the same utility function, but differ in their beliefs. The first agent has $\mathbb{A}$ as his probability measure while the second agent has $\mathbb{B}$ as his probability measure. We refer to the agent who uses $\mathbb{A}$ as agent $\mathbb{A}$ and the agent who uses $\mathbb{B}$ as agent $\mathbb{B}$. It is clear from the context when we refer to an agent as opposed to a probability measure.

Until stated otherwise, we assume that the agents' utility function is time-additive and state-independent with the canonical form

$$\int_0^{\infty} e^{-\rho t} u(C_t) dt$$

where C_t is an agent's consumption at time t, ρ is the time-discount coefficient and $u(\cdot)$ is the utility function. We consider more general forms of the utility function in Section 6. The common utility function $u(\cdot)$ is assumed to be increasing, weakly-concave, and twice continuously differentiable. We assume that $u(\cdot)$ satisfies the standard Inada condition at zero:

$$\lim_{x \to 0} u'(x) = \infty. \tag{3}$$

We use $A(x) \equiv -u''(x)/u'(x)$ and $\gamma(x) \equiv -xu''(x)/u'(x) = xA(x)$ to denote, respectively, an agent's absolute and relative risk aversion at the consumption level x.

Let $C_{\mathbb{A},t}$ and $C_{\mathbb{B},t}$ denote consumption of the two agents. Each agent maximizes his expected utility using his subjective beliefs. Agent i's objective is

$$\mathrm{E}_0^i \left[\int_0^{\infty} e^{-\rho t} u(C_{i,t})\, dt \right] = \mathrm{E}_0^{\mathbb{P}} \left[\int_0^{\infty} e^{-\rho t} \xi_t^i u(C_{i,t})\, dt \right], \quad i \in \{\mathbb{A}, \mathbb{B}\}, \tag{4}$$

where the equality follows from (1). This implies that the two agents are observationally equivalent to the two agents with objective beliefs $\mathbb{P}$ but state-dependent utility functions $\xi_t^{\mathbb{A}} u(\cdot)$ and $\xi_t^{\mathbb{B}} u(\cdot)$ respectively.

The two agents are collectively endowed with a flow of the consumption good. Let the initial share of the total endowment for agent $\mathbb{A}$ and $\mathbb{B}$ be $1-\varpi$ and ϖ, respectively.

Equilibrium

Because the market is complete, if an equilibrium exists, it must be Pareto-optimal. In such situations, consumption allocations can be determined by maximizing a weighted sum of the utility functions of the two agents. The equilibrium is given at each time t by

$$\max_{C_{\mathbb{A},t},\,C_{\mathbb{B},t}} \quad (1-\alpha)\,\xi_t^{\mathbb{A}} u(C_{\mathbb{A},t}) + \alpha\,\xi_t^{\mathbb{B}} u(C_{\mathbb{B},t}) \tag{5}$$

$$\text{s.t.} \quad C_{\mathbb{A},t} + C_{\mathbb{B},t} = D_t$$

where $\alpha \in [0,1]$.

Concavity of the utility function, together with the Inada condition, imply that the equilibrium consumption allocations satisfy the first-order condition

$$\frac{u'(C_{\mathbb{A},t})}{u'(C_{\mathbb{B},t})} = \lambda\,\xi_t, \tag{6}$$

where we denote $\alpha/(1-\alpha)$ by λ.

We define $w_t = \frac{C_{\mathbb{B},t}}{D_t}$ as the share of the aggregate endowment consumed by agent $\mathbb{B}$. The first-order condition for Pareto optimality (6) implies that w_t satisfies

$$-\ln(\lambda\xi_t) = -\ln u'((1-w_t)D_t) + \ln u'(w_t D_t) = \int_{w_t D_t}^{(1-w_t)D_t} A(x)\,dx, \tag{7}$$

since $A(x) = -\frac{d}{dx}\ln u'(x)$. This equation relates belief differences (ξ_t) to individual risk aversion ($A(x)$) and the equilibrium consumption allocation (w_t and D_t), and will be our primary analytical tool.

Definitions of Survival and Price Impact

Without loss of generality, we focus on the survival of agent $\mathbb{B}$ and that agent's impact on security prices in the long run. If one replaces $\lambda\xi_t$ with $\frac{1}{\lambda\xi_t}$ in our analysis, our results instead describe the survival and price impact of agent $\mathbb{A}$.

We first define formally the concepts of survival and price impact to be used in this paper and examine their properties.

Definition 1 [Extinction and Survival] *Agent $\mathbb{B}$ becomes extinct if*

$$\lim_{t\to\infty} \frac{C_{\mathbb{B},t}}{D_t} = 0, \quad \mathbb{P}-\text{a.s.}$$

Agent $\mathbb{B}$ survives if he does not become extinct.

The above definition provides a weak condition for survival: an agent has to consume a positive fraction of the endowment with a positive probability in order to survive.

We define price impact in terms of the state-price density m_t. Our definition formalizes the notion that agent $\mathbb{B}$ has no price impact as long as his beliefs do not affect the state-price density asymptotically. Our definition of price impact in terms of the state-price density is natural for a complete-market economy. Long-lasting distortions of the state-price density imply that some long-lived assets, which can be replicated as portfolios of primitive Arrow-Debreu securities (state-contingent claims), must also be mispriced. However, the reverse implication does not hold, and some portfolios of primitive Arrow-Debreu securities may reveal price impact even if the state-price density is not affected by agent $\mathbb{B}$'s beliefs in the long run.[4] Formally, this reflects the possibility that almost sure convergence of random variables may not imply convergence of their moments. Thus, our definition of price impact

[4]For instance, as we show below, in models with CRRA preferences the agent making inferior forecasts has no price impact according to the above definition. This finding may appear to be at odds with the results in Kogan, Ross, Wang, and Westerfield (2006), who also consider a model with CRRA preferences. However, note that in addition to the differences in settings –Kogan, Ross, Wang, and Westerfield (2006) do not allow for intermediate consumption, while here we do – we adopt a stricter definition of price impact in this paper.

is relatively strict, and the set of economies in which agent $\mathbb{B}$'s beliefs affect prices of some long-lived assets is potentially larger than our definition suggests.

Pareto optimality and the individual optimality conditions imply that

$$m_t = e^{-\rho t} \frac{\xi_t^{\mathbb{A}} u'((1-w_t)D_t)}{u'((1-w_0)D_0)} = e^{-\rho t} \frac{\xi_t^{\mathbb{B}} u'(w_t D_t)}{u'(w_0 D_0)}. \tag{8}$$

In general, m_t depends on λ, the relative weight of the two agents in the economy, through their initial endowments. Thus, we write $m_t = m_t(\lambda)$. We denote by $m_t^\star(\lambda)$ the state-price density in the economy in which both agents have beliefs described by the measure $\mathbb{A}$ and hence $\xi_t = 1$. We define $m_t(0)$ to be the state-price density in an economy in which all wealth is initially allocated to agent $\mathbb{A}$. We identify the price impact exerted by agent $\mathbb{B}$ by comparing m_t to $m^\star$.

Definition 2 [Price Impact] *Agent $\mathbb{B}$ has no price impact if there exists $\lambda^\star \geq 0$, such that for any $s > 0$,*

$$\lim_{t\to\infty} \frac{m_{t+s}(\lambda)/m_t(\lambda)}{m_{t+s}^\star(\lambda^\star)/m_t^\star(\lambda^\star)} = 1, \quad \mathbb{P}-\text{a.s.} \tag{9}$$

Otherwise, he has price impact.

In contrast to the notion of long-run survival, equations (8) and (9) show that price impact is determined by *changes* in consumption over finite time intervals relative to a benchmark economy. In particular, we compare the state price density in the original economy, $m_{t+s}(\lambda)/m_t(\lambda)$, to the one in a reference economy where both agents maintain the same beliefs, but, possibly, have a different initial wealth distribution, $m_{t+s}^\star(\lambda)/m_t^\star(\lambda^\star)$. We define the price impact in this way for two reasons. First, we wish to focus on the price impact of differences in beliefs, not the asymptotic distortion of the wealth distribution caused by differences in beliefs during the earlier time periods. Thus, we allow the relative weight of the two agents in the reference economy, $\lambda^\star$, to be different from that in the original economy. In addition, this definition of price impact remains applicable when both agents survive in the long run.

The above definition may seem difficult to apply because the condition (9) must be verified for all values of $\lambda^\star$. However, for economies in which agent $\mathbb{B}$ does not survive it is often sufficient to verify the definition for $\lambda^\star = 0$. In that case, since

$$\frac{u'(D(1-w))}{u'(D)} = \exp\left(\int_{D(1-w)}^{D} A(x)\,dx\right), \tag{10}$$

a sufficient condition for the absence of price impact is

$$\forall\ s > 0: \quad \lim_{t\to\infty} \int_{D_{t+s}(1-w_{t+s})}^{D_{t+s}} A(x)\,dx - \int_{D_t(1-w_t)}^{D_t} A(x)\,dx = 0, \quad \mathbb{P}-\text{a.s.} \tag{11}$$

When agent $\mathbb{B}$ survives in the long run, it is natural to consider $\lambda^\star > 0$ for the reference economy. We find that the case of $\lambda^\star = 1$ is often sufficient. Under this assumption, the two agents in the reference economy consume equal amounts and we obtain the following sufficient condition for the absence of price impact

$$\forall\ s > 0: \quad \lim_{t\to\infty} \int_{\frac{1}{2}D_{t+s}}^{D_{t+s}(1-w_{t+s})} A(x)\,dx - \int_{\frac{1}{2}D_t}^{D_t(1-w_t)} A(x)\,dx = 0, \quad \mathbb{P}-\text{a.s.} \tag{12}$$

3 Examples

In this section we use a series of examples to illustrate how survival and price impact properties depend on the interplay of the model primitives and to provide basic intuition for the more general results in the next section. Our examples are organized in four sets. The first three sets of examples compare economies differing from each other with respect to only one of the primitives, namely, beliefs, endowment, or preferences. The last set of examples focuses on state-dependence of preferences, such as habit formation.

3.1 Beliefs

Our first set of examples illustrates extinction and survival in an economy with two Bayesian learners.

Example 3.1 *Consider a continuous-time economy with the aggregate endowment given by a geometric Brownian motion:*

$$\frac{dD_t}{D_t} = \mu\, dt + \sigma\, dZ_t, \quad D_0 > 0.$$

Assume that the two agents have logarithmic preferences: $U(c) = \ln(c)$. *Assume that the agents do not know the growth rate of the endowment process. They start with a Gaussian prior belief about* μ, $\mathcal{N}(\hat{\mu}^i, (\nu^i)^2)$, $i \in \{\mathbb{A}, \mathbb{B}\}$, *and update their beliefs based on the observed history of the endowment process according to the Bayes rule. Then, if both agents have non-degenerate priors,* $\min(\nu^{\mathbb{A}}, \nu^{\mathbb{B}}) > 0$, *then both agents survive in the long run. If agent* $\mathbb{A}$ *knows the exact value of the endowment growth rate but agent* $\mathbb{B}$ *does not, i.e.,* $\nu^{\mathbb{B}} > \nu^{\mathbb{A}} = 0$, *then agent* $\mathbb{B}$ *fails to survive.*

In the above example, both agents' beliefs tend to the true value of the unknown parameter μ asymptotically. What determines survival is the rate of learning. If both agents start not knowing the true value of μ, then, regardless of the bias or precision of their prior, they both learn at comparable rates. Formally, the ratio of the agents' belief densities converges to one, $\lim_{t\to\infty} \xi_t = 1$. However, if one agent starts with perfect knowledge of the true parameter value, the rate of learning of the other agent is not sufficient to guarantee that agent's survival.[5] Formal derivations are presented in the Appendix.

Our second example is motivated by Dumas, Kurshev, and Uppal (2008), who study an economy with an irrational (overconfident) agent who fails to account for noise in his signal during the learning process. We do not model the learning process of the overconfident agent explicitly, as Dumas, Kurshev, and Uppal (2008) do, but instead postulate a qualitatively similar belief process exogenously.

Example 3.2 *Consider a discrete-time economy with the aggregate endowment given by*

$$D_t = D_{t-1} \exp\left(\mu_{t-1} + \sigma \epsilon_t\right), \quad D_0 > 0,$$

[5]See Blume and Easley (2006) for further discussion of Bayesian learning and its implications for survival.

where ϵ_t are i.i.d with $\epsilon_t \sim N(0,1)$ and the conditional growth rate of the endowment, μ_{t-1} is a stationary moving average of the shocks $\epsilon_{t-1}, \epsilon_{t-2}, \ldots$. Assume that agent $\mathbb{A}$ knows the true value of μ_t, but agent $\mathbb{B}$, who is an overconfident learner, observes it with noise. Specifically, agent $\mathbb{B}$'s estimate of the current growth rate of the endowment is given by $\mu_{t-1} + \delta_{t-1}$, where δ_t follows a stationary moving average process driven by an independent series of standard normal random variables u_t. Assume that both agents have logarithmic preferences. Then agent $\mathbb{B}$ fails to survive in the long run.

Mathematically, the moving-average representation of the expected growth rate of the endowment in the above example captures the special case of a Bayesian learner confronted with an unobservable expected growth rate following a low-order autoregressive process, which is the setting considered in Dumas, Kurshev, and Uppal (2008). Thus, one can interpret the objective distribution of the endowment process in our example as the beliefs of a Bayesian learner $\mathbb{A}$. In contrast, agent $\mathbb{B}$ does not follow the Bayes rule, and we capture his mistakes by adding noise to his forecasts of endowment growth. Agent $\mathbb{B}$'s errors follow a stationary process and thus do not diminish over time. As we show in the Appendix, agent $\mathbb{B}$'s beliefs diverge from the Bayesian learner's probability measure asymptotically, $\lim_{t\to\infty} \xi_t = 0$, and therefore he fails to survive.

3.2 Endowments

The following set of examples illustrates the dependence of survival and price impact results on the endowment process. We consider three economies with identical preferences and beliefs but different endowment processes. In our first example agent, $\mathbb{B}$ survives and affects prices. In the second example, he does not survive and has no price effect. In the last example, agent $\mathbb{B}$ fails to survive but does exert a long-run impact on prices.

Consider a continuous-time economy with uncertainty described by a Brownian Motion Z_t. Both agents have the same utility function described by the absolute risk aversion function

$$A(x) = \begin{cases} x^{-1} & 0 < x \leq 1 \\ \\ x^{-\alpha} & x > 1 \end{cases} \tag{13}$$

where $0 \leq \alpha < 1$. The utility function thus defined exhibits constant relative risk aversion at low levels of consumption with increasing relative risk aversion at high levels of consumption. We assume that the two agents have constant disagreement δ about the drift of the Brownian motion Z, and therefore the difference in agents' beliefs is described by the density process

$$\xi_t = \exp\left(-\frac{1}{2}\delta^2 t + \delta Z_t\right). \tag{14}$$

Agents' beliefs thus diverge asymptotically, with $\lim_{t\to\infty} \xi_t = 0$. We set the relative utility weight λ to one. The relevant endowment processes are specified in each of the following examples.

It is critical that in the endowment examples, as the economy grows, agents' risk aversion increases. As we show in Corollary 4.2 to Theorem 4.1, agent $\mathbb{B}$ does not survive in an economy with bounded relative risk aversion. If relative risk aversion of the agents is not bounded, one can observe both survival and price impact in the long run, as the following examples illustrate. As we show formally in Sections 4 and 5, survival and price impact results depend on the relation between the asymptotic growth rate of agents' risk aversion relative and the rate of divergence of their beliefs. Our three examples represent the cases when risk aversion rises faster, slower, or at the same rate that the agents' beliefs diverge.

Our first example illustrates that agent $\mathbb{B}$ may survive if risk aversion in the economy rises fast enough compared to the rate at which agents' beliefs diverge. In particular, the agent with the higher consumption share is also more risk averse asymptotically, which explains why agent $\mathbb{A}$ does not dominate the economy in the long run. Informally, as $\mathbb{A}$ pulls ahead of $\mathbb{B}$ in his consumption share, he also becomes sufficiently more risk averse to allow $\mathbb{B}$ to catch up. Thus, the two agents increase their consumption at the same asymptotic rate.

Example 3.3 *Let the endowment process be a Geometric Brownian motion with a positive* *drift,*

$$\frac{dD_t}{D_t} = \mu dt + \sigma dZ_t, \quad \mu > 0, \quad D_0 > 0. \tag{15}$$

Then agent $\mathbb{B}$ survives in the long run, moreover, his consumption share approaches 1/2 *asymptotically. Agent $\mathbb{B}$ exerts long-run impact on prices.*

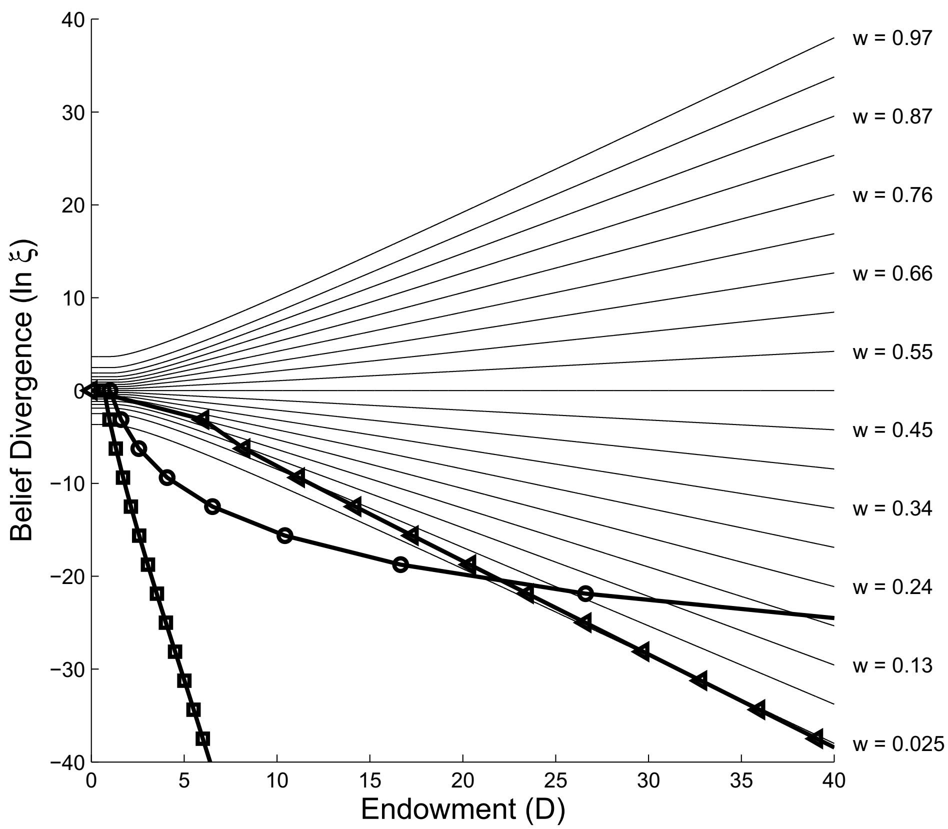

Figure 1: Survival. This figure illustrates survival results in Examples 3.3 through 3.5. Aggregate endowment D is plotted on the horizontal axis, while belief divergence, $\ln(\xi)$ is plotted on the vertical axis. Solid lines are the level curves for the consumption share of agent $\mathbb{B}$, so that each solid line plots pairs $(D, \ln(\xi))$ that give rise to a given consumption share w. These pairs can be found by fixing w and plotting the value of $\ln(\xi)$ as a function of D, with the function given by the Pareto optimality condition (7). Labels for agent $\mathbb{B}$'s consumption share are shown alone the right margin. The marked lines show the median path ($Z_t = Z'_t = 0$) of $(D_t, \ln(\xi_t))$ for each example. The paths corresponding to Examples 3.3, 3.4, and 3.5 are marked with circles, squares, and triangles respectively. We choose the following parameter values: $\alpha = 0$, $\lambda = 1$, $\mu = 0.02$, $\sigma = 0.05$, $\delta = 0.5$. In Example 3.5, we set $X_0 = 1.5$.

Figure 1 shows the *median path* of the economy in Example 3.3, plotted against level curves for the consumption share of agent $\mathbb{B}$ (solid lines). Each level curve represents pairs

$(D, \ln(\xi))$ that give rise to a particular consumption share w. These lines depend only on preferences, and they can be found by fixing w and plotting the value of $\ln(\xi)$ as a function of D, with the function given by the Pareto optimality condition (7). We assume that $\alpha = 0$ so that the agents' utilities have constant absolute risk aversion when consumption is greater than one. This assumption explains why the level curves tend to be spaced wider at higher endowment levels. Because of constant absolute risk aversion at high consumption levels, a given difference in consumption shares requires a larger difference in beliefs at higher endowment levels. In addition, relative risk aversion is asymptotically proportional to consumption, and therefore the growth rate of risk aversion in this economy is the same as the growth rate of the endowment process.

The median path of the economy in Example 3.3 is obtained by setting the driving Brownian motion Z_t equal to zero, and it is shown by a line marked with circles. The economy illustrated in Figure 1 is growing, since $\mu - \sigma^2/2 > 0$, and so the median path is traced from left to right as time passes. We can see that the median path crosses the level curves from below (for large t), which shows that as the economy grows over time, the consumption share of agent $\mathbb{B}$ along the median path increases.

In our next example, aggregate endowment grows slower than in the previous example, and therefore risk aversion of the two agents rises sufficiently slowly compared to the divergence rate of their beliefs that agent $\mathbb{B}$ becomes extinct and has no price impact.

Example 3.4 *Let the endowment process be given by $D_t = \ln(1 + \exp(X_t))$, where X_t is an arithmetic Brownian motion with a positive drift, $X_t = \mu t + \sigma Z_t$, $\mu > 0$. Then agent $\mathbb{B}$ does not survive and has no price impact in the long run.*

In Figure 1, the median path for Example 3.4 (marked with squares) crosses consumption-share level curves from above, showing that as the economy grows, agent $\mathbb{B}$'s consumption share vanishes. Since the rate for belief divergence is identical in all three examples in this section, the only reason why the median paths in our examples have different slopes is because of the different growth rates of the endowment process. Slow endowment growth (and hence slow growth of risk aversion) generates a steep median path, leading to agent $\mathbb{B}$'s extinction.

Together, Examples 3.3 and 3.4 suggest that agent $\mathbb{B}$'s survival depends on the tradeoff between the rate of divergence of beliefs in the market and the growth rate of risk aversion.

Our last example in this set illustrates that survival and price impact are distinct concepts. In this example, risk aversion grows at an intermediate rate: slower than in Example 3.3 but faster than in Example 3.4. With the appropriate choice of the endowment process, we demonstrate that an agent can exert long-term price impact despite of becoming extinct. In addition, we highlight the challenges in obtaining a sharp characterization of economies satisfying the market selection hypothesis: economies for which agents making relatively inaccurate forecasts have no long-run impact on prices.

Example 3.5 *Let X_t be a positive stationary process*

$$dX_t = (X_t - a)(X_t - b)\, dZ'_t, \quad X_0 \in (a, b) \tag{16}$$

where $b > a > 0$ and Z'_t is a Brownian motion independent of Z_t. Assume that the aggregate endowment is given by

$$D_t = \left(\frac{|\ln \xi_t|}{(1-\alpha)^{-1} - X_t^{1-\alpha} |\ln \xi_t|^{\alpha - 1}} \right)^{\frac{1}{1-\alpha}}. \tag{17}$$

Then agent $\mathbb{B}$ becomes extinct asymptotically but maintains long-run impact on prices.

If $\alpha = 0$ and $\mathbb{A} = \mathbb{P}$, then the maximum achievable instantaneous Sharpe ratio (the standard deviation of the state price density) is asymptotically equal to $\left(\delta^2 + \frac{1}{4}(X_t - a)^2 (X_t - b)^2\right)^{\frac{1}{2}}$. In contrast, in the benchmark homogeneous-beliefs economy with $\lambda = 0$, the maximum Sharpe ratio is asymptotically equal to $\left(\delta^2 + (X_t - a)^2 (X_t - b)^2\right)^{\frac{1}{2}}$. Thus, the price impact of agent $\mathbb{B}$'s relatively inaccurate forecasts creates persistent changes in the investment opportunity set.

The construction of Example 3.5 formalizes our intuitive discussion of price impact in the introduction. For t large, $w_t \approx (1-\alpha)^{\frac{1}{1-\alpha}} \frac{X_t}{|\ln \xi_t|}$, which is close to zero because $|\ln \xi_t|$ converges to infinity. However, $w_t D_t$ is volatile enough to affect the volatility of the marginal utility of agent $\mathbb{A}$, $U'(D_t - w_t D_t)$. Because the agents are sufficiently risk averse, a small but variable consumption share of agent $\mathbb{B}$ translates into variable prices, generating price impact without survival.

In Figure 1, the median path for Example 3.5 (marked with triangles) crosses consumption-share level curves from above, showing that as the economy grows, agent $\mathbb{B}$'s consumption share vanishes. The difference between Examples 3.4 and 3.5 is in the rate at which agent $\mathbb{B}$'s consumption share vanishes. This process is slower in Example 3.5, and the relatively slow rate of extinction allows agent $\mathbb{B}$ to retain impact on prices in the long run.

3.3 Preferences

We now illustrate how survival depends on preferences. We consider a family of economies with state-independent preferences that differ only with respect to the agents' utility function.

Example 3.6 *Consider a continuous-time economy with the aggregate endowment given by a geometric Brownian motion:*

$$\frac{dD_t}{D_t} = \mu\, dt + \sigma\, dZ_t, \quad D_0 > 0, \quad \mu, \sigma > 0.$$

Assume that agent $\mathbb{A}$ uses the correct probability measure, $\mathbb{A} = \mathbb{P}$, but agent $\mathbb{B}$ has a constant bias, $\delta \neq 0$, in his forecasts of the growth rate of the endowment:

$$\xi_t = \exp\left(-\frac{\delta^2}{2\sigma^2} t + \frac{\delta}{\sigma} Z_t\right).$$

Let the relative risk aversion function of the two agents be given by

$$\gamma(x) = (1+x)^\alpha, \quad \alpha \leq 1.$$

Then, if relative risk aversion is non-increasing, $\alpha \leq 0$, agent $\mathbb{B}$ does not survive and does not *affect prices asymptotically. If agents' preferences exhibit increasing relative risk aversion, $\alpha \in (0,1]$, agent $\mathbb{B}$ survives and has price impact in the long run.*

In this family of preferences, relative risk aversion is decreasing for negative values of α, increasing for positive values, and constant for $\alpha = 0$. The example shows that the case of constant relative risk aversion (e.g., Yan (2008)) is a knife-edge case: agent $\mathbb{B}$ becomes extinct if risk aversion is decreasing or constant, and survives otherwise. In particular, any positive rate of growth in risk aversion over time can generate survival.

4 Survival

In this section we present general necessary and sufficient conditions for survival. The following theorem shows formally that survival depends on how the asymptotic rate of growth of aggregate relative risk aversion compares to the rate of belief divergence.

Theorem 4.1 *A necessary condition for agent $\mathbb{B}$ to become extinct is that for all $\epsilon \in (0, \frac{1}{2})$,*

$$\limsup_{t\to\infty} \frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{|\ln(\lambda\xi_t)|} \leq 1, \quad \mathbb{P}-\text{a.s.} \tag{18}$$

A sufficient condition for his extinction is that the inequality is strict, i.e., for all $\epsilon \in (0, \frac{1}{2})$,

$$\limsup_{t\to\infty} \frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{|\ln(\lambda\xi_t)|} < 1, \quad \mathbb{P}-\text{a.s.} \tag{19}$$

From the conditions in Theorem 4.1, it is clear that survival depends on the joint properties of aggregate endowments (D_t), preferences (in particular, risk aversion $A(x)$), and beliefs (ξ_t). Survival is determined by the relation between the growth rates of belief divergence and risk aversion. Theorem 4.1 formalizes the informal discussion in the introduction. If risk aversion grows too rapidly, the numerator in (18) dominates and agent $\mathbb{B}$ survives. Intuitively, the numerator captures the relation between differences in consumption and differences in marginal utilities between the two agents. The Pareto optimality condition (6) implies that if beliefs of the two agents diverge, so must their marginal utilities evaluated at their equilibrium consumption. But if risk aversion grows too fast, increasing differences in marginal utilities fail to generate large differences in consumption. Equation (18) provides the precise restriction on the growth rate of risk aversion necessary for agent $\mathbb{B}$'s extinction.

The following straightforward applications of Theorem 4.1 identify a broad class of models in which agent $\mathbb{B}$ does not survive. These results replace the joint condition on the primitives in Theorem 4.1 with an easily verifiable condition on only one of the primitives: the utility function in Corollary 4.2 and the endowment process in Corollary 4.3.

Corollary 4.2 *If relative risk aversion is bounded and* $\lim_{t\to\infty} \xi_t = 0$*, then the* $\mathbb{B}$ *agent never survives.*

If risk aversion is bounded, large differences in marginal utilities imply large differences in consumption, therefore when beliefs diverge, agent $\mathbb{B}$ does not survive. The class of models with bounded relative risk aversion is quite large. It includes, for instance, all utilities of HARA (hyperbolic absolute risk aversion) type, except for the CARA (constant absolute risk aversion) utility.

If the endowment process is bounded, so will be the level of risk aversion, therefore models with bounded endowments have the same survival properties as the models with bounded risk aversion. Sandroni (2000) and Blume and Easley (2006) study models with bounded endowment and diverging beliefs and find that agent $\mathbb{B}$ fails to survive regardless of the exact form of preferences. We replicate this result as a consequence of Theorem 4.1.

Corollary 4.3 *If the aggregate endowment process is bounded away from zero and infinity and* $\lim_{t\to\infty} \xi_t = 0$*, then the* $\mathbb{B}$ *agent never survives.*

If the endowment process is not bounded (away from zero or away from infinity), then the precise relation between the primitives is important in determining agent $\mathbb{B}$'s survival. We simplify the conditions of Theorem 4.1 for the class of utilities with decreasing absolute risk aversion (DARA), which is generally considered to be the weakest a priori restriction on utility functions.

Proposition 4.4 *Suppose that the utility function exhibits DARA. Then, for the $\mathbb{B}$ agent to go extinct it is sufficient that there exists a sequence $\epsilon_n \in (0, \frac{1}{2})$ converging to zero such that for any n*

$$\lim_{t\to\infty} \frac{\gamma(\epsilon_n D_t)}{|\ln \xi_t|} = 0, \quad \mathbb{P}-\text{a.s.} \tag{20}$$

For the $\mathbb{B}$ agent to survive, it is sufficient that for some $\epsilon \in (0, \frac{1}{2})$

$$\limsup_{t\to\infty} \frac{\gamma(\epsilon D_t)}{|\ln \xi_t|} = \infty, \quad \mathbb{P}-\text{a.s.} \tag{21}$$

If, in addition,

$$\lim_{t\to\infty} \frac{\gamma(D_t)}{|\ln \xi_t|} = \infty, \quad \mathbb{P}-\text{a.s.} \tag{22}$$

then $\lim_{t\to\infty} w_t = \frac{1}{2}$, $\mathbb{P}-$a.s.

This result shows that all the survival-relevant information about preferences is captured by the relative risk aversion coefficient evaluated at the constant fractions of aggregate endowment. This formally defines the asymptotic growth rate of risk aversion.

To illustrate how the model primitives jointly determine survival, consider a special family of DARA utilities with unbounded relative risk aversion. Agent $\mathbb{B}$ may or may not survive, as shown in the following corollary of Theorem 4.1.

Corollary 4.5 *Assume that relative risk aversion satisfies*

$$\gamma(x) = k_1 + k_2 x^a, \quad k_1, k_2 > 0, \quad a \in [0, 1].$$

Then agent $\mathbb{B}$ survives if $\lim_{t\to\infty} \frac{D_t^a}{|\ln \xi_t|} = \infty$ *and becomes extinct if* $\lim_{t\to\infty} \frac{D_t^a}{|\ln \xi_t|} = 0$.

To ensure survival, the endowment must grow rapidly enough that the aggregate risk aversion (D_t^a) outpaces the divergence of beliefs ($\ln|\xi_t|$).

We finally consider a generalization of the setting analyzed in Kogan, Ross, Wang, and Westerfield (2006) and Yan (2008), where endowment follows a Geometric Brownian motion and agent $\mathbb{B}$ is persistently optimistic about the growth rate of the endowment. We make a weaker assumption that the endowment and belief differences grow at the same asymptotic rate, i.e. $\lim_{t\to\infty} \frac{\ln D_t}{|\ln \xi_t|} = b < \infty$.

Corollary 4.6 *Consider an economy with* $\lim_{t\to\infty} \frac{\ln D_t}{|\ln \xi_t|} = b < \infty$ *and* $\lim_{t\to\infty} \xi_t = 0$, $\mathbb{P}-$a.s. *Assume that the utility function is of DARA type. Then agent* $\mathbb{B}$ *becomes extinct if*

$$\lim_{x\to\infty} \frac{\gamma(x)}{k_1 + k_2 \ln(1+x)} = 0. \tag{23}$$

for some positive constants k_1 *and* k_2*. The* $\mathbb{B}$ *agent survives if*

$$\lim_{x\to\infty} \frac{\gamma(x)}{k_1 + k_2 \ln(1+x)} = \infty, \tag{24}$$

We thus identify two broad classes of preferences for which survival does and does not take place under the above assumption on the endowment and beliefs. Agent $\mathbb{B}$ becomes extinct if risk aversion at high consumption levels grows slower than logarithmically, and he survives if risk aversion grows faster than logarithmically. This, again, illustrates the interplay between assumptions about endowments, beliefs, and preferences. For instance, if we leave endowments and beliefs unrestricted, then we can have extinction of agent $\mathbb{B}$ for relative risk aversion with faster-than-logarithmic growth, as in Corollary 4.5. However, this is not possible if the endowment and belief differences grow at the same rate, as in Corollary 4.6.

5 Price Impact

We now consider the influence agent $\mathbb{B}$ has on the long-run behavior of prices and how this influence is related to his survival. As pointed out by Kogan, Ross, Wang, and Westerfield (2006), survival and price impact are two separate concepts. We show that in general survival is neither a necessary nor a sufficient condition for price impact. We then provide examples of sufficient and necessary conditions for price impact.

As we have seen, when beliefs diverge and endowments are bounded from above and below, agent $\mathbb{B}$ does not survive (see Corollary 4.3). In this case, agent $\mathbb{B}$ also has no price impact in the long run.

Proposition 5.1 *If relative risk aversion is bounded and* $\lim_{t\to\infty} \xi_t = 0$*, then the* $\mathbb{B}$ *agent has no price impact.*

The relation between Arrow-Debreu prices and marginal utilities means that agent $\mathbb{B}$ can affect prices if he has nontrivial impact on the marginal utility of agent $\mathbb{A}$. When risk aversion is bounded, this requires him to have significant impact on consumption growth of agent $\mathbb{A}$ asymptotically, which is impossible since agent $\mathbb{B}$ does not survive.

As with survival, the result for bounded risk aversion is similar to that of a bounded endowment:

Proposition 5.2 *If aggregate endowment is bounded above and below away from zero and* $\lim_{t\to\infty} \xi_t = 0$*, then agent* $\mathbb{B}$ *has no price impact.*

The above propositions show that for a broad class of economies agent $\mathbb{B}$ agent does not survive and does not exert price impact in the long run.

When endowments or risk aversion are unbounded, the situation becomes more complicated. Now it is possible for agent $\mathbb{B}$ to affect the marginal utility of agent $\mathbb{A}$ without having nonvanishing effect on agent $\mathbb{A}$'s consumption growth, which means that price impact may be consistent with agent $\mathbb{B}$'s extinction. We have described such an economy in Example

3.5. We show below that the reverse is also possible: agent $\mathbb{B}$ may survive without affecting prices. But first, we formulate a set of sufficient conditions under which agent $\mathbb{B}$ both survives and affects prices.

Proposition 5.3 *Consider an economy with DARA preferences in which*

$$\lim_{t\to\infty} \frac{\gamma(\frac{1}{2}D_t)}{(\ln \xi_t)^2} = \infty, \quad \mathbb{P}-\text{a.s.} \tag{25}$$

and $\lim_{t\to\infty} \xi_t = 0$. *Then agent* $\mathbb{B}$ *survives and asymptotically consumes a half of the aggregate endowment.*

As a by-product of our analysis, we obtain a general result for consumption sharing rules in economies with unbounded relative risk aversion: if beliefs are homogeneous, the asymptotic consumption distribution is independent of the initial wealth distribution.

Proposition 5.4 *Consider an economy with homogeneous beliefs, growing endowment process,* $D_t \to \infty$ *as* $t \to \infty$, *and monotonically increasing, unbounded relative risk aversion* $\gamma(x)$. *Then, for any initial allocation of wealth between the agents, their consumption shares become asymptotically equal. Moreover, the state price density in this economy is asymptotically the same as in an economy in which the agents start with equal endowments.*

We now state the necessary and sufficient conditions for price impact.

Proposition 5.5 *In the economy defined in Proposition 5.3, agent* $\mathbb{B}$ *has long-run price impact if and only if the belief process* ξ_t *has non-vanishing growth rate asymptotically, i.e., there exists* $s > 0$ *and* $\epsilon > 0$ *such that*

$$\text{Prob}\left[\limsup_{t\to\infty} |\ln \xi_{t+s} - \ln \xi_t| > \epsilon\right] > 0. \tag{26}$$

Moreover, asymptotically the state price density does not depend on the initial wealth distribution, i.e., does not depend on λ.

Proposition 5.5 shows that survival does not necessarily lead to price impact. Since under condition (25) the asymptotic consumption allocation does not depend on the initial wealth distribution (each agent consumes one half of the aggregate endowment), in the long run, the state price density does not depend on the initial wealth distribution.

Proposition 5.4 sheds light on why the survival and price impact properties of various economies are connected to whether risk aversion of the agents is bounded or increasing. Economies with bounded relative risk aversion exhibit simple behavior: agent $\mathbb{B}$ does not survive and has no asymptotic impact on the state-price density. When relative risk aversion is increasing and therefore unbounded, we know from Proposition 5.4 that in a homogeneous-belief economy consumption shares of the agents tend to become equalized over time no matter how uneven the initial wealth distribution is. Similarly, the state-price density does not depend (asymptotically) on the initial wealth distribution. This mechanism remains at work in economies with heterogeneous beliefs. However, there is another force present now: agent $\mathbb{B}$ tends to mis-allocate his consumption across states due to his distorted beliefs, which reduces his asymptotic consumption share. The tradeoff between these two competing forces is intuitive: distortions in consumption shares caused by belief differences tend to disappear over time, unless the belief differences grow sufficiently rapidly. This explains the result in (26)

Proposition 5.5 completes our taxonomy of models with respect to survival and price impact. Corollaries 4.2 and 4.3 and Propositions 5.1 and 5.2 show general conditions under which there is neither survival nor price impact. Example 3.5 in Section 3.2 describes an economy with price impact but no survival. Completing the set, Proposition 5.5 describes economies in which there is survival both with and without price impact.

6 State-Dependent Preferences

In this section, we generalize our results on survival and price impact to models with state-dependent preferences. Let the utility function take the form $u(C, H)$, where C is agent's consumption and H is the process for state variables affecting the agent's utility. We assume that H is an exogenous adapted process. This specification covers models of external habit formation, or catching-up-with-the-Joneses preferences, as in Abel (1990) and Campbell and Cochrane (1999), in which case the process H is a function of lagged values of the aggregate endowment.

Theorem 4.1 extends to the case of state-dependent preferences. Let $A(C, H)$ and $\gamma(C, H)$ denote, respectively, the coefficients of absolute and relative risk aversion at consumption level C. Then we obtain an analog of Proposition 4.4.

Proposition 6.1 *Assume that the utility function $u(C, H)$ exhibits DARA: the coefficient of absolute risk aversion $A(C, H)$ is decreasing in C. Then, for agent $\mathbb{B}$ to go extinct it is sufficient that there exists a sequence $\epsilon_n \in (0, \frac{1}{2})$ converging to zero such that for any n*

$$\lim_{t\to\infty} \frac{\gamma(\epsilon_n D_t, H_t)}{|\ln \xi_t|} = 0, \quad \mathbb{P}-\text{a.s.} \tag{27}$$

For agent $\mathbb{B}$ to survive, it is sufficient that for some $\epsilon \in (0, \frac{1}{2})$

$$\limsup_{t\to\infty} \frac{\gamma(\epsilon D_t, H_t)}{|\ln \xi_t|} = \infty, \quad \mathbb{P}-\text{a.s.} \tag{28}$$

If, in addition,

$$\lim_{t\to\infty} \frac{\gamma(D_t, H_t)}{|\ln \xi_t|} = \infty, \quad \mathbb{P}-\text{a.s.} \tag{29}$$

then $\lim_{t\to\infty} w_t = \frac{1}{2}, \quad \mathbb{P}-\text{a.s.}$

In a growing economy ($D_t \to \infty$, $\mathbb{P} - a.s.$) with diverging beliefs and state-independent preferences, survival of agent $\mathbb{B}$ requires that the coefficient of relative risk aversion is unbounded at large levels of consumption, namely, that $\limsup_{x\to\infty} \gamma(x) = \infty$ (see Proposition 4.4). This is not the case if preferences are state-dependent. In many common models of external habit formation, the process H_t is such that the process for relative risk aversion in the economy is stationary. In such cases, survival and price impact results are sensitive to the distributional assumptions on the beliefs and endowment. In order for agent $\mathbb{B}$ to survive, the stationary distribution of risk aversion must have a sufficiently heavy right tail, so that the condition (27) is violated. As an illustration, consider the following example of two economies which differ only with respect to the distribution of endowment growth.

Example 6.2 *Consider two discrete-time economies with external habit formation. Let the relative risk aversion coefficient of the two agents be*

$$\gamma(x, H) = 1 + \left(\frac{x}{H}\right),$$

where $H_t = D_{t-1}$. Thus, agents' external habit level equals the lagged value of the aggregate endowment. Assume that the disagreement process follows

$$\ln \xi_t = -\frac{1}{2}t + \sum_{n=1}^{t} Z_n$$

where Z_n are distributed according to a standard normal distribution and are independent of the endowment process. Endowment growth is independently and identically distributed over time in both economies. Assume that the endowment process D_t is independent of the disagreement process ξ_t, which means that agents $\mathbb{A}$ and $\mathbb{B}$ disagree on probabilities of payoff-irrelevant states.[6][7] *In the first economy, endowment growth has bounded support, $0 < \underline{g} \leq \frac{D_t}{D_{t-1}} \leq \overline{g} < \infty$. In the second economy, endowment growth is bounded from below but unbounded from above. Moreover, the distribution of endowment growth is heavy-tailed, namely, there exists a positive constant a such that, for sufficiently large x,*

$$\text{Prob}\left[\frac{D_t}{D_{t-1}} > x\right] > ax^{-1/3}.$$

Then, agent $\mathbb{B}$ becomes extinct in the first economy, and survives in the second economy.

[6] Another example of an economy in which belief differences are independent of the aggregate endowment is a multi-sector economy in which agents agree on the distribution of the aggregate endowment, but disagree about the distribution of sectors' shares in the aggregate endowment.

[7] Survival results in this example do not depend on the joint distribution of the endowment process and the disagreement process. Thus, one may assume that the two agents disagree about the probabilities of payoff-relevant states by specifying $\frac{D_t}{D_{t-1}}$ to be a nonlinear function of Z_t. The assumption of independence of endowment and beliefs makes it easy to establish price impact results below.

In the first economy, it is clear that condition (27) is satisfied for any positive ϵ_n, and thus agent $\mathbb{B}$ does not survive. In the second economy, the distribution of endowment growth is such that relative risk aversion exhibits frequent large spikes, namely

$$\text{Prob}\left(\gamma(\epsilon D_t, H_t) > t^3\right) \geq \text{Prob}\left(\epsilon \frac{D_t}{D_{t-1}} > t^3\right) > a\epsilon^{1/3} t^{-1}$$

Such spikes in risk aversion occur frequently enough that the condition (28) holds. [8]

The following propositions extends our results on price impact to economies with state-dependent preferences. Their proofs follow closely the results of Sections 4 and 5.

Proposition 6.3 *There is no price impact or survival in models with bounded relative risk aversion.*

In the model with state-independent preferences, bounding the endowment implied bounding relative risk aversion. This, in turn, implied a lack of price impact. With state-dependent preferences, a bounded endowment no longer implies that relative risk aversion is bounded.

Proposition 6.4 *Consider a model with the utility function of DARA type, and let the coefficient of relative risk aversion be monotonically increasing in its first argument. Assume that*

$$\lim_{t\to\infty} \frac{\gamma(\frac{1}{2}D_t, H_t)}{(\ln \xi_t)^2} = \infty, \quad \mathbb{P}-\text{a.s.}, \tag{30}$$

Then agent $\mathbb{B}$ survives and asymptotically consumes a half of the aggregate endowment. He has price impact if and only if the disagreement process ξ_t is such that its growth rate does not vanish asymptotically, i.e., there exists $s > 0$ and $\epsilon > 0$ such that

$$\text{Prob}\left[\limsup_{t\to\infty} |\ln \xi_{t+s} - \ln \xi_t| > \epsilon\right] > 0. \tag{31}$$

[8]Since $\sum_{t=1}^{\infty} t^{-1} = \infty$, the Borel-Cantelli lemma implies that $\limsup_{t\to\infty} \frac{D_t}{D_{t-1}} t^{-3} \geq 1 \ \mathbb{P} - a.s.$ Since $\lim_{t\to\infty} |\ln \xi_t| t^{-3} = 0 \ \mathbb{P} - a.s.$, (28) follows.

Moreover, asymptotically the state-price density does not depend on the initial wealth distribution, i.e., does not depend on λ.

Returning to Example 6.2, note that agent $\mathbb{B}$ has no price impact in the first economy but exerts price impact in the second economy. The first result follows immediately from Proposition 6.4, since bounded dividend growth implies bounded relative risk aversion in this economy. The second result can be established using a slight modification of the proof of Proposition 5.5.[9]

7 Conclusion

In this paper we examine the economic mechanism behind the Market Selection Hypothesis and establish necessary and sufficient conditions for its validity in a general setting with minimal restrictions on endowments, beliefs, or utility functions. We show that the MSH holds in economies with bounded endowments or bounded relative risk aversion. The commonly studied special case of constant relative risk aversion preferences belongs to this class of models. However, we show that the MSH cannot be substantially generalized to a broader class of models. Instead, survival is determined by a comparison of the forecast errors to risk attitudes. The price impact of inaccurate forecasts is distinct from survival because price impact is determined by the volatility of traders' consumption shares rather than by their level. Our results also apply to economies with state-dependent preferences, such as habit formation.

[9] As we show above, $\limsup_{t\to\infty} \frac{D_t}{D_{t-1}} t^{-3} \geq 1$, $\mathbb{P}-$a.s., while $\lim_{t\to\infty} |\ln \xi_t|^2 t^{-3} = 0$, $\mathbb{P}-$a.s., implying that $\limsup_{t\to\infty} \gamma(D_t, H_t)/|\ln \xi_t|^2 = \infty$, $\mathbb{P}-$a.s. The price impact result then follows from independence of D_t and ξ_t and the assumption that increments $\ln \xi_t - \ln \xi_{t-1}$ are independent across time.

A Examples

A.1 Example 3.1

Define $\delta_t^{\mathbb{A}} = \hat{\mu}_t^{\mathbb{A}} - \mu$. Then, using the Kalman Filter,

$$d\delta_t^{\mathbb{A}} = -\delta_t^{\mathbb{A}} \nu_t^{\mathbb{A}} dt + \nu_t^{\mathbb{A}} dZ_t \quad \text{and} \quad d\nu_t^{\mathbb{A}} = -\left(\nu_t^{\mathbb{A}}\right)^2 dt$$

and therefore

$$\delta_t^{\mathbb{A}} = \frac{\delta_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} t + 1} + \frac{\nu_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} t + 1} Z_t$$

Next, from the definition of $\xi_t^{\mathbb{A}}$, we have

$$\begin{aligned}
\ln \xi_t^{\mathbb{A}} &= -\frac{1}{2} \int_0^t \left(\delta_s^{\mathbb{A}}\right)^2 ds + \int_0^t \delta_s^{\mathbb{A}} dZ_s \\
&= -\int_0^t \left(\frac{1}{2} \left(\delta_0^{\mathbb{A}}\right)^2 \frac{1}{(\nu_0^{\mathbb{A}} s + 1)^2} + \frac{1}{2} \left(\nu_0^{\mathbb{A}}\right)^2 \frac{1}{(\nu_0^{\mathbb{A}} s + 1)^2} Z_s^2 + \delta_0^{\mathbb{A}} \nu_0^{\mathbb{A}} \frac{1}{(\nu_0^{\mathbb{A}} s + 1)^2} Z_s \right) ds \\
&\quad + \int_0^t \delta_s^{\mathbb{A}} dZ_s
\end{aligned}$$

In addition, direct integration by parts shows us that

$$\begin{aligned}
\int_0^t \delta_s^{\mathbb{A}} dB_s &= \int_0^t \left(\frac{\delta_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} s + 1} + \frac{\nu_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} s + 1} Z_s \right) dZ_s \\
&= \frac{1}{2} \frac{\nu_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} t + 1} Z_t^2 + \frac{\delta_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} t + 1} Z_t \\
&\quad + \int_0^t \left(-\frac{1}{2} \frac{\nu_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} s + 1} + \frac{1}{(\nu_0^{\mathbb{A}} s + 1)^2} \left(\frac{1}{2} \left(\nu_0^{\mathbb{A}}\right)^2 Z_s^2 + \delta_0^{\mathbb{A}} \nu_0^{\mathbb{A}} Z_s \right) \right) ds
\end{aligned}$$

Plugging the last equation into the expression for $\ln \xi_t^{\mathbb{A}}$ leaves us with

$$\ln \xi_t^{\mathbb{A}} = \frac{1}{2} \frac{\nu_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} t + 1} Z_t^2 + \frac{\delta_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} t + 1} Z_t + \int_0^t \left[-\frac{1}{2} \left(\delta_0^{\mathbb{A}}\right)^2 \frac{1}{(\nu_0^{\mathbb{A}} s + 1)^2} - \frac{1}{2} \frac{\nu_0^{\mathbb{A}}}{\nu_0^{\mathbb{A}} s + 1} \right] ds$$

Since the sum of the first three does not converge to a constant, but the fourth grows as $\ln(t)$, we have that $\nu^{\mathbb{A}} \neq 0$ implies

$$\lim_{t \to \infty} \frac{\xi_t^{\mathbb{A}}}{-2 \ln(t)} = 1$$

with the same result for agent $\mathbb{B}$.

If $\min(\nu^{\mathbb{A}}, \nu^{\mathbb{B}}) > 0$, then $\xi_t \to 1$, and so (7) with $A(x) = \frac{1}{x}$ implies that both agents survive. If $\nu^{\mathbb{B}} > \nu^{\mathbb{A}} = 0$, then $\xi^{\mathbb{B}} = \xi_t \to 0$, and so Proposition 4.4 implies that $\mathbb{B}$ does not survive.

A.2 Example 3.2

Agent B's beliefs are characterized by the density process

$$\xi_t = \exp\left(\sum_{s=1}^{t} -\frac{\eta_{s-1}^2}{2} + \eta_{s-1}\epsilon_s\right)$$

where $\eta_t = \delta_t/\sigma$. The process $M_t = \sum_{s=1}^t \eta_{s-1}\epsilon_s$ is a martingale. Since $\lim_{t\to\infty} \frac{\sum_{s=1}^t \eta_{s-1}^2}{t} = \mathrm{E}[\eta_t^2]$, the quadratic variation process of M_t converges to infinity almost surely under $\mathbb{P}$, and therefore $\lim_{t\to\infty} M_t / \left(\sum_{s=1}^t \eta_{s-1}^2\right) = 0 \ \mathbb{P} - a.s.$ (see Shiryaev 1996, §7.5, Th. 4). This implies that $\lim_{t\to\infty} \xi_t = 0$ $a.s.$ and hence the condition (20) in Proposition 4.4 is satisfied. We conclude that agent $\mathbb{B}$ does not survive in the long run.

A.3 Example 3.3

The sufficient condition for survival (22) in Proposition 4.4 is satisfied, since $\gamma(D_t) = D_t^{1-\alpha}$ grows exponentially, thus, according to the Proposition, agent $\mathbb{B}$ survives with asymptotic consumption share equal to $\frac{1}{2}$. Proposition 5.4 implies that agent $\mathbb{B}$ exerts price impact asymptotically, since the condition (26) is clearly satisfied by the belief process (14) and (25) follows from the exponential growth rate of $\gamma(D_t)$.

A.4 Example 3.4

The Pareto optimality condition (7) cannot be satisfied for $w_t D_t \geq 1$ for large enough t. To see this, assume the contrary. Then, (7) implies that

$$|\ln \xi_t| = \frac{1}{1-\alpha} D_t^{1-\alpha}\left((1-w_t)^{1-\alpha} - w_t^{1-\alpha}\right)$$

which is impossible, since the $|\ln \xi_t|$ increases asymptotically linearly in time, while $D_t^{1-\alpha}$ grows at the rate of $t^{1-\alpha}$. Thus, we conclude that $\limsup_{t\to\infty} w_t D_t \leq 1$ and therefore $\lim_{t\to\infty} w_t = 0 \ \mathbb{P} - a.s.$ and agent $\mathbb{B}$ does not survive.

To show that there is no price impact in this economy, we consider the reference economy with $\lambda^\star = 0$ and verify the condition (11). Since asymptotically $w_t D_t < 1$,

$$\int_{(1-w_t)D_t}^{D_t} A(x)dx = \frac{1}{(1-\alpha)} D_t^{1-\alpha} \left(1 - (1-w_t)^{1-\alpha}\right) = D_t^{1-\alpha}(w_t + o(w_t^2)), \tag{A1}$$

which clearly converges to zero, since $D_t^{1-\alpha} w_t = (w_t D_t) D_t^{-\alpha}$, D_t tends to infinity, and $w_t D_t$ is asymptotically bounded.

A.5 Example 3.5

We look for the solution to (7) under the assumption that $w_t D_t > 1$. Assuming $w_t D_t > 1$, for large enough t, (7) implies that

$$(1-w_t)^{1-\alpha} - w_t^{1-\alpha} = 1 - (1-\alpha) X_t^{1-\alpha} |\ln \xi_t|^{\alpha-1}$$

and therefore $w_t \to 0$, $\mathbb{P} - a.s.$ Using the Taylor expansion in w_t around zero, we find that

$$\frac{w_t^{1-\alpha} - (1-\alpha)w_t + o(w_t^2)}{(1-\alpha)X_t^{1-\alpha}|\ln \xi_t|^{\alpha-1}} = 1,$$

which in turn implies that $\lim_{t\to\infty} w_t |\ln \xi_t| / X_t = (1-\alpha)^{1/(1-\alpha)}$, $\mathbb{P} - a.s.$ This implies that agent $\mathbb{B}$ becomes extinct. To complete the first part of the proof we must verify that, asymptotically, $w_t D_t$ exceeds one. This follows immediately since $|\ln \xi_t| D_t^{\alpha-1} \to (1-\alpha)^{-1}$, $\mathbb{P} - a.s.$ and $X_t \geq a > 0$.

To verify that agent $\mathbb{B}$ exerts price impact, we consider separately the case reference economies with $\lambda^\star = 0$ and $\lambda^\star > 0$.

To see that there is price impact in this economy relative to the reference economy with $\lambda^\star = 0$, note that, for large enough t,

$$\int_{(1-w_t)D_t}^{D_t} A(x)dx = \frac{1}{1-\alpha} D_t^{1-\alpha} \left(1 - (1-w_t)^{1-\alpha}\right) = D_t^{1-\alpha}(w_t + o(w_t))$$

From the limiting results established above, we conclude that, asymptotically, $D_t^{1-\alpha} w_t$ behaves as $(1 - \alpha^{\alpha/(1-\alpha)})\, X_t$. Since the process X_t is stationary and has non-vanishing variance, this implies that the condition (9) is violated and hence there is price impact relative to the reference economy with $\lambda^\star = 0$. The case of $\lambda^\star > 0$ follows a similar argument.

A.6 Example 3.6

Survival results follow from Proposition 4.4, since the logarithm of the belief density ratio $\ln \xi_t$ exhibits linear grows, while the aggregate endowment D_t grows exponentially. Price impact results follow from Proposition 5.5.

B Proofs

B.1 Proof of Theorem 4.1

Suppose that the agent with beliefs $\mathbb{Q}$ becomes extinct, i.e., $w_t = \frac{C_{n,t}}{D_t}$ converges to zero almost surely. For each element of the probability space for which w_t vanishes asymptotically, one can find $T(\epsilon)$, such that $w_t < \epsilon$ for any $t > T(\epsilon)$. Since $\int_{wD}^{(1-w)D} A(x)\,dx$ is a decreasing function of w, the first-order condition (7) implies that for all $t > T(\epsilon)$

$$1 = \frac{\int_{w_t D_t}^{(1-w_t)D_t} A(x)\,dx}{-\ln(\lambda \xi_t)} \geq \frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{-\ln(\lambda \xi_t)}.$$

Thus, the desired result follows by applying $\limsup_{t\to\infty}$ to both sides of the inequality.

We now prove the sufficient condition. Consider the elements of the probability space for which $\limsup_{t\to\infty} \frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{|\ln(\lambda \xi_t)|} < 1$ for any $\epsilon > 0$. For each such realization, we can define $T(\epsilon)$ and $\delta > 0$, such that

$$\frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{|\ln(\lambda \xi_t)|} \leq 1 - \delta$$

for all $t > T(\epsilon)$. If $\limsup_{t\to\infty} w_t \neq 0$, then one can always find $\epsilon > 0$ and $t > T(\epsilon)$, such that $w_t > \epsilon$. But then

$$1 = \frac{\int_{w_t D_t}^{(1-w_t)D_t} A(x)\,dx}{-\ln(\lambda \xi_t)} \leq \frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{-\ln(\lambda \xi_t)}.$$

Taking $\limsup_{t\to\infty}$ on both sides, implies $1 \leq 1 - \delta$, which is a contradiction.

B.2 Proof of Corollary 4.2

Let $\gamma(x) = xA(x) < \bar{\gamma}$ for all x and $\bar{\gamma} > 0$. Then,

$$\frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{|\ln(\lambda\xi_t)|} \leq \bar{\gamma}\frac{\ln\epsilon - \ln(1-\epsilon)}{|\ln(\lambda\xi_t)|}$$

which converges to zero almost surely as $t \to \infty$.

B.3 Proof of Corollary 4.3

Let D_m and D_M denote the upper and lower bound of D_t. Then, $0 < D_m \leq D_M$. Let $\bar{A}$ denote the max of $A(x)$ on $[D_m,\, D_M]$. We then have $\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)dx \leq (1-2\epsilon)(D_M - D_m)\bar{A}$ which is finite. Given that $\xi_t \to 0$ as $t \to \infty$ and hence $|\ln \xi_t|$ goes to infinity, we immediately conclude that (19) holds, and agent $\mathbb{B}$ does not survive.

B.4 Proof of Proposition 4.4

Since $A(x)$ is a non-increasing function,

$$\int_{\epsilon' D}^{(1-\epsilon')D} A(x)\,dx \geq \int_{\epsilon' D}^{\epsilon D} A(x)\,dx \geq A(\epsilon D)D(\epsilon - \epsilon') = \gamma(\epsilon D)\frac{\epsilon - \epsilon'}{\epsilon},$$

where $0 < \epsilon' < \epsilon$. Condition (21) then implies that

$$\limsup_{t\to\infty} \frac{\int_{\epsilon' D_t}^{(1-\epsilon')D_t} A(x)\,dx}{|\ln(\lambda\xi_t)|} = \infty, \quad \mathbb{P}-\text{a.s.},$$

and hence a necessary condition for extinction is violated. Thus, agent $\mathbb{B}$ survives.

Next, for any $\epsilon \in (0, 1/2)$, find $\epsilon_n < \epsilon$. Then, since since $A(x)$ is a non-increasing function,

$$\int_{\epsilon D}^{(1-\epsilon)D} A(x)\,dx \leq A(\epsilon D)D(1-2\epsilon) \leq A(\epsilon_n D)D(1-2\epsilon)$$

Then,

$$\frac{\int_{\epsilon D_t}^{(1-\epsilon)D_t} A(x)\,dx}{|\ln(\lambda \xi_t)|} \leq \frac{A(\epsilon_n D_t)\epsilon_n D_t(1-2\epsilon)}{\epsilon_n |\ln(\lambda\xi_t)|} = \frac{\gamma(\epsilon_n D_t)(1-2\epsilon)}{\epsilon_n|\ln(\lambda\xi_t)|},$$

and the result follows from (19).

Lastly, since the utility function is of DARA type, condition (7) implies that

$$|\ln(\lambda\xi_t)| \geq A((1-w_t)D_t)D_t(1-2w_t) \geq \gamma(D)(1-2w_t)$$

and therefore, using condition (22), $\lim_{t\to\infty} w_t = 1/2$.

B.5 Proof of Corollary 4.5

Follows directly from Proposition 4.4.

B.6 Proof of Corollary 4.6

Consider a set (of measure one) of ωs for which $\lim_{t\to\infty} \frac{\ln D_t}{|\ln \xi_t|} = b$ and $\lim_{t\to\infty} \xi_t = 0$. On this set,

$$\begin{aligned}
\lim_{t\to\infty} \frac{\gamma(\epsilon D_t)}{|\ln \xi_t|} &= \lim_{t\to\infty} \frac{k_1 + k_2 \ln(1+\epsilon D_t)}{|\ln \xi_t|} \frac{\gamma(\epsilon D_t)}{k_1 + k_2\ln(1+\epsilon D_t)} \\
&= bk_2 \lim_{t\to\infty} \frac{\gamma(\epsilon D_t)}{k_1 + k_2 \ln(1+\epsilon D_t)} = 0
\end{aligned}$$

for any positive ϵ. Thus, by Propsition 4.4, agent $\mathbb{B}$ becomes extinct as long as the risk aversion coefficient satisfies (23). According to the the same corollary, if the risk aversion coefficient satisfies (24), then then agent $\mathbb{B}$ survives.

B.7 Proof of Propositions 5.1 and 5.2

As we show in corollary 4.2, there is no survival in models with bounded relative risk aversion. Thus, w_t converges to zero almost surely. Consider now the first term in (11). By the mean value theorem, this term equals

$$A(x^\star_{t+s})D_{t+s}w_{t+s} = \gamma(x^\star_{t+s})\frac{D_{t+s}w_{t+s}}{x^\star_{t+s}},$$

for some $x^\star_{t+s} \in [(1-w_{t+s})D_{t+s}, D_{t+s}]$. Since, almost surely, the ratio $\frac{D_{t+s}}{x_{t+s^*}}$ converges to one, w_{t+s} converges to zero, and the relative risk aversion coefficient $\gamma(x^\star_{t+s})$ is bounded, we conclude that the first term in (11) converges to zero. The same argument implies that the second term converges to zero almost surely, and therefore there is no price impact. This proves proposition 5.1. Proposition 5.2 follows from the fact that bounding the endowment implies bounding relative risk aversion.

B.8 Proof of Propositions 5.3, 5.4, and 5.5

Since the utility function is of DARA type, condition (7) implies that

$$|\ln(\lambda\xi_t)| \geq A((1-w_t)D_t)D_t(1-2w_t) \geq \gamma(D_t)(1-2w_t) \tag{A2}$$

and therefore, using condition (25), $\lim_{t\to\infty} w_t = 1/2$. Thus, agent $\mathbb{B}$ survives. This proves Proposition 5.3.

Continuing, condition (6) and $\lim_{t\to\infty}\xi_t = 0$ a.s. imply that there exists a T so that for $t > T$, we have $w_t \leq \frac{1}{2}$ a.s. We will consider in this proof such times $t > T$.

Finally, recall that $A(x)$ is decreasing and $\gamma(x)$ is increasing in x.

Next, to show that there is price impact, we verify that the difference

$$PI(t,s) \equiv \int_{\frac{1}{2}D_{t+s}}^{D_{t+s}(1-w_{t+s})} A(x)\,dx - \int_{\frac{1}{2}D_t}^{D_t(1-w_t)} A(x)\,dx$$

does not converge to zero almost surely. This corresponds to $\kappa = 1$ in the definition of price impact.

First, the upper bound: DARA and condition (7) imply that

$$\int_{\frac{1}{2}D}^{D(1-w)} A(x)\,dx \leq \frac{1}{2}\int_{Dw}^{D(1-w)} A(x)\,dx = \frac{1}{2}|\ln(\lambda\xi)|$$

Second, the lower bound. A preliminary from (A2):

$$|\ln(\lambda\xi)| \geq \gamma\left(D\right)\left(1-2w\right) \geq \gamma\left(\frac{D}{2}\right)2\left(\frac{1}{2}-w\right)$$

implies

$$0 \leq \left(\frac{1}{2} - w\right)^2 \gamma\left(\frac{D}{2}\right) \leq \frac{1}{4}\frac{|\ln(\lambda\xi)|^2}{\gamma\left(\frac{D}{2}\right)} \to 0 \text{ a.s.} \tag{A3}$$

Next, a Taylor expansion shows that

$$\begin{aligned}
\int_{\frac{1}{2}D}^{D(1-w)} A(x)\,dx &= \int_{\frac{1}{2}D}^{D(1-w)} \gamma(x)\frac{1}{x}\,dx \geq \int_{\frac{1}{2}D}^{D(1-w)} \gamma\left(\frac{D}{2}\right)\frac{1}{x}\,dx \\
&\geq \gamma\left(\frac{D}{2}\right)\ln\left(1 + 2\left(\frac{1}{2} - w\right)\right) \\
&\geq \gamma\left(\frac{D}{2}\right)\left[2\left(\frac{1}{2} - w\right) - \frac{2}{\left(1 + 2\left(\frac{1}{2} - w^\star\right)\right)^2}\left(\frac{1}{2} - w\right)^2\right]
\end{aligned} \tag{A4}$$

where $w^\star \in \left[w, \frac{1}{2}\right]$. However, (A3) implies that the last term in the third line of (A4) approaches zero almost surely as t approaches ∞, and so

$$\int_{\frac{1}{2}D}^{D(1-w)} A(x)\,dx \geq 2\gamma\left(\frac{D}{2}\right)\left(\frac{1}{2} - w\right) + o(t) \tag{A5}$$

Next, another Taylor expansion shows that

$$\begin{aligned}
|\ln(\lambda\xi)| &= \int_{Dw}^{D(1-w)} A(x)\,dx \leq A\left(\frac{D}{2}\right)\left(\frac{1}{2} - w\right)D + \int_{Dw}^{\frac{D}{2}} \frac{\gamma\left(\frac{D}{2}\right)}{x}\,dx \\
&\leq 2\gamma\left(\frac{D}{2}\right)\left(\frac{1}{2} - w\right) - \gamma\left(\frac{D}{2}\right)\ln\left(1 - 2\left(\frac{1}{2} - w\right)\right) \\
&\leq 2\gamma\left(\frac{D}{2}\right)\left(\frac{1}{2} - w\right) - \gamma\left(\frac{D}{2}\right)\left[-2\left(\frac{1}{2} - w\right) - \frac{2}{\left(1 - 2\left(\frac{1}{2} - w^\star\right)\right)^2}\left(\frac{1}{2} - w\right)^2\right]
\end{aligned} \tag{A6}$$

where $w^\star \in \left[w, \frac{1}{2}\right]$. However, (A3) implies that the last term in the third line of (A6) approaches zero almost surely as t approaches ∞, and so

$$|\ln(\lambda\xi)| \leq 4\gamma\left(\frac{D}{2}\right)\left(\frac{1}{2} - w\right) + o(t) \tag{A7}$$

We combine (A5) and (A7) to obtain our lower bound

$$\int_{\frac{1}{2}D}^{D(1-w)} A(x)\,dx \geq \frac{1}{2}|\ln(\lambda\xi)| + o(t)$$

We can now impose bounds on price impact:

$$PI(t,s) \geq \frac{1}{2}\left(\ln(\lambda\xi_{t+s}) - \ln(\lambda\xi_t)\right) + o(t) \tag{A8a}$$

$$PI(t,s) \leq \frac{1}{2}\left(\ln(\lambda\xi_{t+s}) - \ln(\lambda\xi_t)\right) + o(t) \tag{A8b}$$

We therefore conclude that there is an asymptotic difference in state-price densities between the original economy and the benchmark economy with $\kappa = 1$ if and only if the increments of $\ln(\lambda\xi_t)$ do not vanish. Moreover, asymptotically, the state-price density does not depend on λ (or κ). This also shows that it is impossible to find a value of κ that would eliminate the differences between the state-price density in our economy and the benchmark. This proves Proposition 5.5.

Proposition 5.4 follows by setting $\xi_t \equiv 1$.

References

Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* 80, 38–42.

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

Blume, Lawrence, and David Easley, 2006, If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets, *Econometrica* 74, 929–966.

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: a consumption based explanation of aggragate stock market behavior, *Journal of Political Economy* 107, 205–251.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldman, 1991, The survival of noise traders in financial markets, *Journal of Business* 64, 1–19.

Dumas, Bernard, Alexander Kurshev, and Raman Uppal, 2008, Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility, *Journal of Finance* forthcoming.

Fedyk, Yuriy, and Johan Walden, 2007, High-speed natural selection in financial markets with large state spaces, Working Paper, SSRN.

Friedman, Milton, 1953, The case for flexible exchange rates, in *Essays in Positive Economics* (University of Chicago Press, Chicago).

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield, 2006, The price impact and survival of irrational traders, *Journal of Finance* 61, 195–229.

Sandroni, Alvaro, 2000, Do markets favor agents able to make accurate predictions?, *Econometrica* 68, 1303–134.

Shiryaev, Albert, 1996, *Probability* vol. 95 of *Graduate Texts in Mathematics*. (Springer-Verlag New York) second edn.

Xiong, Wei, and Hongjun Yan, 2008, Heterogeneous expectations and bond markets, Working Paper, NBER.

Yan, Hongjun, 2008, Natural selection in financial markets: does it work?, *Management Science* 54, 1935–1950.