

The Market Organism: Long Run Survival in Markets with Heterogeneous Traders

Lawrence Blume and David Easley
Cornell University



September 2006

First version, November 2000. Research support from the National Science Foundation for research support under grant SES 9808690 and the Cowles Foundation for Research in Economics at Yale University is gratefully acknowledged. This draft has been prepared for the European Central Bank Conference, 'Learning, Asset Markets and Monetary Policy', and is very incomplete.

‘Amazingly, the stock market knows better than the analysts do.’

Henry Herrmann, Chief Investment Officer, Waddell & Reed

1 Introduction

Economists, businessmen and lay people all talk about the knowledge of the market: ‘The market learns...’ and ‘the market knows ...’ are all accepted explanations for observed pricing phenomena. The ultimate expression of this idea is the wide and recent interest in prediction markets. Despite the embarrassment of the United States Defense Department’s FutureMAP program, today companies like Microsoft, Eli Lilly and Hewlett-Packard use information markets as a way of eliciting information from workers and managers in order to guide decisionmaking.¹ Even very early studies of market efficiency compared market performance to the performance of experts, and most found that the market did at least as well (Figlewski 1979, Snyder 1978). There are three ways in which a market can predict: The market can ‘balance’ the different beliefs of traders. On this account the market could be more accurate than any single trader’s information. The market can ‘select’ beliefs; that is, markets favor traders with more accurate information, and as these traders grow in wealth, market prices come to reflect their views. This is an old Chicago School argument often attributed (incorrectly) to Milton Friedman, and its implications for asset markets were drawn out by Fama (1965) and Cootner (1964). The market can exchange information among traders; that is traders can learn what others know from market prices. This is the idea behind rational expectations and the literature on learning from prices.

The literature on informational exchange in markets is huge. Market balancing and market selection, on the other hand, are much less studied. Here we will build some simple dynamic equilibrium models to investigate the the long run behavior of asset prices in markets with heterogeneous beliefs. We build upon the market selection results of Blume and Easley (1992), Sandroni (2000) and Blume and Easley (2006). Along the way we will extend the analysis of these papers. In particular, we will show that the necessary conditions for traders’ long-run survival developed in these three papers are not sufficient. But our main purpose is to illustrate the implications of our (extended) selection analysis for the behavior of asset prices.

We study complete markets. The assets we price are Arrow securities. More complex assets can be priced by arbitrage from these assets. We do not allow traders to learn. Blume and Easley (2006) conduct a detailed examination of the market selection hypothesis when traders learn, and the implications of that analysis for asset prices could be traced out. The interaction of selection and asymmetric information is addressed in Mailath and Sandroni (2003). Here we prefer to study the effects of balancing and selection in isolation, without the interesting but confounding effects of information sharing.

2 The Model

Our model is an infinite horizon exchange economy which allocates a single commodity. Our method is to examine Pareto optimal consumption paths and the prices which support them. The first welfare theorem applies to the economies we study, so every competitive path is Pareto optimal. Thus any property of all optimal paths is a property of any competitive path. In this section we establish basic notation, list the key assumptions and characterize Pareto optimal allocations.

We assume that time is discrete and begins at date 0. The possible states at each date form a finite set $S = \{1, \dots, \mathbf{s}\}$, with cardinality $\mathbf{s} = |S|$. The set of all infinite sequences of states is Σ with representative sequence $\sigma = (\sigma_0, \dots)$, also called a *path*. σ_t denotes the value of σ at date t , and $\sigma^t = (\sigma_0, \dots, \sigma_t)$ denotes the partial history through date t of the path σ . Let H_t denote the set of all partial histories through date t , let $H_0 = \{\sigma^0\}$, the set containing the null history, and let $H = \cup_{t=0,1,\dots} H_t$ denote the set of all partial histories. Since the processes and beliefs are iid, counts will be important. Let $n_t^s(\sigma^t) = |\{\tau \leq t : \sigma_\tau = s\}|$.

The set Σ together with its product sigma-field is the measurable space on which everything will be built. Let p denote the “true” probability measure on Σ . It is the distribution on sequences consistent with iid draws from probability distribution ρ on S . The ‘true probability’ of state s is $\rho(s)$.

Expectation operators without subscripts intend the expectation to be taken with respect to the measure p . For any probability measure p' on Σ , $p'_t(\sigma)$ is the

(marginal) probability of the partial history $\sigma^t = (\sigma_1, \dots, \sigma_t)$. That is, $p'_t(\sigma) = p'(\{\sigma_1\} \times \dots \times \{\sigma_t\} \times S \times S \times \dots)$.

In the next few paragraphs we introduce a number of random variables of the form $x_t(\sigma)$. All such random variables are assumed to be date- t measurable; that is, their value depends only on the realization of states through date t . Formally, \mathcal{F}_t is the σ -field of events measurable through date t , and each $x_t(\sigma)$ is assumed to be \mathcal{F}_t -measurable.

2.1 Traders

An economy contains I traders, each with consumption set \mathbf{R}_{++} . A *consumption plan* $c : \Sigma \rightarrow \prod_{t=0}^{\infty} \mathbf{R}_{++}$ is a sequence of \mathbf{R}_{++} -valued functions $\{c_t(\sigma)\}_{t=0}^{\infty}$ in which each c_t is \mathcal{F}_t -measurable; that is, $c_t : H_t \rightarrow \mathbf{R}_{++}$. Each trader is endowed with a particular consumption plan e^i , called the *endowment stream*.

Trader i has a utility function $U_i(c)$ which assigns to each consumption plan the expected present discounted value of the plan's payoff stream with respect to some beliefs. Specifically, trader i has beliefs about the evolution of states, which are represented by a probability distribution p^i on Σ . She in fact believes that states are iid draws from probability distribution ρ^i on S . She also has a payoff function $u_i : \mathbf{R}_{++} \rightarrow \mathbf{R}$ on consumptions and a discount factor β_i strictly between 0 and 1. The utility of a consumption plan is

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

We will assume throughout the following properties of payoff functions:

A. 1. *The payoff functions u_i are C^1 , strictly concave, strictly monotonic, and satisfy an Inada condition at 0.*

Each trader's endowment is a consumption plan. We assume that endowments are strictly positive and that the aggregate endowment is uniformly bounded. Let $e_t(\sigma^t) = \sum_i e_t^i(\sigma^t)$ denote the aggregate endowment at date t on path σ .

A. 2. For all traders i , all dates t and all paths σ , $e_t^i(\sigma^t) > 0$. Furthermore, there are numbers $F \geq f > 0$ such that $f \leq \inf_{t,\sigma} e_t(\sigma^t) \leq \sup_{t,\sigma} e_t(\sigma^t) < F < \infty$.

The upper bound in particular is important to the derivation of our results. The conclusions hold when F grows slowly enough, but may fail when F grows too quickly.

The following assumption about beliefs will be maintained for convenience throughout the paper. Any trader who violates this axiom would not survive, so there is no cost to discarding them now.

A. 3. For each trader i and $s \in S$, if $\rho(s) > 0$ then $\rho^i(s) > 0$.

2.2 Pareto Optimality

Standard arguments show that in this economy, Pareto optimal consumption allocations can be characterized as maxima of weighted-average social welfare functions. If $c^* = (c^{1*}, \dots, c^{I*})$ is a Pareto optimal allocation of resources, then there is a non-negative vector of welfare weights $(\lambda^1, \dots, \lambda^I) \neq \mathbf{0}$ such that c^* solves the problem

$$\begin{aligned} & \max_{(c^1, \dots, c^I)} \sum_i \lambda^i U_i(c) \\ & \text{such that } \sum_i c^i - e \leq \mathbf{0} \\ & \forall t, \sigma \ c_t^i(\sigma^t) \geq 0 \end{aligned} \tag{1}$$

where $e_t = \sum_i e_t^i$. The first order conditions for problem 1 are:

For all t there is a positive \mathcal{F}_{t-1} -measurable random variable η_t such that

$$\lambda^i \beta_i^t u_i'(c_t^i(\sigma^t)) \prod_s \rho_i(s)^{n_s^i(\sigma^t)} - \eta_t(\sigma^t) = 0 \tag{2}$$

almost surely, and

$$\sum_i c_t^i(\sigma^t) = e_t(\sigma^t) \tag{3}$$

These equations will be used to characterize the long-run behavior of consumption plans for individuals with different preferences, discount factors and beliefs.

2.3 Competitive Equilibrium

A price system is a price for consumption in each state at each date such that the value of each trader's endowment is finite.

Definition 1. *A function $p : \Sigma \rightarrow \prod_{t=0}^{\infty} \mathbf{R}_{++}$ such that each p_t is \mathcal{F}_{t-1} -measurable is a present value price system if, for all traders i , $\sum_{\sigma^t \in H} p_t(\sigma^t) \cdot e_t^i(\sigma^t) < \infty$.*

As is usual, a competitive equilibrium is a price system and, for each trader, a consumption plan which is affordable and preference maximal on the budget set such that all the plans are mutually feasible. The existence of competitive equilibrium price systems and consumption plans is straightforward to prove. See Peleg and Yaari (1970).

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 account in partial history (σ^t, s) and zero otherwise. The price of the state s Arrow security in units of consumption at partial history σ^t is the price of consumption at partial history (σ^t, s) in terms of consumption at partial history σ^t , which is $\tilde{q}_t^s(\sigma) \equiv p_{t+1}(\sigma^t, s)/p_t(\sigma)$. Under our assumptions, every equilibrium present-value price system will be strictly positive (because every partial history is believed to have positive probability, and because conditional preferences for consumption in each possible state are non-satiated), and so all current value prices are well defined. We will be particularly interested in normalized current-value prices: $q^s(\sigma^t) = \tilde{q}^s(\sigma^t) / \sum_{\nu} \tilde{q}^{\nu}(\sigma^t)$.

It is not obvious what it means to price an Arrow security (or any other asset) correctly. The literature contains notions such as (for long-lived assets): Prices should equal the present discounted value of the dividend stream. But in a world in which traders' discount factors are not all identical, it is not intuitively obvious what the discount rate should be; and to say that the 'correct' discount rate is the 'market' discount rate is to beg the question. Is the market discount rate, after all, correct? With Arrow securities, it seems that prices should be related to the likelihood of the states. But in a market with endowment risk in which attitudes to risk are not all identical, risk premia should matter too, and again in a market in which not all traders have the same attitude to risk, it is not obvious what the correct risk premium is. So that we can meaningfully talk about correct prices, we make the following assumption:

A.4. *There is an $e > 0$ such that for all paths σ and dates t , $e_t(\sigma^t) \equiv e$.*

That is, there is no aggregate risk. The only risk in this economy is who gets what, not how much is to be gotten. The reason for this assumption is the following result:

Theorem 1. *Assume A.1–4.*

1. *If all traders have identical beliefs ρ' , then for all dates t and paths σ , $q_t^s(\sigma^t) = \rho'$.*
2. *On each path σ at each date t and for all $\epsilon > 0$ there is a $\delta > 0$ such that if $|c_t^i(\sigma^t) - e| < \delta$, then $\|q_t(\sigma^t) - \rho^i\| < \epsilon$.*

A consequence of the first point is that in a rational expectations equilibrium, the Arrow securities spot prices will be ρ , the true probabilities of the state realizations. Thus we now know what it means for assets to be ‘correctly’ priced. The second point asserts that when one trader is dominant in the sense that his demand is very large relative to that of the other traders, the equilibrium will primarily reflect her beliefs. The proof of both points is elementary, in the first case from a calculation and in the second from a calculation and the upper hemi-continuity of the equilibrium correspondence.

3 Selection

By ‘selection’ we mean the idea that markets identify those traders with the most accurate information, and the market prices come to reflect their beliefs. We illustrate this idea with an example.

3.1 A Leading Example

Consider an economy with two states of the world, $S = \{A, B\}$. States are iid draws, and the probability of state A at any date t is ρ . Arrow securities are traded for each state at each date, so markets are complete. Trader i has endowment $e^i(\sigma)$, and suppose that the aggregate endowment is a constant $e > 0$ at each date and event. Traders have

logarithmic utility, and have identical discount factors, $0 < \beta < 1$. Trader i knows that the state process is iid, and believes that A will occur in any given period with probability ρ_i . This is basically just a big Cobb-Douglas economy, and equilibrium is easy to compute. Let w_0^i denote the present discounted value of trader i 's endowment stream, and let $w_t^i(\sigma^t)$ denote the amount of wealth which is transferred to partial history σ^t , measured in current units. The optimal consumption plan for trader i is to spend fraction $(1 - \beta)\beta^t \rho_i^{n_t^A(\sigma)} (1 - \rho_i)^{n_t^B(\sigma)}$ of w_0^i on consumption at date-event σ^t . This can be described recursively as follows: In each period, eat fraction $1 - \beta$ of beginning wealth, w_t^i , and invest the residual, βw_t^i , in such a manner that the fraction ρ_i of date- t investment which is allocated to the asset which pays off in state A is ρ_i . let q_t^A denote the prices of the security which pays out 1 in state A at date t and 0 otherwise; let q_t^B denote the corresponding price for the other date- t Arrow security. Given the beginning-of-period wealth and the market price, trader i 's end-of-period wealth is determined only by that period's state:

$$\begin{aligned} w_{t+1}^i(A) &= \frac{\beta \rho_i w_t^i}{q_t^A} \\ w_{t+1}^i(B) &= \frac{\beta(1 - \rho_i) w_t^i}{q_t^B} \end{aligned}$$

Each unit of Arrow security pays off 1 in its state, and the total payoff in that state must be the total wealth invested in that asset. Thus in equilibrium,

$$\sum_i \frac{\beta \rho^i w_t^i}{q_t^s} = \sum_j \beta w_t^j,$$

and so the price of asset s at date t is

$$\begin{aligned} q_t^A &= \frac{\sum_i \rho^i w_t^i}{\sum_j w_t^j} \\ &= \sum_i \rho^i r_t^i \\ q_t^B &= \sum_i (1 - \rho^i) r_t^i \end{aligned}$$

where r_t^i is the *share* of date t wealth belonging to trader i . That is, the price of asset s at date t is the *wealth share weighted* average of beliefs. So at any date, the market

prices states by averaging traders' beliefs. Of course there is no reason for this average to be correct since the initial distribution of wealth was arbitrary. But the process of allocating the assets and then paying them off reallocates wealth. The distribution of wealth evolves through time, and the limit distribution of wealth determines prices in the long run. We can work this out to see how the market 'learns'. In this model it should be clear what "correct" asset pricing means. If all traders had rational expectations, then the price of the A Arrow security at any point in the date-event tree would be ρ , and the price of the B Arrow security would be $1 - \rho$.

Let $1_A(s)$ and $1_B(s)$ denote the indicator functions for states A and B , respectively. Along any path σ of states,

$$w_{t+1}^i(\sigma^{t+1}) = \beta_i \left(\frac{\rho^i}{q_t^A(\sigma^t)} \right)^{1_A(\sigma_{t+1})} \left(\frac{1 - \rho^i}{q_t^B(\sigma^t)} \right)^{1_B(\sigma_{t+1})} w_t^i(\sigma^t),$$

and so the ratio of i s wealth share to j s evolves as follows:

$$\frac{r_{t+1}^i(\sigma^{t+1})}{r_{t+1}^j(\sigma^{t+1})} = \left(\frac{\rho^i}{\rho^j} \right)^{1_A(\sigma_{t+1})} \left(\frac{1 - \rho^i}{1 - \rho^j} \right)^{1_B(\sigma_{t+1})} \frac{r_t^i(\sigma^t)}{r_t^j(\sigma^t)}.$$

This evolution is more readily analyzed in its log form:

$$\log \frac{r_{t+1}^i(\sigma^{t+1})}{r_{t+1}^j(\sigma^{t+1})} = 1_A(\sigma_{t+1}) \log \left(\frac{\rho^i}{\rho^j} \right) + 1_B(\sigma_{t+1}) \log \left(\frac{1 - \rho^i}{1 - \rho^j} \right) + \log \frac{r_t^i(\sigma^t)}{r_t^j(\sigma^t)}. \quad (4)$$

To understand how the market can learn, consider a Bayesian whose prior beliefs about state evolution contain I iid models in her support, $\{\rho^1, \dots, \rho^I\}$, and let r_t^i denote the probability she assigns to model i posterior to the first $t - 1$ observations. The Bayesian rule for posterior revision is exactly that of equation (4). The market is a Bayesian learner. The evolution of the distribution of wealth parallels the evolution of posterior beliefs. Market prices are wealth share-weighted averages of the traders' models, and so the pricing function for assets is identical to the rule which assigns a predictive distribution on outcomes to any prior beliefs on states. In other words, the price of asset A in this example is the probability the Bayesian learner would assign to the event that the next state realization will be A .

From these observations we can draw several conclusions. If some trader holds correct beliefs, then in the long run his wealth share will converge to 1, and the market price will converge to ρ . The assets are priced correctly in the long run. Second, if no model is correct, the posterior probability of any model whose Kullback-Leibler distance from the true distribution is not minimal converges a.s. to 0. In this example, selection cannot make the market do better than the best-informed trader. In particular, if there is a unique trader whose beliefs ρ^i are closest to the truth, then prices converge in the long run to ρ^i almost surely, and so assets are mispriced.

3.2 Selection in Complete IID Markets

Traders are characterized by three objects: A payoff function u_i , a discount factor β_i and a belief ρ_i . We will see that, so long as payoff functions satisfy the Inada condition, they are irrelevant to survival. Only beliefs and discount factors matter. We would expect that discount factors matter in a straightforward way: Higher discount factors reflect a greater willingness to trade present for future consumption, and so they should favor survival. Similarly, traders will be willing to trade consumption on unlikely paths for consumption on those they think more likely. Those traders who allocate the most to the highest-probability paths have a survival advantage. This advantage, as we will see, can be measured by the *Kullback-Leibler* distance of beliefs from the truth, the relative entropy of ρ with respect to ρ^i :

$$I_\rho(\rho_i) = \sum_s \rho^s \log \frac{\rho^s}{\rho_i^s}$$

The Kullback-Leibler distance is not a true metric. But it is non-negative, and 0 iff $\rho^i = \rho$.² Assumption A.3. ensures that $I_\rho(\rho^i) < \infty$ (and this is its only role).

Our results will demonstrate several varieties of asymptotic experience for traders in iid economies. Traders can vanish, they can survive, and the survivors can be divided into those who are negligible and those who are not. Definitions are as follows:

Definition 2. *Trader i vanishes on path σ if $\lim_t c_t^i(\sigma^t) = 0$. She survives on path σ if $\limsup_t c_t^i(\sigma^t) > 0$. A survivor i is negligible on path σ if for all $0 < r < 1$, $\lim_{T \rightarrow \infty} (1/T) |\{t \leq T : c_t^i(\sigma^t) > r e_t(\sigma^t)\}| = 0$. Otherwise she is non-negligible.*

In the long run, traders can either vanish or not, in which case they survive. There are two distinct modes of survival. A negligible trader is someone who consumes a given positive share of resources infinitely often, but so infrequently that the long-run fraction of time in which this happens is 0. The definitions of vanishing, surviving and being negligible are reminiscent of transience, recurrence and null-recurrence in the theory of Markov chains.

3.3 The Basic Equations

Our method uses the first order conditions to solve for the optimal consumption of each trader i in terms of the consumption of some particular trader, say trader 1. We then use the feasibility constraint to solve for trader 1's consumption. The fact that we can do this only implicitly is not too much of a bother.

Let $\kappa_i = \lambda_1/\lambda_i$. From equation 2 we get that

$$\frac{u'_i(c_t^i(\sigma^t))}{u'_1(c_t^1(\sigma^t))} = \kappa_i \left(\frac{\beta_1}{\beta_i} \right)^t \prod_{s \in S} \left(\frac{\rho_s^1}{\rho_s^i} \right)^{n_t^s(\sigma^t)} \quad (5)$$

It will sometimes be convenient to have this equation in its log form:

$$\log \frac{u'_i(c_t^i(\sigma^t))}{u'_1(c_t^1(\sigma^t))} = \log \kappa_i + t \log \frac{\beta_1}{\beta_i} - \sum_s n_t^s(\sigma^t) \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right).$$

We can decompose the evolution of the ratio of marginal utilities into two pieces: The mean direction of motion, and a mean-0 stochastic component.

$$\begin{aligned} \log \frac{u'_i(c_t^i(\sigma^t))}{u'_1(c_t^1(\sigma^t))} &= \log \kappa_i + t \log \frac{\beta_1}{\beta_i} - t \sum_s \rho_s \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right) + \\ &\quad \sum_s (n_t^s(\sigma^t) - t\rho_s) \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right) \\ &= \log \kappa_i + t(\log \beta_1 - I_\rho(\rho^1)) - t(\log \beta_i - I_\rho(\rho^i)) + \\ &\quad \sum_s (n_t^s(\sigma^t) - t\rho_s) \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right) \end{aligned}$$

The mean term in the preceding equation gives a first order characterization of traders' long run fates.

Definition 3. *Trader i 's survival index is $\mathbf{s}_i = \log \beta_i - I_\rho(\rho^i)$.*

Then

$$\log \frac{u'_i(c_t^i(\sigma^t))}{u'_1(c_t^1(\sigma^t))} = \log \kappa_i + t(\mathbf{s}_1 - \mathbf{s}_i) + \sum_s (n_t^s(\sigma^t) - t\rho_s) \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right) \quad (6)$$

3.4 Who Survives? — Necessity

Necessary conditions for survival have been studied before, notably by Blume and Easley (2006) and Sandroni (2000). In this economy, a sufficient condition guaranteeing that trader i vanishes is that trader i 's survival index is not maximal among the survival index of all traders. Consequently, a necessary condition for survival is that the survival index be maximal.

Theorem 2. *Assume A.1–3. If $\mathbf{s}_i < \max_j \mathbf{s}_j$, then trader i vanishes.*

The analysis compares one trader, say trader 1, to other traders in the economy. We use equation (6) to show that if trader i has a larger survival index, than trader 1, trader 1 must vanish. The first step is to relate long-run survival outcomes to the ratios of traders marginal utilities, the lhs of (6).

Lemma 1. *If on a sample path σ , $\log u'_i(c_t^i(\sigma^t))/u'_1(c_t^1(\sigma^t)) \rightarrow -\infty$ for some trader i , then $\lim_t c_t^1(\sigma^t) = 0$. If $\limsup \min_i \log u'_i(c_t^i(\sigma^t))/u'_1(c_t^1(\sigma^t)) > -\infty$, then $\limsup_t c_t^1(\sigma^t) > 0$.*

Proof: Suppose first that the limit of the log of the ratio of marginal utilities converges to $-\infty$ along a path σ . This can happen in one of two ways: If either the denominator converges to 0 or the numerator diverges to infinity. It must be the latter, because the numerator is bounded below by $u'_i(F) > 0$. Consequently, on any such path, $c_t^i(\sigma^t) \rightarrow 0$.

In every period t , there is a trader $i(t)$ who consumes at least $c_{i(t)}(\sigma^t) \geq f/I$. If trader 1 were to vanish, then $\lim_t \log u'_{i(t)}(c_{i(t)}^{i(t)}(\sigma^t))/u'_1(c_t^1(\sigma^t))$ converges to $-\infty$ (since the number of traders is finite). But if the limsup condition is satisfied, then there is an ϵ such that $\log u'_{i(t)}(c_{i(t)}^{i(t)}(\sigma^t))/u'_1(c_t^1(\sigma^t)) > \epsilon$ infinitely often. \square

Proof of Theorem 2. We prove this theorem by examining equation (6). Take time averages of both sides and observe that for each s , $t^{-1}(n_t^s - \rho_s)$ converges p -almost surely to 0, to conclude that for almost all paths σ^t ,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \frac{u'_i(c_t^i(\sigma^t))}{u'_1(c_t^1(\sigma^t))} = \mathbf{s}_1 - \mathbf{s}_i.$$

If \mathbf{s}_1 is not maximal, there is an i for which $\mathbf{s}_1 - \mathbf{s}_i < -\epsilon < 0$. For almost all paths σ^t there is a T such that if $t > T$, then $u'_i(c_t^i(\sigma^t))/u'_1(c_t^1(\sigma^t)) < -\epsilon t$. According to lemma 1, trader 1 vanishes. \square

This result is a consequence of the SLLN. Blume and Easley (2006) extend this result to identify necessary conditions for survival in many different, non-IID settings.

3.5 Market Equilibrium — Selection

The implications for long-run asset pricing are already illustrated in the example which began this section.

Corollary 1. *If there is a unique trader i with minimal survival index \mathbf{s}_i among the trader population, then market prices converge to ρ^i almost surely.*

This Corollary is an immediate consequence of Theorems 1 and 2. If only trader i has maximal survival index, then almost surely all other traders vanish and q_t converges to ρ^i . The beliefs of the trader with minimal survival index may not be correct, in which case Arrow securities are incorrectly priced in the long run. This may happen because no trader has correct beliefs, or because a trader's incorrect beliefs are compensated

for by a higher discount factor. In the latter case, allowing for heterogeneous discount factors, the prices need not converge to the most accurate beliefs present in the market.

Using the tools of Blume and Easley (2006) we can extend this result in various ways. For instance, we can provide a survival index analysis of (finite state) Markov economies with traders who hold Markov models of the economy or even traders who hold (misspecified) iid models. If all traders are Bayesian learners satisfying certain regularity conditions, and the truth is in the support of their beliefs, then all will eventually learn the true state distribution and so prices will ultimately be correct. But those traders with low-dimensional belief supports will learn faster than those with higher-dimensional belief supports, and prices will converge to the true prices at the faster rate.

4 Balancing

When a single trader (type) has the highest survival index, market prices converge to his view of the world. There is no room for balancing of different beliefs because, in the long run, there is only one belief and discount factor present in the market. But if the market process is more complicated than the world view of any single trader so that no trader has correct beliefs, or if traders are asymmetrically informed, it is possible that multiple traders could have maximal survival index. Will all such traders survive, and what are the implications for sufficiency?

4.1 Who Survives? — Sufficiency

Theorem 2 shows that traders with survival indices that are less than maximal in the population vanish. This does not imply that all those with maximal survival indices survive. The rhs of equation (6) is a random walk, and the analysis of the previous section is based on an analysis of the mean drift of the rhs of equation (6). Theorem 2 shows that a non-zero drift has implications for the survival of some trader. When two traders with maximal survival indices are compared, the drift of the walk is 0, and further analysis of equations (5) and (6) is required. Since the long run behavior of

a random walk depends upon the dimension of the space being walked through, our results will depend upon the number of states \mathbf{s} .

More definitions are required. For a probability distribution θ on S , define the vector of log-probabilities: $\text{lo}(\theta) = (\log \theta(s)/\theta(\mathbf{s}))_{s=1}^{\mathbf{s}-1}$. Let Sur denote the set of traders with maximal survival index. Theorem 1 indicates that these are the only potential survivors. The fate of a trader in Sur is determined by how her beliefs, as represented by $\text{lo}(\rho^i)$, are positioned relative to the beliefs of the other traders in Sur . Denote by $C\{\text{lo}(\rho^j)\}_{j=1}^I$ the closed convex cone generated by the log-probability vectors of the traders.

Definition 4. *Trader i is interior if $\text{lo}(\rho^i)$ is in the relative interior of $C\{\text{lo}(\rho^j)\}_{j=1}^I$. She is extremal if $\text{lo}(\rho^i)$ is an extreme point, that is, not a non-negative linear combination of the other $\text{lo}(\rho^j)$, and boundary otherwise.*

We are interested in markets with heterogeneous potential survivors. To simplify the statements of theorems it will be convenient to focus on the case where Sur contains at least two traders, and no two traders have identical beliefs.³

Theorem 3. *Assume A.1–3, and suppose $\mathbf{s} \leq 3$ and $0 < r < 1$.*

1. *If $i \in \text{Sur}$, then trader i survives.*
2. *Trader i is negligible if and only if she is not extremal.*
3. *For each trader i , $\lim_{T \rightarrow \infty} (1/T) |t \leq T : c_t^i < re_t| > 0$ a.s.*
4. *If trader i is extremal, $\lim_{T \rightarrow \infty} (1/T) |\{t \leq T : c_t^i > re_t\}| > 0$ a.s.*

When $\mathbf{s} \leq 3$, a maximal survival index is sufficient (as well as necessary) for survival. But how one survives depends upon one's position in the group of survivors. Boundary and interior survivors are negligible. The fraction of time they consume a positive share of aggregate endowment is 0. Extremal traders, on the other hand, have highly volatile consumption. The fraction of time each consumes an arbitrarily small share of aggregate endowment is positive, as is the fraction of time each consumes nearly all of the aggregate endowment.

When $\mathbf{s} > 3$, the picture is even more stark. Interior traders vanish. Maximality of a trader's survival index is no longer a sufficient condition for survival.

Theorem 4. Assume A.1–3, and suppose $\mathbf{s} > 3$ and $0 < r < 1$.

1. If i is extremal or boundary, then trader i survives.
2. Trader i vanishes if she is interior, and is negligible if she is boundary.
3. For each trader i , $\lim_{T \rightarrow \infty} (1/T) |\{t \leq T : c_t^i < re_t\}| > 0$ a.s.
4. If trader i is extremal, $\lim_{T \rightarrow \infty} (1/T) |\{t \leq T : c_t^i > re_t\}| > 0$ a.s.

Proof. Suppose that trader 1 is in **Sur**, and consider trader $i \neq 1$. Since 1 and i are both in **Sur**, they have identical survival indices, and so the right hand side of equation (6) becomes

$$\hat{z}_{it}(\sigma) = \log \kappa_i + \sum_s (n_t^s(\sigma^t) - tp_s) \log \frac{\rho_s^1}{\rho_s^i}$$

We want to investigate if the limit of \hat{z}^i is infinitely often arbitrarily large, converges to $-\infty$, or is bounded above but also bounded away from $-\infty$ infinitely often. To answer this question the constants on the right can be discarded, and we are left with

$$z_{it}(\sigma) = \sum_s (n_t^s(\sigma^t) - tp_s) \log \frac{\rho_s^1}{\rho_s^i}.$$

Thus the vector $z_t(\sigma) = (z_t(\sigma)_{i \in \text{Sur}})$ is a random walk in $\mathbf{R}^{\mathbf{I}}$. Rewriting,

$$z_t(\sigma) = \tilde{A} \cdot \tilde{w}_t(\sigma)$$

where $\tilde{w}_t(\sigma)$ is the random walk in $\mathbf{R}^{\mathbf{s}}$ whose s th term is $n_t^s(\sigma^t) - tp_s$, and the $(i-1)$ st row of the $(I-1) \times \mathbf{s}$ matrix \tilde{A} is the vector $(\log(\rho_s^1) - \log(\rho_s^i))_{s=1}^{\mathbf{s}}$. The random walk $\{\tilde{w}_t(\sigma)\}_{t=1}^{\infty}$ is an $\mathbf{s} - 1$ dimensional random walk that lives in the subspace W of $\mathbf{R}^{\mathbf{s}}$ consisting of all vectors v such that $\sum_s v_s = 0$. Simple algebra shows that the random walk can be rewritten as

$$z_t(\sigma) = A \cdot w_t(\sigma)$$

where the $(i-1)$ st row of the $(I-1) \times \mathbf{s}$ matrix A is the vector $(\log(\rho_s^1) - \log(\rho_s^i))_{s=1}^{\mathbf{s}-1}$ and $w_t(\sigma)$ the vector in $\mathbf{R}^{\mathbf{s}-1}$ whose coordinates are the first $\mathbf{s} - 1$ coordinates of $\tilde{w}_t(\sigma)$.

We consider three cases of matrices A . Case 1: There is a direction x such that $Ax \gg 0$. Case 2: For all directions x there is a row a_i such that $a_i x < 0$. Case 3: There

is a direction x such that $Ax \geq 0$ but no direction x for which Ax is strictly positive. First we establish conditions under which each of these cases occurs.

Lemma 2. *The inequality system $Ax \gg 0$ has a solution iff $\text{lo}(\rho^1)$ is an extreme point of the convex hull of the $\text{lo}_i(\rho^i) \in \text{Sur}$, that is, iff trader 1 is extremal. If trader 1 is interior, then for all directions x there is an a_i such that $a_i x < 0$. If neither of these two cases occurs, $\text{lo}(\rho^1)$ is on the boundary of the convex hull of the $\text{lo}(\rho^i)$.*

Proof of Lemma 2. A theorem of the alternative (see the Appendix) states that there is an x such that $Ax \gg 0$ iff no non-trivial, non-negative linear combination of the rows is 0. That is, there is no non-trivial and non-negative set of weights $\{\lambda_i\}_{i \geq 2}$ such that $\sum_{i \geq 2} \lambda_i \text{lo}(\rho^i) = \text{lo}(\rho^1)$. In other words, a strictly positive direction x exists iff $\text{lo}(\rho^1)$ is extremal in $C\{\text{lo}(\rho^j)\}_{j=1}^I$.

A Theorem of the Alternative (Gale 1960) also shows that either $Ax \geq 0$ (and not equal to 0 in every component) has a solution, or the a_i are linearly dependent with strictly positive weights. Thus if trader 1 is interior, there is no non-negative direction, and in the remaining case, where $\text{lo}(\rho^1)$ is on the boundary of the convex hull of the $\text{lo}(\rho^i)$ but not an extreme point, a non-negative but not strictly positive direction must exist. \square

Now we examine the three types of traders.

Extremal Traders: Suppose that trader 1 is extremal. Then there is an open cone C_1 of directions such that for all $w \in C_1$ and for all traders $i \neq 1$, $w \cdot (\text{lo}(\rho^1) - \text{lo}(\rho^i)) > 0$. Furthermore, for all consumption shares $r < 1$ there is a bound $b(r)$ such that if, for all $i \neq 1$, $w_t(\sigma) \cdot (\text{lo}(\rho^1) - \text{lo}(\rho^i)) > b(r)$, then $c_t^1(\sigma^t) > re_t(\sigma^t)$. The set of such w values is the open cone less a compact set containing the origin. When $\mathbf{s} \leq 3$ (so that $\mathbf{s} - 1 \leq 2$) such sets are positive recurrent. Consequently $\lim_{T \rightarrow \infty} (1/T) |\{t \leq T : c_t^1(\sigma^t) > re_t(\sigma^t)\}|$ exists and is positive.

Interior Traders: If trader 1 is interior, then for any direction of the walk, there is a trader i such that $\log u'_i(c_t^i)/u'_1(c_t^1)$ is arbitrarily negative when the walk is far enough out in that direction. Consumption for trader 1 is bounded away from 0 only on compact sets containing the origin. Such sets are recurrent but rare for two dimensional walks,

and transient for walks in dimension three or higher. So when $s \leq 3$, trader 1 is negligible; otherwise trader 1 vanishes.

Boundary Traders: For boundary traders, the set of directions D for which z_t does not grow negatively in some direction is lower dimensional. Hence the set of vectors w for which z_t is bounded from below is an open neighborhood of D . Such sets are in the plane and recurrent but rare in higher dimensions. \square

The following figures demonstrate the geometry of Theorem 3 when $s = 3$. The left illustration in Figure 1 plots the log-odds ratios of five surviving beliefs. The discount factors for all traders cannot be all the same. The log-odds vector of the true distribution can be anywhere in the plane, but it is most entertaining to think of it as being inside the triangle, and perhaps even coincident with E . Figure 1 plots the log-

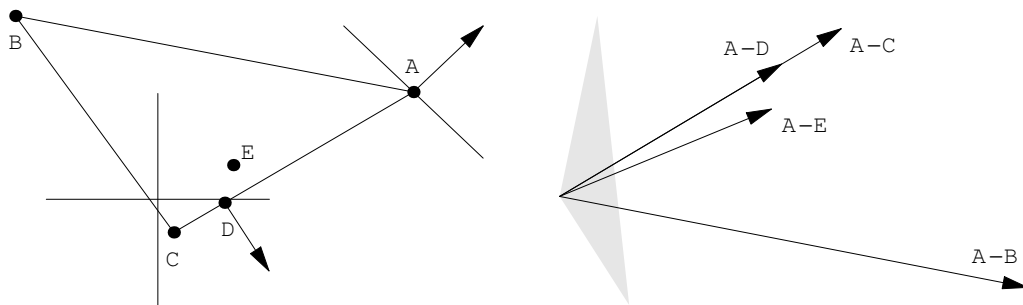
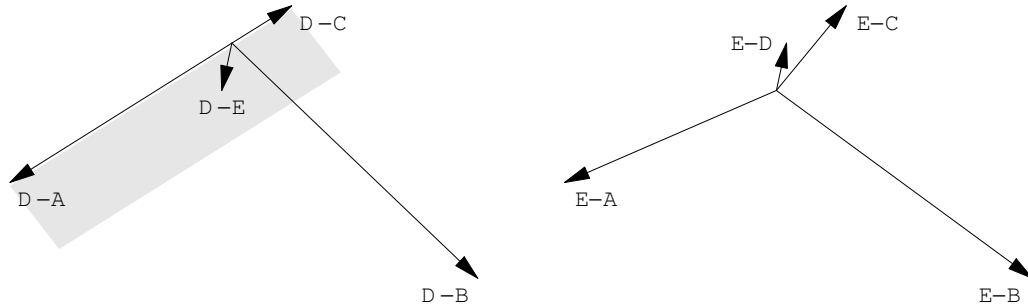


Figure 1: Five Surviving Beliefs

odds ratios of five surviving beliefs. (The discount factors for all traders cannot be all the same.) The cone on the right indicates the directions in which the random walk can move so as to increase the wealth share of extremal trader A relative to all other traders. In Figure 2, the direction of increase for boundary trader D is the intersection of three half spaces while on the right there is no direction of increase for interior trader E . Trader E nonetheless survives because her wealth share is positive on a recurrent set, such as any neighborhood of the origin.

Figure 2: Survival of Beliefs D and E

4.2 Market Equilibrium — Balancing

The implication for market equilibrium from the existence of multiple survivors is perhaps surprising:

Corollary 2. *If multiple traders have maximal survival index, then for all extremal traders i and all $\epsilon > 0$, $|q_t - \rho^i| < \epsilon$ infinitely often. If $\mathfrak{s} > 3$ it is possible that for $\epsilon > 0$ sufficiently small, the event $|q_t - \rho|$ is transient, even if some survivor has rational expectations.*

With multiple survivors, asset prices are volatile. Furthermore, asset prices need be approximately right; specifically, approximately right prices may be transient. One might hope that, nonetheless, the time average of prices is approximately correct. We believe that this weaker notion of correct asset pricing may fail, and we hope to have a proof shortly.

Figure 4.2 illustrates some of the possibilities for prices with multiple survivors in the leading example of the previous section with log utility. In this figure the true distribution is ρ . The closed curve connecting points P , Q and R is a curve of constant relative entropy, in this case 0.18. Suppose all traders have identical factors, and all have beliefs which are on or outside the curve. Those traders with beliefs outside the curve will vanish. Suppose now that Surv contains three traders with beliefs P , Q and R . All three will survive. The equilibrium price will wander around inside the convex hull of these three points. As the three points are drawn, ρ is in their convex hull, and it

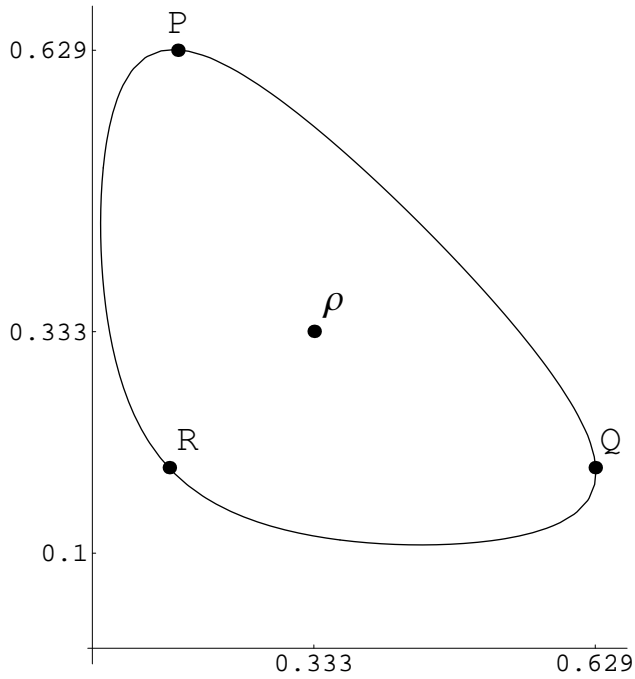


Figure 3: Multiple survivors, $s = 3$.

is at least possible that the average behavior of prices over time could be approximately correct. On the other hand, suppose Q and R were higher up on the iso-relative entropy curve, nearer to P . It is possible to arrange them so that ρ is no longer in the convex hull, and so the long-run time average of prices would be nowhere near ρ . Finally, consider moving point Q off the curve. If it moves in, this trader is the unique survivor, and selection dictates that prices converge to Q . On the other hand, if Q moves out, this trader is no longer a survivor. The two survivors are P and R , and in the long run prices will move up and down on the line segment connecting these two points. Again there is no connection between the long-run behavior of prices and ρ .

5 Conclusion

This analysis suggests that, contrary to Henry Herrman’s view in the epigram which begins this paper, the market knows not better but only as well as the analysts do (unless there is someone in the market who knows better than the analysts). The market can be no better informed than the most fit trader according to the fitness index metric, and if there are several most-fit traders with distinct beliefs, then the market beliefs as expressed in limit equilibrium prices may fail to converge.

The necessity of a maximal survival index for long-run selection has frequently been noticed, including Sandroni (2000) and Blume and Easley (2006). The observation that it is not sufficient, and the sufficient condition derived here for the iid economy, are new. It is not surprising that only beliefs matter for the sufficient condition. Since discount factors are non-stochastic, they are part of the “mean term” which gives the necessary condition. Were discount factors stochastic, the analysis of necessary conditions for survival would remain essentially unchanged, but deviations from the mean log-discount factor would appear in the sufficient condition.

It is clear how this analysis extends to Markov and other stochastic complete market environments in which the necessary conditions for survival described in Blume and Easley (2006) can be summarized in a real-valued index. We have not attempted to find sufficient conditions for survival at the level of generality of that analysis.

Perhaps an even more compelling question is an asymptotic analysis of wealth shares and prices when markets are incomplete. (Blume and Easley 2006) have some simple examples of how incomplete markets can select for the wrong trader, which makes the market, in the limit, less smart than its smartest trader. In the most compelling example, an excessively optimistic trader oversaves, and thus comes to dominate in the limit. Becker et. al. (2006) analyse a market in which the only assets are money and one risky asset, so that (with enough states) the market is incomplete. They too find that long-run price volatility with multiple survivors. In particular, if two or more traders survive in the long run, then the each trader consumes arbitrarily little infinitely often.

In studying prediction markets like those contracts traded on Iowa Electronic Market which make book on political races, it is important to take account of learning through prices, and to entertain the possibility that the accurate performance of these

markets is due at least as much to trader learning from prices (as opposed to more ‘outside’ information) as it is to market selection. In our view, this is less important when it comes to large markets for securities and other financial assets. This is not to say that learning is not important; surely it is. But these markets are sufficiently complicated, and trading occurs for so many diverse motives, that the possibility of consistent learning rules seems, to us, remote. This leaves room for the market to be, in the long run, smarter than its traders; and so we are led to ask, how is the market’s learning experience different than that of its traders.

Appendix: Linear Algebra

Let A be an $n \times m$ matrix.

Theorem 1. *One and only one of the following equation systems has a solution:*

$$Ax \gg 0 \tag{7}$$

$$\begin{aligned} yA &= 0, \\ y &\geq 0, \quad y \neq 0 \end{aligned} \tag{8}$$

This Theorem is an immediate consequence of the following theorem, due to Fan, Glicksberg, and Hoffman (1957), concerning m convex functions, each mapping the non-empty convex set K to \mathbf{R} .

Theorem 2 (Fan et. al.). *One and only one of the following alternatives holds:*

1. *The system of inequalities $f_i(x) < 0$, $i = 1, \dots, m$, $x \in K$ has a solution;*
2. *There are non-negative scalars λ_i , not all 0, such that $\sum_i \lambda_i f_i(x) \geq 0$ for all $x \in K$.*

Proof of Theorem 1. Take $f_i(x) = -a_i x$, where a_i is the i th row of the matrix A . The f_i are convex functions and W is a convex set. If (7) has no solution, then according

to Fan et. al., there are non-negative scalars y_i not all 0 such that for all $x \in W$,

$$\sum_i y_i(-a_i x) \geq 0. \quad (9)$$

In particular, $\sum_i y_i a_i x = 0$ for all x , because if not it will be possible to make this term arbitrarily negative by suitable choice of x , and so the inequality will be violated for some x . This will be true if and only if $\sum_i y_i a_i = 0$. \square

Notes

¹“The End of Management”, Time Magazine Bonus Section, August 2004, <http://www.time.com/660965-1,00.html>. For a popular account of prediction markets, see Surowiecki (2003).

²In fact, it is jointly convex in (ρ, ρ^i) , but we will not need to make use of this fact.

³If two traders i and j in \mathbf{Sur} have identical beliefs, then they have identical discount factors, and the ratio of their marginal utilities is a constant on every path. This fact and assumption A.2. imply that there are constants $0 < k < K$ such that almost surely, $kc_t^i(\sigma^t) < c_t^j(\sigma^t) < Kc_t^i(\sigma^t)$.

References

- BLUME, L., AND D. EASLEY (1992): “Evolution and Market Behavior,” *Journal of Economic Theory*, 58(1), 9–40.
- 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.
- COOTNER, P. (1964): *The Random Character of Stock Market Prices*. MIT Press, Cambridge, MA.
- FAMA, E. (1965): “The Behavior of Stock Market Prices,” *Journal of Business*, 38(1), 34–105.
- FAN, K., I. GLICKSBERG, AND A. J. HOFFMAN (1957): “Systems of Inequalities Involving Convex Functions,” *Proceedings of the American Mathematical Society*, 8(3), 617–622.
- FIGLEWSKI, S. (1979): “Subjective Information and market Efficiency in a Betting Model,” *Journal of Political Economy*, 87(1), 75–88.
- GALE, D. (1960): *Linear Economic Models*. McGraw-Hill, New York.
- MAILATH, G., AND A. SANDRONI (2003): “Market Selection and Asymmetric Information,” *Review of Economic Studies*, 70(2), 343–368.
- PELEG, B., AND M. E. YAARI (1970): “Markets with Countably Many Commodities,” *International Economic Review*, 11(3), 369–377.
- SANDRONI, A. (2000): “Do Markets Favor Agents Able to Make Accurate Predictions?,” *Econometrica*, 68(6), 1303–42.
- SNYDER, W. W. (1978): “Horse Racing: Testing the Efficient Markets Model,” *Journal of Finance*, 33(4), 1109–1118.
- SUROWIECKI, J. (2003): “Decisions, Decisions,” *The New Yorker*, March 24, p. 29.