Heterogeneity, Selection and Wealth Dynamics

Lawrence Blume
Department of Economics, Cornell University
IHS Vienna
The Santa Fe Institute
&
David Easley
Department of Economics, Cornell University

September 6, 2009

Abstract

The market selection hypothesis states that, among expected utility maximizers, competitive markets select for agents with correct beliefs. In some economies this holds, while in others it fails. It holds in complete market economies with a common discount factor and bounded aggregate consumption. It can fail when markets are incomplete, when consumption grows too quickly, or when discount factors and beliefs are correlated. These insights have implication for the analysis of the heterogeneous agent stochastic dynamic general equilibrium models common in finance and macroeconomics.

> “The trading floor is a jungle,” he went on, “and the guy you end up working for is your jungle leader. Whether you succeed here or not depends on knowing how to survive in the jungle.”
>
> Lewis (1989, pp. 39–40.)

1 INTRODUCTION

“It’s a jungle out there.” That is the common understanding of markets, shared by economists and laymen alike. People with different tastes, skills, access to resources and beliefs compete for scarce resources. It is easy to find colorful jungle metaphors in the business and financial press, and this vision has also illuminated the thinking of economists. Marshall argued that biological rather than mechanical systems offer a better analogy for the analysis of economic dynamics (Thomas, 1991). The competition for scarce resources is the essence of Robbins (1932, p. 16) famous definition of economics as “. . . the science which studies human behavior as a relationship between ends and scarce means which have alternative uses.” Allocation through competition is the central role of markets. We understand that markets allocate resources among different productive activities, and among individuals with different endowments of skills, knowledge and materials. Most of us understand that just as the market redistributes resources among those who produce guns and those who produce butter, so too does it redistribute across those with different tastes and beliefs. Thus it is especially surprising that this insight is hardly visible in contemporary economics.

The survival of the fittest in the market jungle has long been used in finance to defend the efficient markets hypothesis. According to Cootner (1964, p. 80): ‘Given the uncertainty of the real world, the many actual and virtual traders will have many, perhaps equally many, forecasts. . . If any group of traders was consistently better than average in forecasting stock prices, they would accumulate wealth and give their forecasts greater and greater weight. In this process, they would bring the present price closer to the true value.’ The chief architect of the efficient markets hypothesis puts it slightly differently in Fama (1965, p. 38):

> . . . dependence in the noise generating process would tend to produce ‘bubbles’ in the price series. . . If there are many sophisticated

> traders in the market, however, they will be able to recognize situations where the price of a common stock is beginning to run up above its intrinsic value. If there are enough of these sophisticated traders, they may tend to prevent these 'bubbles' from ever occurring.

According to Fama (p. 40), 'A superior analyst is one whose gains over many periods of time are *consistently* greater than those of the market.'

In spite of this rhetoric about the dynamics of markets populated by heterogeneous agents, most asset pricing intuition is based on representative agent models. Modern asset pricing uses a no-arbitrage condition to derive a stochastic discount factor relating random future payoffs on assets to current asset prices. This relationship between future values and current prices does not require a representative agent. It has the virtue, as a building block for asset pricing, of being so general that its hard to reject, but without further structure it tells us little about prices. The standard approach to giving the stochastic discount factor some content is to determine it from the Euler equation for a representative consumer as in Lucas (1978). It must consume the aggregate consumption so its marginal utilities are tied down, and if it has correct expectations, its beliefs are tied down. This provides structure on asset prices, but it actually provides too much structure. It is responsible for the equity premium puzzle, the apparent excess volatility of returns, and momentum effects, all of which are empirical regularities inconsistent with equilibrium in an economy with a representative agent who has correct beliefs and psychology plausible preferences.

These puzzles have led to many attempts to fix up the analysis (Campbell, 2000). Some authors change the representative consumer's preferences to allow for habit formation or for time varying discount factors. Others take a behavioral approach in which the representative consumer is irrational, or both rational and irrational consumers are present, but constraints prevent the rational consumer from determining prices, and allow the irrational consumer to have an effect on prices. Still others consider consumers who are heterogeneous in their incomes or preferences. It is the last approach that we explore more fully in this essay.

Over the last forty years macroeconomists have given with one hand while taking away with the other. The early attempts by Hicks and his contemporaries to provide a rigorous analytical framework for Keynes' insights were largely a failure, and even while subsequent large-scale forecasting models built on Keynesian lines provided the tool kit for macroeconomic policy evaluation, macroeconomists' inability to reconcile Keynesian ideas with equilibrium notions such as market clearing was driving a wedge into the Keynesian consensus of the 50s and 60s. The now-classic book, *Microeconomic Foundations of Employment and Inflation Theory* (Phelps and et. al., 1970) was, from its opening sentence, an irresistible call for a new macroeconomic paradigm based on sound microeconomic foundations:

> The conventional neoclassical theory of the supply decisions of the household and of the firm, the theory that we all teach and rarely practice, is well known to be inconsistent with Keynesian models of employment and with post-Keynesian models of inflation.

The gifts of contemporary macroeconomics are two: First is the setting of macroeconomic analysis within a market-clearing framework. The advent of dynamic stochastic general equilibrium models requires macroeconomists to reckon with market-clearing. The use of macroeconomic relationships not derived from optimizing behavior of firms and individuals with appropriate information assumptions came to be regarded, not least by Sargent and Wallace (1975, p. 241) in describing their own work, as "deplorable". The second gift of macroeconomics is the reconsideration of expectations. While Keynes recognized the importance of expectations in *The General Theory*, only forty years later in the work of Sargent (1971), Sargent and Wallace (1975) and Lucas (1976) did they achieve their current centrality. Furthermore, the New Classical Macroeconomists understood that the hypothesis of market-clearing has implications for what kinds of beliefs can be found in the market. They imposed equilibrium restrictions on beliefs by assuming the *rational expectations hypothesis,* that beliefs are correct. As Sargent put it in an interview (Evans and Honkapohja, 2005, p. 566), "In rational expectations models, people's beliefs are among the outcomes of our theorizing. They are not inputs."

Even as they gave, however, the new classical macroeconomists took back. The simple dynamic stochastic general equilibrium models which formed the basis for their analysis are general equilibrium in name only. The jungle had only one beast, the representative consumer. Since it had no one to trade with, equating supply and demand amounted to finding prices for which the representative consumer would not want to trade. The picture is not much enriched by adding a technology, since prices now only play the role of decentralizing an optimal or constrained-optimal choice of production and consumption plans. The conditions under which aggregate excess demand can be described by optimizing behavior are well-known, are non-generic, and impose severe constraints on the kinds of time series of consumption and prices that can be realized. Second, the casting of equilibrium constraints on expectations in such a strong form ruled out consumers' diversity of beliefs. Rational expectations practitioners reintroduced animal spirits into the market, only to cage them with equilibrium restrictions which are much stronger than market-clearing alone would require.

It is possible to back away from a single economic agent and rational expectations, and contemporary macroeconomics and contemporary finance often allows limited heterogeneity. In this essay we describe a research program which builds heterogeneous agent dynamic stochastic general equilibrium models. Equilibrium restrictions on beliefs derive from the redistribution of wealth by the market. The long-run distribution of wealth will typically look quite different from the initial wealth distribution, and long-run prices will be determined by who has the resources to invest in the long run.

This research program was initially an attempt to understand the efficient markets and rational expectations hypotheses. Early work on learning by ourselves and others (Blume and Easley, 1982; Bray, 1982) convinced us that rational expectations were not globally stable under the dynamic processes that arise from the learning behavior of economic agents. This view was reinforced by the exchange between Kalai and Lehrer (1993) and Nachbar (1997), who examined a related question for Nash equilibria of dynamic games: What does it take to "learn" to play a Nash equilibrium? Kalai and Lehrer provided a description of what rational learning to play Nash entails, and Nachbar showed that the Kalai-Lehrer conditions for rational learning are often mutually inconsistent — that is, for some repeated games, rational learning is impossible. If efficient markets and rational expectations are to be well founded, then, we could not rely entirely on learning by the market participants. But in rewarding some agents while punishing others through the redistribution of wealth, the market's reweighing of beliefs and tastes is a kind of learning dynamic for the market as a whole. The *market selection hypothesis* is the claim that, through wealth redistribution, the market will come to be dominated by those with the most correct beliefs. In the next sections we build a model to test this hypothesis, demonstrate when it is true, and explore its limits.

In the next section we lay out the basic model of the complete markets, pure exchange economy that is studied in the selection literature. Section 3 provides the selection analysis for economies populated by subjective expected utility maximizers and derives survival indices for a variety of economies. Section 4 discusses the asset pricing implications of market selection. Here the question is whether in the long run assets are priced correctly and markets are efficient. In Section 5 we discuss wealth dynamics in macro models. In Section 6 we provide various extensions of the analysis to economies in which some traders are not subjective expected utility maximizers. Section 7 provides an extension of these results to economies with production. Here the focus is on whether the market selects for profit maximizing firms. We conclude in Section 8.

2 A BASIC MODEL

We begin by describing wealth dynamics in a complete market exchange economy in which each of a finite number of traders is an subjective expected utility (SEU)-maximizer. At each date $t \in \{0, 1 \ldots, \infty\}$ the agents consume the single good that is available in the economy. At each of these dates a state of the world is drawn from the set of possible states $S = \{1, \ldots, \mathbf{S}\}$, which has cardinality $\mathbf{S} < \infty$. This process generates an infinite history of states which is called a *path*, $\sigma = (\sigma_0, \ldots) \in \Sigma$. Partial histories, in addition to complete histories, will be important. The *partial history* through date t on path σ is denoted $\sigma^t = (\sigma_0, \ldots, \sigma_t)$.

The true probability distribution on infinite histories is p. Any given trader may not know p, and may instead have some other probability q on infinite histories. For any such probability, $q_t(\sigma)$ is the marginal probability of the partial history σ^t. When expectations are computed using any $q \neq p$ the expectations operator will be subscripted by q, when they are computed using p they will not be subscripted. We consider a number of random variables of the form $x_t(\sigma)$. All of these random variables are assumed to be date-t measurable; that is, their value depends only on the realization of states through date t. We will sometimes write $x_t(\sigma^t)$ to emphasize this.

In this section, we consider economies with a rich enough set of assets to yield dynamically complete markets. The simplest set of such assets are Arrow securities. We assume that at each partial history σ^t, and for each state s, there is an Arrow security which trades at partial history σ^t, and which pays off one unit of the consumption good in partial history (σ^t, s) and zero otherwise.

There are I individuals, each with consumption set $\mathbf{R}_{++}$. A *consumption plan* for an individual is a stochastic process $c : \Sigma \to \prod_{t=0}^{\infty} \mathbf{R}_{++}$. Each individual's *endowment stream* is an initial consumption plan e^i. In this section, we assume that each individual is a subjective expected utility maximizer. We also assume that each individual's payoff function is time separable and exhibits geometric discounting. In short, we assume that each individual i's preferences have an SEU representation: A payoff function u_i, a discount factor β_i and a probability distribution p^i on paths such that the utility of a consumption plan is

$$U_i(c) = E_{p^i}\left\{ \sum_{t=0}^{\infty} \beta_i^t u_i\big(c_t(\sigma)\big) \right\}, \tag{1}$$

the expected present discounted value of the stochastic process of payoffs. The three components of the SEU representation can be thought of as characterizing tastes, time preferences (more or less) and beliefs. The selection literature examines the relative importance of these three components for determining long-run survival in markets, and long-run asset prices.

The assumption that individuals are subjective expected utility maximizers implies that they have beliefs over paths, but it places no restrictions on these beliefs other than the obvious requirement that they are a probability distribution. It permits fixed beliefs, right or wrong, and it allows for learning behavior according to Bayes rule, whether or not the true distribution is in the support of prior beliefs. SEU traders act as if they update their beliefs using Bayes rule, but this is not in fact restrictive. In particular, it places no restrictions on their sequences of one-period-ahead forecasts. So even if we can observe, or infer from his behavior, an individual's entire sequence of such forecasts, there is no data that can disconfirm the hypothesis of Bayesian updating unless we hold additional assumptions about prior beliefs.

Each individual is characterized by an endowment stream e^i, beliefs p^i, a payoff function u_i, and a discount factor β_i. The analysis of competitive equilibria for these (complete markets) economies requires only a few assumptions about these objects: i) The payoff functions u_i are assumed to be C^1, strictly concave, strictly monotonic, and to satisfy an Inada condition at 0. ii) Individual endowments are assumed to be strictly positive and the aggregate endowment is assumed to be uniformly bounded. The aggregate endowment is $e = \sum_i e^i$ and $e_t(\sigma) = \sum_i e^i_t(\sigma)$ is the aggregate endowment at date t on path σ. So more precisely, the assumption is that there are numbers $\infty > F \geq f > 0$ such that for all dates t and all paths σ, $f \leq \inf_{t,\sigma} e_t(\sigma) \leq \sup_{t,\sigma} e_t(\sigma) \leq F$. iii) For all partial histories through date t there is positive probability on every state at date $t+1$ and every individual believes that this is true. That is, for all individuals i, all dates t and all paths σ, $p_t(\sigma) > 0$ and $p^i_t(\sigma^t) > 0$.

3 EQUILIBRIUM PRICES AND ALLOCATIONS

There are two approaches to describing competitive equilibrium prices and allocations. First, one could characterize competitive equilibria directly through the Euler equations that arise from individual optimization and the market clearing conditions required of equilibria. This is the approach taken in Sandroni (2000). Alternatively, one could characterize all Pareto optimal allocations from the first order conditions for the Pareto problem, and then use the first welfare theorem to note that any property satisfied by all Pareto optimal paths and the prices that support them is satisfied by any competitive equilibrium path. This is the approach taken in Blume and Easley (2006). In this article we will follow the Pareto approach as it is easier to illustrate, but in fact these two approaches are very similar and its easy to translate from one to the other.

3.1 Price Systems

A *present-value price system* is a price for consumption in each state at each date measured in units of the date 0 consumption good. In equilibrium the value of each individual's endowment must be finite, and so the value of the aggregate endowment must be finite as well. A present value price system is an $\mathbf{R}_{++}$-valued stochastic process π such that $\sum_t \sum_{\sigma^t} \pi_t(\sigma^t) \cdot e_t(\sigma^t) < \infty$. The budget set for individual i is the set $B^i(\pi, e^i)$ of all consumption plans such that $\sum_t \sum_{\sigma^t} \pi(\sigma^t)\big(c^i(\sigma^t) - e^i(\sigma^t)\big) \leq 0$.

A *current value price system* measures the price of any security available at date t in units of date t (rather than date 0) consumption. That is, the price of the state s Arrow security in units of consumption at partial history σ^t is $\tilde{q}_t^s(\sigma) \equiv \pi_{t+1}(\sigma^t, s)/\pi_t(\sigma^t)$. Arrow security prices are sometimes called *state prices* because the current value price of Arrow security s at time t is the price in partial history σ^t of one of the good in state s at time $t+1$. We will sometimes analyze normalized current-value prices: $q_t^s(\sigma) = \tilde{q}_t^s(\sigma)/\sum_\nu \tilde{q}_t^\nu(\sigma)$, which are the prices of consumption in state s relative to the price of sure consumption.

3.2 Equilibrium and Optimality

A competitive equilibrium is a tuple $(\pi, c^1, \ldots, c^I)$ wherein π is a present-value price system, and the c^i are consumption plans, one for each individual, satisfying two properties: i) Each plan c^i is utility-maximal on the individual i's budget set, and ii) Supply equals demand at each partial history; that is, for each partial history σ^t, $\sum_i c^i(\sigma^t) = e(\sigma^t)$. The economy just described is a conventional exchange economy but for the fact that there are a countable number of goods. Fortunately the complications this introduces are well understood, and the existence of a competitive equilibrium is guaranteed by, for example, the theorem of Peleg and Yaari (1970).

As we shall see below, the hypothesis of SEU rationality imposes few restrictions on equilibrium prices. Restrictions on state prices require restrictions on individuals beliefs. There are two obvious possible sources for belief restrictions. First, although individuals are not endowed with correct beliefs, they could change their beliefs over time so that they become more correct; that is, individuals may learn correct beliefs. Second, the market could select for individuals whose beliefs are for whatever reason more nearly correct. This is the market selection story.

3.3 Optimal Allocations

In any competitive equilibrium the value of the aggregate endowment is finite, and so the textbook argument proves that the first welfare theorem holds: Every competitive allocation is Pareto optimal. The first-order conditions for Pareto optimality characterize optimal consumption plans, and they can be used to study the optimal allocation of resources in the long run (Blume and Easley, 2006; Sandroni, 2000).

Optimal allocations are characterized by the property that between any two commodities, all consumers share a common marginal rate of substitution. This implies

$$\frac{u_i'\big(c^i(\sigma^t)\big)}{u_j'\big(c^j(\sigma^t)\big)} = \frac{u_i'(c_0^i)}{u_j'(c_0^j)}\left(\frac{\beta_j}{\beta_i}\right)^t \frac{p^j(\sigma^t)}{p^i(\sigma^t)} \tag{2}$$

which is the fundamental equation for characterizing long-run equilibrium consumption.

We are interested in the long run fate of individuals who may differ in their beliefs, discount factors and payoff functions. First, lets be precise about what is mean by the "long run fate" of an individual. We say that an individual i *vanishes* on path σ if the limit of i's consumption is 0 on this path, i.e. $\lim c_t^i(\sigma) = 0$. We say that individual i *survives* on path σ if i's consumption does not converge to 0 on this path, i.e. $\limsup c_t^i(\sigma) > 0$. In the long run individuals can either survive or vanish, but survival encompasses many different long run fates. A surviving individual's consumption could converge to the aggregate endowment, as it would if he was the only survivor, or it may not converge at all, in which case it must infinitely often be at least some given positive amount of the good.

The assumptions in section 2 imply that if the ratio on the left of (2) diverges along some path through the date-event tree, then along that path the consumption of individual i is vanishing. Remarkably, then, there is a sufficient condition for individual i's consumption to disappear which is independent of her payoff function. In the bounded economy of this section tastes do not matter for long-run survival; only time preferences and beliefs matter.

3.4 An IID Economy

There are many ways in which equation (2) can be manipulated to understand the long-run properties of consumption and prices. A particulary simple case arises when the state process is iid. Suppose that states are drawn independently each period, and let ρ_s denote the probability that state s is realized in any draw. Counts of the number of occurrences of states will be important in the iid case. Let $n_t^s(\sigma^t)$ be number of occurrences of state s by date t given partial history σ^t. Suppose that each individual i has iid beliefs which assign probability ρ_s^i to state s at any node in the date-event tree. Then $p^i(\sigma^t) = \prod_s (\rho_s^i)^{n_t^s(\sigma^t)}$, and equation (2) becomes

$$\frac{u_i'\big(c^i(\sigma^t)\big)}{u_j'\big(c^j(\sigma^t)\big)} = \frac{u_i'(c_0^i)}{u_j'(c_0^j)}\left(\frac{\beta_j}{\beta_i}\right)^t \prod_s \left(\frac{\rho_s^j}{\rho_s^i}\right)^{n_t^s(\sigma^t)}$$

In the IID world, this equation can be analyzed with elementary tools from probability theory. Rewriting,

$$\frac{1}{t}\log\frac{u_i'\big(c^i(\sigma^t)\big)}{u_j'\big(c^j(\sigma^t)\big)} = \frac{1}{t}c_{ij} + \log\frac{\beta_j}{\beta_i} + \frac{1}{t}\sum_s n_t^s(\sigma_t)\log\frac{\rho_s^j}{\rho_s^i} \tag{3}$$

where c_{ij} is the ratio of period 0 marginal utilities. According to the strong law of large numbers, the right hand side converges to a limit which is independent of path:

$$\frac{1}{t}\log\frac{u_i'\big(c^i(\sigma^t)\big)}{u_j'\big(c^j(\sigma^t)\big)} \xrightarrow{\text{a.s.}} \log\frac{\beta_j}{\beta_i} + \sum_s \rho_s \log\frac{\rho_s^j}{\rho_s^i}$$

Define for each trader i a *survival index*:

$$\kappa_i = \log\beta_i - I_i(\rho)$$

where $I_i(\rho) = \sum_s \rho_s \log(\rho_s/\rho_s^i)$ is the relative entropy of the true probability ρ with respect to individual i's beliefs ρ^i. Relative entropy serves as a distance. It is a non-negative convex function whose value is 0 only when $\rho = \rho^i$. The survival index in the iid economy neatly expresses a tradeoff between accuracy of beliefs and patience. Accuracy measured with relative entropy trades off one for one with patience measured with the log of the discount factor.

Equation (3) can be rewritten in terms of these indices:

$$\frac{1}{t}\log\frac{u_i'\big(c^i(\sigma^t)\big)}{u_j'\big(c^j(\sigma^t)\big)} \xrightarrow{\text{a.s.}} \kappa_j - \kappa_i.$$

This final relationship tells nearly the whole story: If $\kappa_i < \kappa_j$, then the right hand side is positive, and so along any path where convergence holds, the ratio of marginal utilities $u_i'\big(c^i(\sigma^t)\big)/u_j'\big(c^j(\sigma^t)\big)$ diverges. Thus $c^i \to 0$ almost surely. An individual whose survival index is not maximal in the population vanishes almost surely.

It is easy to underestimate the importance of the iid economy. The underlying stochastic process on states is iid, but endowments need not be iid. Endowments depend in principle on the entire past history of the economy. They are iid only if one imposes the additional restriction that $e(\sigma^t)$ depends only on σ_t. Even with iid endowments, consumption paths and prices will not typically be iid unless, of course, there is only one (representative) individual.

3.5 Long Run Heterogeneity

In the world of IID states, the long-run evolution of discounted likelihood ratios is characterized at least in part by the difference of survival indices. If an individual's survival index is not maximal, then her long-run consumption almost surely converges to 0; maximality of the index is a necessary condition for survival. Is it sufficient? A more detailed examination of equation (2) with iid state and belief processes, and equal survival indices, gives

$$\begin{aligned}\log \frac{u_i'\big(c_t^i(\sigma^t)\big)}{u_j'\big(c_t^j(\sigma^t)\big)} &= c_{ij} + t\log(\frac{\beta_j}{\beta_i}) - t\sum_s \rho_s \left(\log\frac{\rho_s^i}{\rho_s} - \log\frac{\rho_s^j}{\rho_s}\right) - \\ &\qquad \sum_s \big(n_t^s(\sigma^t) - t\rho_s\big)\left(\log\frac{\rho_s^i}{\rho_s} - \log\frac{\rho_s^j}{\rho_s}\right) \\ &= c_{ij} + \sum_s \big(n_t^s(\sigma^t) - t\rho_s\big)\left(\log\frac{\rho_s^j}{\rho_s} - \log\frac{\rho_s^i}{\rho_s}\right)\end{aligned}$$

since the omitted term is just $t(\kappa_j - \kappa_i) = 0$. Thus for two individuals with identical survival indices, the log of the marginal utility ratios is a mean-0 random walk.

From this it follows that the $\liminf$ of each individual's consumption is 0. Survivors fall into three types. Represent each survivor's beliefs by the vector of log-odds ratios with respect to, say, state 1, and now imagine the convex hull of these vectors in $\mathbf{R}^{\mathbf{S-1}}$. An individual whose log-odds vector is an extreme point of the this set is *extremal*. An individual whose log-odds vector is in the relative interior of this set is *interior* The remaining individuals are *boundary*. If the number of states exceeds 3, then for extremal individuals, $\limsup_t c_t^i / e_t = 1$, and for interior individuals, $\limsup_t c_t^i / e_t = 0$ (Blume and Easley, 2009a). In summary, a necessary condition for survival is that the survival index is maximal and that the individual does not have interior beliefs. A sufficient condition is that, in addition, beliefs are extremal.

3.6 Learning

Individuals who learn about the distribution of states can be expected eventually to have identical and correct beliefs if their learning rules are consistent. The time averages of predicted probabilities will be identical, so analyses like that of equation (3) will not help. Nonetheless, Bayesian learners can be distinguished on a finer analysis (Blume and Easley, 2006). Suppose all individuals know, correctly, that the state process is iid. Each individual believes the true model is in a set Δ_i, which is a sub-manifold without boundary of $\mathbf{R}^{\mathbf{S-1}}$; an element of Δ_i is a vector of iid state probabilities. Assume too that each individual i has prior beliefs on models which have a density q_i with respect to Lebesgue measure on Δ_i. Then $\rho^i(\sigma^t) = \int_\delta \prod_{s=1}^{\mathbf{S}} \delta_s^{n_t^s(\sigma^t)} q_i(\delta) d\delta$. This analysis departs from the previous two sections in that now beliefs are no longer iid.

Rewrite equation (2):

$$\log \frac{u_i'\big(c_t^i(\sigma^t)\big)}{u_j'\big(c_t^j(\sigma^t)\big)} = c_{ij} + t \log \frac{\beta_j}{\beta_i} + \log \frac{p(\sigma^t)}{p^i(\sigma^t)} - \log \frac{p(\sigma^t)}{p^j(\sigma^t)}.$$

The behavior of the terms $\log p(\sigma^t)/p^k(\sigma^t)$ has been studied in the statistics literature (Clarke and Barron, 1990; Phillips and Ploberger, 2003; Rissanen, 1986). They diverge as t grows, at rate $(\dim \Delta_k/2) \log t$; that is, these terms diverge slowly at a rate determined by the dimension of the space of models which the Bayesian learner considers. This has two implications. First, discount factor differences trump any differences in prior beliefs. The individual with the lower discount factor will vanish if the true distribution is in the support of the other individual's prior belief. Second, when discount factors are identical, the individual with the higher dimensional model space will vanish (again assuming that the other individual can learn the truth). Finally, if both individuals' model spaces have the same dimension, then they have a common fate; either both survive or both vanish.

These results demonstrate a tradeoff concerning model estimation. A higher-dimensional model space is more likely to contain the true distribution, but contingent on containing the truth, a lower-dimensional model space enables learning at a faster rate, and this can have an impact on long-run survival and welfare.

3.7 Endowment Growth

The conclusion that tastes do not matter for long run survival has to be modified for economies with growing endowments, see Yan (2008) and Blume and Easley (2009b). The derivation of survival properties from equation (2) relies on the fact that if $u_i'\big(c_t^i(\sigma^t)\big)/u_j'\big(c_t^j(\sigma^t)\big)$ diverges, it must be because the marginal utility of consumer i is diverging. If the endowment is growing without bound, however, the denominator could be shrinking as individual js consumption grows, so that analysis fails. Modify equation 2 to account for endowment growth as follows:

$$\frac{u_i'\big(c^i(\sigma^t)\big)/u_i'\big(e(\sigma^t)\big)}{u_j'\big(c^j(\sigma^t)\big)/u_j'\big(e(\sigma^t)\big)} = \frac{u_i'(c_0^i)}{u_j'(c_0^j)}\left(\frac{\beta_j}{\beta_i}\right)^t \frac{p^j(\sigma^t)}{p^i(\sigma^t)}\frac{u_j'\big(e(\sigma^t)\big)}{u_i'\big(e(\sigma^t)\big)} \tag{4}$$

Under a number of conditions, including boundedness of payoff functions or non-positivity of their fourth derivative, divergence of the ratio on the left implies that $u_i'\big(c^i(\sigma^t)\big)$ diverges.

A simple example in the iid economy shows how this matters. Suppose utility is in the CRRA class, $u_i(c) = \gamma_i^{-1}c^{\gamma_i}$ with $\gamma_i < 1$. Suppose that $e(\sigma^t) = r(\sigma_t)e(\sigma^{t-1})$, where the stochastic growth rate is always positive and has mean r. Then

$$\frac{1}{t}\log\frac{u_i'\big(c^i(\sigma^t)\big)}{u_j'\big(c^j(\sigma^t)\big)} = \frac{1}{t}c_{ij} + \log\frac{\beta_j}{\beta_i} + \frac{1}{t}\sum_s n_t^s(\sigma_t)\log\frac{\rho_s^j}{\rho_s^i} + \\ \frac{1}{t}\sum_s n_t^s(\sigma^t)(\gamma_j - \gamma_i)\log r(s) + \frac{1}{t}\log e_0$$

Taking limits as before gives a new survival index:

$$\tilde{\kappa}_i = \log \beta_i + I_i(\rho) + \gamma_i \log r.$$

The rate at which the marginal utility of the aggregate endowment converges to 0 depends upon the intertemporal marginal rate of substitution between present and future consumption, and so now tastes matter to this degree. For CRRA utility the marginal rate of substitution between present and future consumption depends (except at a steady state) on the coefficient of relative risk aversion. This fact is reflected in the survival index.

As simple as it is, this example is important for understanding recent results in evolutionary finance. A number of authors (Yan (2008), Kogan, Ross, Wang, and Westerfield (2006)) study long-run survival in continuous time trading markets where consumption growth is a geometric brownian motion and utility is CRRA. The present example is the discrete-time analog of that model.

Continuous-time asset pricing models are delicate objects. The analytic demands of the models admit only a limited class of endowment processes. Endowment processes are assumed to be geometric Brownian motions, the continuous time version of the iid growth rate process in discrete time. The more accommodating discrete time framework admits a richer class of endowment processes. Suppose that the mean growth rate of the aggregate endowment is time-dependent: $e(\sigma^t) = r(\sigma^t)e(\sigma_{t-1})$ where $r(\sigma^t) = r_t + \epsilon_t(\sigma_t)$. Suppose that the ϵ_t all have mean 0 and collectively satisfy conditions for a strong-law of large numbers — for instance, the variance can grow with time, but not too fast. Then the survival “index” for this economy becomes

$$\tilde{\kappa}_i = \log \beta_i + I_i(\rho) + \gamma_i \lim_{t \to \infty} \frac{1}{t} \sum_{\tau=1}^{t} \log r_\tau. \tag{5}$$

Index is in quotes because the last time average may not exist. Geometric Brownian motion corresponds to r_t constant, independent of time. If $r_t \equiv 1$ there is no drift; endowments neither shrink nor grow on average.

When the $\log(r_t)$ sequence is not constant, there are several possibilities. If their time average converges to 0, tastes do not matter, and the economy looks like those of sections 3.4 and 3.5. If their time average converges to a non-zero constant, tastes and beliefs matter as captured in the index $\tilde{\kappa}_i$. If the time average diverges to $+\infty$, then discount factors and beliefs do not matter, and survival is determined by tastes alone — only the most risk-averse survive. More interesting still is the observation that when the time average of the logged means does not converge, the right hand side of equation (5) can still be used to study the long-run behavior of the market even though no single number now characterizes survival.

4 AGGREGATION

The chief descriptive use of the general competitive model is to characterize equilibrium prices. The Debreu-Mantel-Mas Collel-Sonnenschein theorem states that any closed set of strictly positive prices can be the set of equilibrium prices in a finite good exchange economy (if there are enough consumers). Although no such theorem is known for models with an infinite number of goods, it does not seem likely that the hypothesis of equilibrium alone imposes structure on the equilibrium price set. Thus restrictions on prices will have to come from ancillary assumptions about the nature of preferences and endowments.

4.1 Price Implications of Competitive Equilibrium

In the economy of section 2, all traders have SEU preferences, and so it is natural to ask if the hypothesis of SEU rationality by itself has any implications for the relationship between endowments and present value prices across nodes. It is not hard to see that the answer to this question is “no”; by suitably choosing discount factors and beliefs, even in a single-agent economy with a given endowment process, any present value price system can be made an equilibrium by a suitable choice of discount factors and beliefs. It is hardly surprising that macroeconomists and financial economists need belief restrictions like rational expectations to make much headway; the SEU hypothesis alone is not enough.

SEU rationality without further assumptions about beliefs and discount factors has little to say about the nature of an equilibrium price system, but even with weaker assumptions, equilibrium has asset-pricing implications. If markets are complete, then assets can be priced by arbitrage: Any asset can be described as a linear combination of Arrow securities, so the price of the asset must be the corresponding linear combination of Arrow security prices. All assets can be priced by arbitrage given prices of Arrow securities. So assets are priced correctly if Arrow security prices are correct. Although the existence of correct asset prices is perhaps the key question in finance, this concept is quite slippery. If the aggregate endowment is certain (independent of state) and if individuals have identical beliefs and if these beliefs are correct, then the price of Arrow security s in state σ^t will equal $p(s|\sigma^t)$, the probability that state s will occur at date $t+1$ if the history through date t is σ^t. These prices, most economists agree, are correct. If there is social risk, that is, state-dependent endowments, prices will no longer reflect just the likelihood of tomorrow's states. Prices will reflect a risk premium as well as the likelihood of a state's occurrence, and the magnitude of these risk premia will depend upon the distribution of wealth, attitudes towards risk of the traders and discount factors. Absent this information, how could one know from prices alone if they are "correct"? (When individuals are asymmetrically informed about the state of the world and are learning from prices, there are several extant definitions of "correct" prices, but this will not concern us here.) However, the concept of correct pricing is not our immediate concern. Our central interest is in economies in which individuals have heterogeneous beliefs and endowments. If asset prices are ever correct in such economies, they are only correct in some long-run limit.

4.2 The Log Economy

To illustrate the effects of differences in beliefs and endowments we examine an economy in which all individuals have log payoff functions. This simple example is illuminating because even in an infinite horizon economy with heterogeneous beliefs and discount factors, everything can be computed. In addition, to keep things simple, we assume that the true state process is iid, and that individuals' beliefs are iid as well. This is not, as we shall see, enough to ensure that prices are iid or even Markovian.

The true probability p on Σ is generated by iid draws from a probability distribution ρ on S. Similarly, individual i's beliefs are built up from his probability ρ^i on states. The 'true probability' of state s is ρ_s and individual i's probability of state s is ρ_s^i.

The log economy is a Cobb-Douglas economy with many goods and for such economies computing the competitive equilibrium is a straightforward exercise. Let w_0^i denote the present discounted value of individual i's endowment stream, and let $w_t^i(\sigma)$ denote the amount of wealth that i transfers to partial history σ^t, measured in current units. The optimal consumption plan for individual i is to spend fraction $(1-\beta_i)\beta_i^t Pr^i(\sigma^t)$ of w_0^i on consumption at date-event σ^t. This consumption plan also has a simple description in terms of savings and portfolio choice in the Arrow securities economy. In each period, the individual consumes fraction $1-\beta_i$ of beginning wealth, w_t^i, and invests the residual, $\beta_i w_t^i$, so that the fraction α_s^i of date-t investment that is allocated to the asset which pays off in state s is ρ_s^i. Thus the individual's portfolio rule is his beliefs — he invests his beliefs.

Each unit of an Arrow security pays off 1 in its state, and the total payoff in that state must be the total invested wealth. Thus in equilibrium, supply equals demand for Arrow security s is

$$\sum_i \frac{\rho_s^i \beta_i w_t^i(\sigma)}{q_t^s(\sigma)} = \sum_j \beta_j w_t^j(\sigma),$$

and so the (normalized) price of Arrow security s at date t is

$$q_t^s(\sigma) = \sum_i \frac{\rho_s^i \beta_i w_t^i(\sigma)}{\sum_j \beta_j w_t^j(\sigma)}. \qquad (6)$$

This expression is easier to interpret if we define the date t market savings rate $\beta_t(\sigma) = \sum_i \beta_i w_t^i(\sigma)/\sum_j w_t^j(\sigma)$, individual i's date t relative savings rate $\beta_{it}(\sigma) = \beta_i/\beta_t(\sigma)$, and individual i's date t wealth share $r_t^i(\sigma) = w_t^i(\sigma)/\sum_j w_t^j(\sigma)$. Then equation (6) simplifies to

$$q_t^s(\sigma) = \sum_i \rho_s^i \beta_{it}(\sigma) r_t^i(\sigma).$$

The sum across individuals of the product of relative savings rates and wealth shares is 1. So the price of Arrow security s at date t is a weighted average of beliefs in which the weight on individual i's beliefs is the product of i's relative savings rate and wealth share. This weight corresponds to i's influence in the market as it is the share of invested wealth that is controlled by individual i.

If all individuals have identical beliefs (correct or not), then, of course, the state price will reflect this correct belief. If this common belief happens to be correct then assets are priced correctly. The more interesting case is one in which individuals have differing opinions about the probability of states. In this case the state prices, which are often interpreted as the *market's beliefs*, are an average of beliefs. But contrary to the sentiment expressed in the "Wisdom of Crowds" idea (see Surowiecki (2004)) there is no reason for this average to be correct. It would be correct if wealth shares were identical and beliefs were randomly distributed with a mean equal to the true probability. But neither of these conditions are compelling. Why should the mean belief be correct? Why should wealth shares be identical (or at leat uncorrelated with beliefs)? Initial wealth shares might be identical, but as the economy evolves they will change. These changes are correlated with portfolio choices, and thus with beliefs, so there is no hope that wealth shares and beliefs are uncorrelated. But in fact as the analysis of Section 3 shows this long run correlation, and not simple averaging of opinions, is why markets work.

In an economy in which the individuals have differing beliefs they will hold different portfolios and so wealth shares will change over time. As a result of this wealth share dynamic, state prices also evolve over time. This occurs even if endowments are IID and individual beliefs are fixed. As a result of the wealth dynamic prices are history dependent in an economy in which all of the fundamentals are iid. This occurs not because of any irrationality on the part of the individuals, but rather because wealth flows over time between heterogeneous individuals.

These wealth flows and the state price process they create have a simple interpretation (at least when there is a common discount factor). Blume and Easley (1993) shows that the market can be interpreted as a Bayesian learner with the beliefs of the individuals as the Bayesian's models and the initial wealth shares as the Bayesian's prior on these models. The equilibrium wealth share process is identical to the posterior probabilities that the Bayesian would have on the models, and the state prices are identical to the Bayesian's predicted probabilities on states. Note that this also implies that the prices are indistinguishable from those that would result in a single agent economy in which the agent was learning over this set of models. Brav and Heaton (2002) analyzes the asset pricing implications of such an economy and shows that the equilibrium asset prices that result from a representative learning individual can mimic asset pricing anomalies.

In the simple log economy there is a representative individual. He has a logarithmic utility function, his beliefs at partial history σ^t are $q_t^s(\sigma) = \sum_i \rho_s^i \beta_{it}(\sigma) r_t^i(\sigma)$ for each state s, and his discount factor is $\beta_t(\sigma) = \sum_i \beta_i w_t^i(\sigma) / \sum_j w_t^j(\sigma)$; that is, he is an equilibrium construction. No one in the economy has his beliefs or discount factor, and in fact his discount factor is stochastic. If there was a shock to wealth shares there would be a different representative consumer. It would still have logarithmic utility, but if beliefs are heterogeneous its beliefs would be different, and if discount factors are heterogeneous its discount factor would be different.

This construction used (normalized) Arrow security prices. The same analysis can be done using the stochastic discount factor approach. Individuals have a common intertemporal marginal rate of substitution arising from their Euler equations. This intertemporal marginal rate of substitution, or stochastic discount factor, is the (un-normalized) price of the relevant Arrow security (or equivalently the relevant state price). Of course, individual consumptions are heterogeneous so price data alone is not very helpful in pinning down this stochastic discount factor. But there is a representative consumer, the one constructed above, and he consumes the entire endowment. However, as he is an equilibrium construction, he can only be used in a particular equilibrium.

5 WEALTH DYNAMICS IN MACROECONOMIC MODELS

In this section we address two questions. What is the role of wealth effects in macroeconomic models, and second, how does wealth dynamics compare with other adaptive macroeconomic models.

5.1 Bewley Models

Macroeconomists have been interested in exploring the limits of the single-agent, aggregate consumer paradigm for a decade or more. Krusell and Smith (1998) approach this issue directly. They ask if "...representative-agent models of the macroeconomy might be justified by showing that models with consumer heterogeneity give rise to aggregate time series that are in fact close to those of the representative-agent models (p. 889.)" This paper, and many others as well, work in a class of heterogeneous agent macroeconomic models that Ljunqvist and Sargent (2000) labeled "Bewley models". Surprisingly, perhaps, the evolution of the distribution of wealth does not play a role in these models. This fact is puzzling, since the evolution of wealth determines the long-run behavior of the very conventional general equilibrium model of section 3.

The typical application of Bewley models has a continuum of otherwise identical individuals subject to idiosyncratic shocks, and perhaps an aggregate shock as well. The reason for studying economies with a continuum of agents is not because they are descriptively accurate, but because they are a good approximation to finite agent models which are descriptively accurate, and because they are analytically much simpler. This simplicity comes from the fact that in an appropriately constructed continuum model, a law of large numbers says that the idiosyncratic shocks average out, and the evolution of macro state variables depends only on the aggregate shocks (or is deterministic if there are no aggregate shocks).

The first question to ask of Bewley models is, in what way are they good approximations to a large finite economy with idiosyncratic risk? We claim that Bewley models are useful for studying short-run behavior, but not for the long run. We will illustrate this point in the simplest possible setting, a partial equilibrium example where it will be clear that the problem is simply a misleading application of the law of large numbers.

Our exchange economy has a finite number I of individuals. Trader i has an initial endowment of $w_0^i = 1$ unit of wealth, and 0 endowment subsequently. Each individual bets on iid coin flips of *his* coin, which can be either H (heads) or T (tails). For each coin, $\Pr\{s_t^i = H\} = p$. The return per unit correctly bet is 2. Each individual receives utility $u_i(c_t^i)$ from consumption. We will assume that $u_i(c) = (1-\gamma)^{-1}c^{1-\gamma}$; $\gamma > 0$ is the coefficient of relative risk aversion.

Individuals discount the future identically at rate β. Betting is the only way to move wealth from date 0 through the date-event tree. At date t on the current path individual i must choose a fraction δ_t^i to save and eat the rest. Of the wealth to be saved, fraction α_t^i is bet on H, the remainder on T. This can be recast as an investment portfolio of two Arrow securities, one which pays off on H and the other on T. Assets are perfectly elastically supplied at price 1.

Suppose now that there are $2N$ individuals; that is, $I = 2N$. The first N individuals have correct beliefs; they are p-individuals. The remaining N investors are q individuals, with extreme beliefs. To fix matters, suppose that $p > 1/2$ and $\alpha(q) > \alpha(p)$. Denote the ratio of the wealth of the p-individuals to the q-individuals in the $2N$-sized population at time t by R_t^N. We will suppose the exogenous R_0^N to be independent of N.

First we construct a continuum approximation. In this case it is very simple: What is the mean evolution of group wealth for the two classes of individuals?

$$\begin{aligned}\log R_1^N &\equiv \lim_{N\to\infty} \log \frac{\sum_{i=1}^N w_1^i}{\sum_{i=N+1}^{2N} w_1^i} \\ &= \lim_{N\to\infty} \log \frac{\frac{1}{N}\sum_{i=1}^N w_1^i}{\frac{1}{N}\sum_{i=N+1}^{2N} w_1^i} \\ &= \log \frac{pp^{1/\gamma} + (1-p)(1-p)^{1/\gamma}}{pq^{1/\gamma} + (1-p)(1-q)^{1/\gamma}} + \log R_0^N\end{aligned}$$

The averaging eliminated the idiosyncratic risk, and so the relationship between R_0^N and R_1^N is deterministic. Iterating this relationship, we see that

$$\log R_t^N = t \log \frac{pp^{1/\gamma} + (1-p)(1-p)^{1/\gamma}}{pq^{1/\gamma} + (1-p)(1-q)^{1/\gamma}} + \log R_0^N$$

Because the q-individuals have a higher mean than do the p individuals, the coefficient of t is negative, and so $\log R_t^N \to -\infty$ and the wealth share of the p-individuals, those with accurate beliefs, converges to 0. That is,

$$\lim_{t\to\infty} \lim_{N\to\infty} R_t^N = 0. \tag{7}$$

This fact, however, is not what we want to know. Instead, we care about

$$\lim_{N\to\infty} \lim_{t\to\infty} R_t^N = \infty. \tag{8}$$

Pair off the individuals, the first p-investor against the first q-investor, and so on. Calculations like those in section 3 show that $\lim_{t\to\infty} \log w_t^i / w_t^{N+i} \to +\infty$ almost surely, at a rate determined by γ and the relative entropy of p with respect to q, and from this fact equation (8) follows.

Analyzing asymptotic behavior using a continuum approximation is equivalent to taking the limit as in equation (7): Compute the transitions for an infinite agent economy, and then iterate them an infinite number of times. The limit we want to measure is that in equation (8). In this example it is easy to compute, and the answers are different. In a working macroeconomic model, the limit (8) is much harder to determine. This observation does not mean that continuum approximation models are useless. One can show that over a given time horizon and for a given level of tolerance, a large enough economy approximates the continuum path in probability. How large and for how long are questions as important as they are difficult to answer. The one class of economies for which the order of limits clearly does not matter has complete markets and identical homothetic individuals.

The Bewley model is important for studying the implications of income distribution dynamics for macroeconomic phenomena. If all individuals are identical, the equal treatment property of competitive equilibria implies that the distribution of income will be degenerate, so the existence of a non-trivial income distribution requires heterogeneity of individuals' characteristics. It is often alleged that heterogeneity has no interesting implications in complete markets (Heathcote, Storesletten, and Violante, 2009). But Lucas' 1992 assertion that, "If the children of Noah had been able and willing to pool risks, Arrow-Debreu style, among themselves and their descendants, then the vast inequality we see today, within and across societies, would not exist," contradicts the biblical text. The text attributes the long-run differences in wealth between the descendants of Ham, Shem and Japheth to their personal attributes and not to market structure. The analysis of section 3 shows that different characteristics can have effects which can persist for a long time, even asymptotically. This is not to deny the importance of incomplete markets. The fluctuations we observe in complete markets may appear in incomplete markets too, but will they are masked by the averaging procedure and their effect will be missed. Do these effects amplify, or do they perhaps even offset, the incomplete markets effects documented in the literature? This is an interesting question for future research.

5.2 Learning

In his presidential address, Sargent (2008) contrasts two sources of ideas about the workings of monetary policy: Rational expectations equilibrium, and adaptive equilibrium processes. The adaptive process Sargent refers to here is learning by individuals and policymakers. Rational expectations and learning are closely connected. If rational expectations are a justifiable equilibrium concept, there must be a story of how equilibrium is achieved. For rational expectations, justifications typically appeal to learning, and yet there is considerable doubt that the outcome of learning must be rational expectations (Blume and Easley, 1982; Bray, 1982).

The evolution of the distribution of wealth through repeated trade in a population heterogeneous with respect to tastes, beliefs and discount factors, is another adaptive process. Individuals may or may not be adapting through learning, but the system adapts by redistributing resources among individuals with different characteristics. This raises an obvious question: How does market adaptation through redistribution differ from individual adaption through learning?

In some sense they need not differ. As we have already observed, it is possible to construct an economy wherein individuals have fixed but different beliefs, and identical discount factors and utility functions, such that the price and aggregate consumption data look like that of an economy with a representative Bayesian-learning consumer. Each individual has a fixed model of the world, and the economy selects among models in a Bayesian fashion. There is reason to believe, however, that in more realistic economies wealth dynamics and learning dynamics would not select the most correct model, and perhaps behave quite differently.

The Bayesian learning models we study have one unfortunate characteristic. They are simple enough that Bayesian learning (and many other learning models) are consistent. That is, given enough time, individuals will know the correct model to be true. This is an artifact of the small size of both the model and the sample spaces, and the regularity of likelihood functions. Freedman (1965) observes that in complicated problems the typical behavior of posterior beliefs is to wander around; in fact, at different points of time a Bayesian learner will be nearly sure that the true model lies in any given open set in model space. Whether realistic learning rules will exhibit the behavior Freedman describes is an open empirical question. The toy models we successfully analyze certainly do not.

A related issue is the existence of what we call a “complexity gap”. For a model of the economy to be correct, it has to describe the stochastic evolution of the economy both in some long run state when beliefs have settled down (assuming they do) and in the intermediate runs when beliefs have not converged. This means that individuals’ models have to account not just for their own learning, but also for the learning behavior of others. The assumption that the supports of individuals’ prior beliefs contains the true model seems excessively strong.

Finally, the ability to learn depends upon the amount and quality of available information. The models we have described in section 3 make all information available to all individuals. Every individual knows the full history of state evolution from the beginning through the present. Better models might relax this assumption. The market selection mechanism, however, works on all the information generated by the economy and state evolution process. In this sense the market could indeed be smarter than any individual trader.

6 BEHAVIORS AND SURVIVAL

Market selection forces are at work in any heterogeneous agent economy, not just those containing subjective expected utility maximizing traders. In this section we briefly discuss three interesting classes of economies to which these ideas have been applied. First, we extend the analysis of market selection to decision theories that generalize SEU. Second, we investigate the survival of noise traders in financial markets. Finally, we discuss a more general framework for market selection that replaces decision theories of any kind with demand, thereby offering a behavioral view of market selection.

It is useful to first note that regardless of how an individual’s behavior is motivated if this behavior is consistent with SEU maximization for some beliefs, payoff function and discount factor then the analysis of Section 3 can be applied. In a bounded economy with complete markets all that matters is survival indices; an individual whose survival index is not maximal cannot survive. It does not matter how these survival indices are generated. As SEU maximization provides very little structure on behavior many non-SEU motivated rules may be consistent with SEU maximization, and the fate of individuals following these rules is the same the fate of their observationally equivalent SEU maximizer.

6.1 Non-SEU Behavior

Behavior in experiments and in asset markets that is inconsistent with expected utility maximization has lead to the creation of various alternative decision theories. The most famous of the experiments is the Ellsberg (1961) (thought) experiment in which many individuals prefer bets with known odds to bets with unknown odds, even when expected utility maximization using natural distribution over the unknown odds makes the unknown odds gamble as least as attractive as the known odds gamble. This behavior cannot be rationalized by standard expected utility theory. Empirical evidence from asset markets about the equity premium, home bias in portfolios and non-participation is also difficult to reconcile with standard expected utility theory. This collection of results has been interpreted as evidence that some individuals are averse to ambiguity.

Schmeidler (1989) and Gilboa and Schmeidler (1989) provide similar axiomatic foundations for behavior which encompass this ambiguity aversion. The Gilboa-Schmeidler theory represents an individual's preferences with a set of beliefs and a payoff function such that an individual evaluates any gamble according to its expected utility using the belief in his set of beliefs that gives the gamble its lowest expected utility. This theory has been used to understand Ellsberg's results as well as empirical asset pricing and portfolio choice puzzles. Typically explanations of the asset market puzzles consider economies populated entirely by individuals who behave according to the alternative theory. It is also instructive to ask what happens if some individuals are ambiguity averse and others are not.

Condie (2008) asks this question in a complete markets economy which has at least one individual who is an expected utility maximizer with correct beliefs. He shows that if there is no aggregate risk, then ambiguity averse traders with correct beliefs in the set of beliefs they consider can survive, but they have no affect on asset prices. To see how this occurs consider the Arrow securities prices that would be set if only the expected utility individual existed. These prices are determined by this individual's marginal rates of substitution evaluated at the risk-free aggregate endowment, so they are just ratios of the probabilities of the states. At these (correct) asset prices ambiguity averse individuals also do not want to hold risky portfolios. In effect their indifference curves have a kink on the diagonal of the Edgeworth Box and this is where the equilibrium must lie.

Condie (2008) also shows that if there is aggregate risk in the economy then ambiguity averse individuals whose set of beliefs containing the truth in its interior cannot survive. If an ambiguity averse individual holds a risk-free portfolio, then the expected utility person is earning the return to holding the aggregate risk and will drive the ambiguity averse person out of the market. Alternatively, if the ambiguity averse individual holds risk, then he is acting as if he is an expected utility individual with incorrect beliefs (those that minimize expected utility for the portfolio) and he cannot survive. In either case ambiguity averse individuals have no effect on prices in the long run.

It would be interesting to ask whether similar results hold for other types of non-SEU motivated behavior. For example, one could consider Bewley's 2002 incomplete preference approach to ambiguity aversion, prospect theory, rank dependent expected utility and so forth. Moreover, the market test should be applied to behavioral decision rules. Does hyperbolic discounting survive?

6.2 Noise Traders

Black (1986) defines noise traders as individuals who trade using uninformative signals. Two papers, DeLong, Shleifer, Summers, and Waldmann (1990, 1991), have seeded a literature with their investigation of the survival of noise traders in financial markets. The results surveyed in section 3, however, state that Black's noise traders do not survive. They vanish in the long run when markets are complete. Suppose these traders live in an economy in which a complete set of Arrow securities are traded and some individuals have SEU preferences with correct beliefs. Others are noise traders — individuals with SEU preferences and arbitrary beliefs that could be informed by noise, or anything else. Then one of two things happens: Either their beliefs converge fast enough to the correct beliefs, or their share of aggregate consumption goes to 0. In either case, the noise traders disappear from the market, and have no impact on long-run prices.

The tension between these two results, the survival and death of noise traders, lies in their derivation. DeLong, Shleifer, Summers, and Waldmann (1990) provides an overlapping generations model with noise traders and expected utility traders and shows that the noise traders earn higher expected returns than do expected utility traders by taking on more risk. Since the demographic structure of the model holds the group sizes constant, survival is not an issue. Nonetheless this argument is often cited as a justification for the survival of noise traders. A more sophisticated version of this argument is contained in DeLong, Shleifer, Summers, and Waldmann (1991). This general equilibrium model has a continuum of each of two types of infinitely-lived agents: noise traders and informed traders. Here the noise traders' higher mean return is alleged to imply their long-run survival. The higher variance to noise trader returns is made to disappear by an appeal to the law of large numbers. In section 5.1 we discussed the problem with this line of argument. In general, higher mean returns are not important for long-run survival. Breiman (1961) first noted that in favorable gambling situations it is *expected rates of returns* and not *expected returns* that determines both long run wealth and the likelihood of survival. The argument that higher expected returns guarantees long-run survival confuses $\mathrm{E}\{\log w_t\}$ with $\log \mathrm{E}\{w_t\}$.

The analysis of section 3 does not necessarily spell the death of noise traders. Noise traders could easily survive in incomplete markets. They disappear in complete markets because other individuals, including those with more accurate beliefs, can bet against them. When markets are incomplete, they may provide no opportunity for individuals to trade based on their diverse beliefs. Noise traders may thus be constrained from making losing bets. Blume and Easley (2006) provides an example of an incomplete markets economy in which, due to market incompleteness, noise traders survive, individuals with correct beliefs vanish, and the noise traders effectively set the long run (incorrect) asset prices.

The analysis described in section 3 only applies to individuals whose behavior can be represented by the maximization of additively separable intertemporal SEU preferences. Noise trading in fact has many sources, including trading behavior that cannot be rationalized by any preference-based choice model. In such cases a direct analysis of their survival in the competitive equilibrium induced by their rules of behavior is necessary. The next section addresses this topic.

6.3 Rules

Individual's characteristics matter for survival only through their influence on behavior, so our survival results can be reinterpreted as showing which rules survive within the class of behavioral rules that can be generated by subjective expected utility maximization. The survival question can also be asked directly about rules of behavior which allows consideration of rules that may not be SEU maximizing.

Suppose in the economy that we have been studying (but one with a possibly incomplete set of Arrow securities) that each individual is described by an endowment process (as before) and by portfolio and savings rules, rather by beliefs and a payoff function. The individual's savings rule determines the fraction of his wealth that the individual saves at each date given any partial history of states. The individual's portfolio rule determines the fraction of his savings that the individual invests in each Arrow security. These rules could be consistent with SEU maximizing behavior, but they need not be within this class.

Blume and Easley (1992) shows that an individual whose savings rate is maximal, and whose portfolio rule is always the conditional probability of states, has a maximal expected growth rate of wealth share. This analysis also shows that selection over rules is determined by the expected growth rate of wealth share, so if there is a single individual with maximal expected growth rate of wealth share he is selected for, and all others vanish. This most fit rule is consistent with the individual being a SEU maximizer with logarithmic utility, correct beliefs and a discount factor that is as large as any trader's savings rate. So this analysis extends the results of Kelly (1956) and Breiman (1961) about betting with exogenous odds to a market setting. Blume and Easley (1992) also notes that there are portfolio rules consistent with SEU behavior which do not maximize fitness. So SEU individuals can be driven out of the market by individuals whose rules cannot be generated from SEU maximization.

Rather than asking about survival in the entire class of behavioral rules one could ask about survival within restricted classes of rules. This is one way to interpret the selection question applied to SEU maximizing rules. There are other classes of rules that are interesting. Amir, Evstigneev, Hens, and Schenk-Hoppé (2005) ask if there are particular simple portfolio rules that are selected for when all behavioral rules are simple. In their analysis a simple rule does not depend on current asset prices. Their focus is only on portfolio rules, so all individuals are assumed to save at the same rate. Amir, Evstigneev, Hens, and Schenk-Hoppé (2005) show that an individual whose portfolio rule allocates his wealth across assets according to their conditional expected relative payoffs drives out all other traders as long as none of the other traders end up holding the market. The log optimal portfolio rule from Blume and Easley (1992) selects the same portfolio as does the conditional expected relative payoff rule when only these two rules exist in the market. So both of these rules end up holding the market in the limit and both survive.

7 SELECTION AND WELFARE

The first application of market selection arguments in post-war economic thought was to the profit maximization debate. Do marginal analysis and the profit maximization hypothesis provide good models of firm behavior? Some argued that firms do not know their profit functions, are unable to calculate, or use rules of thumb which are inconsistent with profit-maximization. (See Machlup (1946) for a summary of these arguments.) Although many economists (including Machlup) justified the profit maximization hypothesis by adaptationist arguments, the classic such statement is that of Friedman (1953, p. 22), who wrote:

> Whenever this determinant (of business behavior) happens to lead to behavior consistent with rational and informed maximization of returns, the business will prosper and acquire resources with which to expand; whenever it does not the business will tend to lose resources and can be kept in existence only by the addition of resources from the outside. The process of natural selection thus helps to validate the hypothesis (of profit maximization) or, rather, given natural selection, acceptance of the hypothesis can be based largely on the judgment that it summarizes appropriately the conditions for survival.

This argument is used to justify two claims: First, that firms which appear to be profit maximizers drive others from the market, so that the profit maximization hypothesis describes the long-run market composition. Second, it is claimed that the long run behavior of markets is in fact characterized by competitive market equilibrium with profit-maximizing firms. Koopmans (1957, p. 140) criticized the second claim:

> Here a postulate about individual behavior is made more plausible by reference to the adverse effect of, and hence penalty for, departures from the postulated behavior . . . But if this is the basis for our belief in profit maximization, then we should postulate that basis itself and not the profit maximization which it implies in certain circumstances.

Blume and Easley (2002) investigates both claims in an intertemporal market model with production. Some individuals are workers who buy goods with income earned from selling their labor and other primary factor inputs from their endowments. The remaining individuals are entrepreneurs who own technologies for transforming inputs into consumer goods. Production is intertemporal; inputs today produce goods for sale tomorrow. Technologies are conventional, that is, convex. However, firms may not profit maximize. In all cases, however, their decision rules — factor demands and output supply — satisfy necessary conditions on firms for the existence of market-clearing prices. There are no intermediate goods. All inputs are primary and all outputs are consumer goods only.

In a conventional complete-market intertemporal competitive equilibrium, selection is a one-shot event. Suppose, for simplicity, that for every technology there are some profit maximizing firms and perhaps other firms with different decision rules. In any competitive equilibrium, only the profit maximizing firms will produce. They will borrow money to buy inputs, and repay the loans with revenues from the sale of outputs. Non-profit maximizing firms will not produce. This, of course, is no surprise. Inefficient firms will not operate in a competitive equilibrium.

Arguably, this is what Friedman had in mind. Inefficient firms cannot attract capital to fund their operations, and so they disappear. However, his argument seems not to rest upon the existence of complete markets and outside investors. In the competitive equilibrium it is investors and not "the market" which is doing the sorting. The following incomplete market structure also seems to be consistent with many selection arguments, including Friedman's. Firms are initially endowed with cash as well as with their technologies. Firms cannot lend or borrow; they must purchase inputs from cash on hand. Firms pay a dividend to their owner, so all operations are funded out of retained earnings. There are no capital markets. Workers eat their current earnings, and the only savings opportunities afforded entrepreneurs are through the firm.

Entrepreneurial decisions are made with an eye to consumption. Nonetheless, firm survival, as measured by share on input expenditures, is determined only by the firm's decision rule and the discount factor, and not by the entrepreneur's utility function. A few regularity assumptions on decision rules guarantees that only profit maximizing firms survive in the long run. So the Chicago school was right to argue that market forces favor profit maximization under some rather broad general conditions. This is not surprising: Profit maximizing firms give higher rates of return on investment. These extra returns accumulate so that ultimately the profit maximizers have an ever larger share of retained earnings among those firms with access to the same technology.

This is not the end of the story. We care about profit-maximizing firms because profit maximization is an essential part of the argument that competitive outcomes are optimal. Koopmans hints, however, that the selection dynamics may lead to non-optimal outcomes even as it favors profit maximization. This would not be surprising in the retained earnings economy since markets are incomplete. More surprising, however, is the possibility that long run production can be producer-inefficient; that is, the aggregate production plans can ultimately lie in the interior of the aggregate production set.

Blume and Easley (2002) constructed an economy to demonstrate this point. In this economy there are two consumer goods, one input, and four firms with Leontief production functions. The "per unit" production possibility set can be described by the convex hull of the firms' unit output vectors, the vectors of what can be produced per unit of input, and the origin. It is drawn in figure 1. In it's interior are the unit output vectors of two more firms, the inefficient firms. Efficient production utilizes only the efficient firms. Any output bundle achieved through the operation in part of one of the inefficient firms can be improved upon by reallocating its input between the efficient firms. Suppose now that the two outputs are very complementary in the consumers' preferences. Then equilibrium paths will have extreme price oscillations. The firms with unit output vectors in the lower right of figure 1 make large profits when the price of good 1 is high,

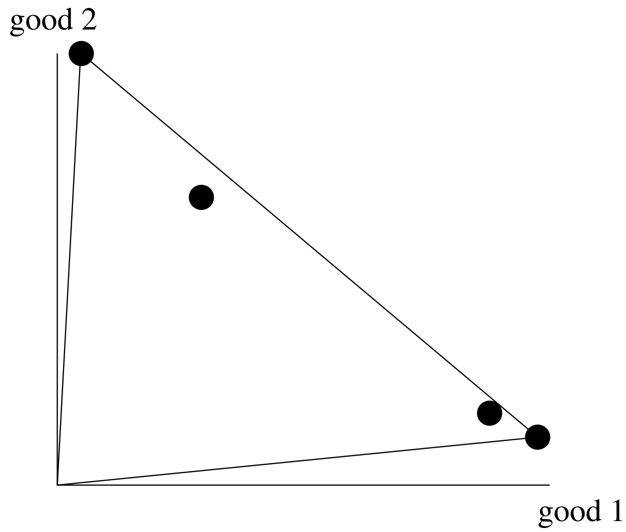

Figure 1: Production Possibility Frontier

and little profit when the price is low, and vice versa for the firms in the upper left. Although each inefficient firm does less well than the corresponding efficient firm when the price is in its favor, it does better when the price is not in its favor. The inefficient firms do better enough that their profits through the cycle are higher. In the long run the two efficient firms disappear, and it is as if the production possibility set has contracted inward to the convex hull of the origin with the two inefficient unit output vectors.

The intuition that evolution leads to optimal outcomes, that selection is an invisible hand, is flatly contradicted by this example. This should not surprise us because we do not expect equilibrium to be optimal in incomplete markets. More surprising, though, is the long run failure of producer efficiency. This, of course, is a consequence of missing markets. There is no market in which the inefficient firm's capital can be reallocated to the efficient firms, and the dynamics of retained earnings cannot accomplish this reallocation on their own. If the two efficient firms were to merge, capital could be reallocated internally between the two efficient technologies, and then the inefficient firms would vanish in the long run.

8 CONCLUSION

The study of heterogeneous agent general equilibrium models has led to important conclusions about which characteristics are selected for in economies with complete markets. For economies composed of subjective expected utility maximizers there is a survival index which can be computed for each agent from that agent's characteristics and, if the economy is growing, the growth rate. Agents with survival indices that are not maximal do not survive. In bounded economies, the only characteristics that matter are discount factors and beliefs; risk aversion is irrelevant. In growing economies, the growth rate relative to the curvature of an agent's utility function also matters as this curvature affects the agent's intertemporal marginal rate of substitution.

This analysis of survival indices shows that, controlling for discount factors, an economy with complete markets and bounded aggregate consumption selects for traders with correct expectations. Among traders who are learning it selects for Bayesians with the truth in the support of their priors; and among Bayesians it selects for those whose priors have lower dimensional supports. This collection of results provides some support for the market selection hypothesis—that competitive markets select for agents who maximize expected utility using correct beliefs. As a result it provides some support for the hypothesis that, at least in the long run, assets are priced correctly and markets are efficient, even if initially some traders do not behave as predicted by models consisting of only agents who maximize objective expected utility. This support is qualified however as the analysis makes strong assumptions about the economy. Most important is the assumption of complete markets. Blume and Easley (2006) provides examples that demonstrate what can go wrong in economies with incomplete markets. In these economies it is possible for agents with incorrect expectations to drive those with correct expectations out of the market and for asset prices to be wrong even in the long run. The assumption that the economy is bounded is also important. Survival indices can be computed for unbounded economies, but the strong selection force for those with correct expectations is lost as now utility functions also matter. Finally, discount factors matter. If they are correlated with beliefs in the right way then its plausible that agents with correct beliefs may not come to dominate the market.

A more general implication of this literature is that analyzing infinite horizon, stochastic general equilibrium economies with heterogeneous agents is possible. Analyses of this sort are beginning to appear in both finance and macroeconomics (Cogley and Sargent, 2008) in response to the inability of representative agent models to fit asset pricing and macro data. We believe that this approach is more promising than approaches which instead retain representative agents but make them strongly irrational. In addition to our methodological objection to these approaches we don't expect them to stand up to market selection. In a heterogeneous agent, complete markets economy in which some agents are objective expected utility maximizers and others follow behavioral rules the results that are known so far strongly suggest that either the behavioral traders disappear or that they persist but don't affect anything. We believe that exploring selection in complex economies of this sort is a promising direction for future research.

References

AMIR, R., I. EVSTIGNEEV, T. HENS, AND K. R. SCHENK-HOPPÉ (2005): “Market selection and the survival of investment strategies,” *Journal of Mathematical Economics*, 41, 105–22.

BEWLEY, D. L. (2002): “Knightian decision theory. Part I,” *Decisions in Economics and Finance*, 25, 79–110.

BLACK, F. (1986): “Noise,” *Journal of Finance*, 41, 529–543.

BLUME, L., AND D. EASLEY (1992): “Evolution and Market Behavior,” *Journal of Economic Theory*, 58(1), 9–40.

——— (1993): “Economic Natural Selection,” *Economics Letters*, 42(2-3), 281–289.

BLUME, L. E., AND D. EASLEY (1982): “Learning to be Rational,” *Journal of Economic Theory*, 26(2), 340–351.

——— (2002): “Optimality and Natural Selection in Markets,” *Journal of Economic Theory*, 107(1), 95–135.

BLUME, L. E., AND D. EASLEY (2006): “If you’re so smart, why aren’t you rich? Belief selection in complete and incomplete markets,” *Econometrica*, 74(4), 929–966.

BLUME, L. E., AND D. EASLEY (2009a): “The Market Organism: Long-Run Survival in Markets with Heterogeneous Traders,” *Journal of Economic Dynamics and Control*, 5(5), 1023–1035.

BLUME, L. E., AND D. EASLEY (2009b): “Market Selection and Asset Pricing,” in *Handbook of Financial Markets: Dynamics and Evolution*, ed. by T. Hens, and K. Schenk-Hoppe. Elsevier.

BRAV, A., AND J. B. HEATON (2002): “Competing Theories of Financial Anomalies,” *The Review of Financial Studies*, 15(2), 575–606.

BRAY, M. (1982): “Learning, Estimation, and the Stability of Rational Expectations,” *Journal of Economic Theory*, 26(2), 318–339.

BREIMAN, L. (1961): "Optimal gambling systems for favorable games," in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, ed. by J. Neyman, vol. 1, pp. 65–78, Berkeley. University of California Press.

CAMPBELL, J. (2000): "Asset Pricing at the Millennium," *Journal of Finance*, LV(4), 1515–1567.

CLARKE, B. S., AND A. R. BARRON (1990): "Information Theoretic Asymptotics of Bayes Methods," *IEEE Transactions on Information Theory*, 36(3), 453–71.

COGLEY, T., AND T. J. SARGENT (2008): "Diverse Beliefs, Survival, and the Market Price of Risk," Unpublished, New York University.

CONDIE, S. (2008): "Living with ambiguity: prices and survival when investors have heterogeneous preferences for ambiguity," *Economic Theory*, 36(1), 81–108.

COOTNER, P. H. (1964): *The Random Character of Stock Market Prices*. MIT Press, Cambridge, MA.

DELONG, J. B., A. SHLEIFER, L. SUMMERS, AND R. WALDMANN (1990): "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98(4), 703–38.

——— (1991): "The Survival of Noise Traders in Financial Markets," *Journal of Business*, 64(1), 1–19.

ELLSBERG, D. (1961): "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics*, 75, 643–69.

EVANS, G., AND S. HONKAPOHJA (2005): "An Interview with Thomas Sargent," *Macroeconomic Dynamics*, 9, 561–583.

FAMA, E. (1965): "The Behavior of Stock Market Prices," *Journal of Business*, 38(1), 34–105.

FREEDMAN, D. A. (1965): "Onthe Asymptotic Behavior of Bayes Estimates in the Discrete Case II," *Annals of Mathematical Statistics*, 36(6), 454–56.

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

GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18(2), 141–153.

HEATHCOTE, H., K. STORESLETTEN, AND G. VIOLANTE (2009): “Quantitative Macroeconomics with Heterogeneous Households,” *Annual Reviews of Economics*, 1, 319–54.

KALAI, E., AND E. LEHRER (1993): “Rational Learning Leads to Nash equilibrium,” *Econmetrica*, 61(5), 1019–45.

KELLY, J. L. (1956): “A new interpretation of information rate,” *Bell System Technical Journal*, 35, 917–926.

KOGAN, L., S. ROSS, J. WANG, AND M. WESTERFIELD (2006): “The Price Impact and Survival of Irrational Traders,” *Journal of Finance*, 61, 195–229.

KOOPMANS, T. C. (1957): *Three Essays on the State of Economic Science*. McGraw–Hill, New York.

KRUSELL, P., AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *The Journal of Political Economy*, 106(5), 867–96.

LEWIS, M. (1989): *Liar’s Poker*. W. W. Norton, New York.

LJUNQVIST, L., AND T. J. SARGENT (2000): *Recursive Macroeconomic Theory*. MIT Press, Cambridge, MA, 2nd edn.

LUCAS, R. (1978): “Asset Prices in an Exchange Economoy,” *Econometrica*, 46, 1429–1446.

LUCAS, R. (1992): “On Efficiency and Distribution,” *Economic Journal*, 102(411), 233–47.

LUCAS, R. E. (1976): “Econometric Policy Evaluation: A Critique,” in *The Phillips Curve and Labour Markets*, ed. by K. Brunner, and A. H. Meltzer, Amsterdam. North Holland.

MACHLUP, F. (1946): “Marginal Analysis and Empirical Research,” *American Economic Review*, 36(4), 519–54.

NACHBAR, J. (1997): “Prediction, optimization, and learning in repeated games,” *Econometrica*, 65(2), 275–309.

PELEG, B., AND M. E. YAARI (1970): "Markets with Countably Many Commodities," *International Economic Review*, 11(3), 369–377.

PHELPS, E. S., AND ET. AL. (1970): *Microeconomic Foundations of Employment and Inflation Theory*. W. W. Norton, New York.

PHILLIPS, P. C. B., AND W. PLOBERGER (2003): "Empirical Limits for Time Series Econometric Models," *Econometrica*, 71(2), 627–73.

RISSANEN, J. (1986): "Stochastic Complexity and Modelling," *Annals of Statistics*, 14(3), 1080–1100.

ROBBINS, L. (1932): *An Essay on the Nature and Significance of Economic Science*. MacMillan and Co., London, 2nd edn.

SANDRONI, A. (2000): "Do Markets Favor Agents Able to Make Accurate Predictions," *Econometrica*, 68(8), 1303–42.

SARGENT, T. J. (1971): "A note on the accelerationist controversy," *Journal of Money, Credit and Banking*, 3(721–725).

——— (2008): "Evolution and Intelligent Design," *American Economic Review*, 98(1), 5–37.

SARGENT, T. J., AND N. WALLACE (1975): "Rational expectations, the optimal monetary instrument, and the optimal money supply rule," *Journal of Political Economy*, 83(241–254).

SCHMEIDLER, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.

SUROWIECKI, J. (2004): *The Wisdom of Crowds: Why the Many are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Random House, New York.

THOMAS, B. (1991): "Alfred Marshall on economic biology," *Review of Political Economy*, 3(1), 1–14.

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

List of Figures