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

Lawrence E. Blume
David Easley

SANTA FE INSTITUTE

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

Lawrence Blume[†][‡] and David Easley[†]

March 2008

† Cornell University
‡ Santa Fe Institute

Abstract: The information content of prices is a central problem in the general equilibrium analysis of competitive markets. Rational expectations equilibrium identifies conditioning simultaneously on contemporaneous prices and private information as the mechanism by which information enters prices. Here we look to the ecology of markets for an explanation of the information content of prices. Markets could select across traders with different beliefs, or, reminiscent of 'the wisdom of crowds', markets could balance the diverse information of many participants. We provide theoretical support in favor of the first mechanism, and against the second. Along the way we demonstrate that the necessary condition for long-run survival in complete markets found in Sandroni (2000) and Blume and Easley (2006) is not sufficient for long run survival. We also demonstrate some surprising behavior of market prices when several trader types with different beliefs survive. This paper continues the research program of Blume and Easley (1992) and (2006).

Corresponding Author:
  Lawrence Blume
  Department of Economics
  Uris Hall
  Cornell University
  Ithaca NY 14853
  USA

  lb19@cornell.edu

'Amazingly, the stock market knows better than the analysts do.'

Henry Herrmann, Chief Investment Officer, Waddell & Reed

'The stock market has predicted nine out of the last five recessions.'

Paul Samuelson, *Newsweek*, 1966.

1   Introduction

Economists, businessmen and the laiety 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, 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 could be said to 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. This seems to be the idea behind arguments for "the wisdom of crowds". 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. Its implications for asset markets were drawn out by Fama (1965) and Cootner (1964). Finally, 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 role of balancing and selection in 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.

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 $S = \{1, \ldots, \mathtt{s}\}$, with cardinality $\mathtt{s} = |S|$. The set of all infinite sequences of states is Σ with representative sequence $\sigma = (\sigma_0, \ldots)$, also called a *path*. σ_t denotes the value of σ at date t, and $\sigma^t = (\sigma_0, \ldots, \sigma_t)$ 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 $H = \cup_{t=0,1,\ldots} H_t$ denote the set of all partial histories.

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 $\rho(s)$. Since the processes and beliefs are iid, counts will be important. Let $n_t^s(\sigma) = |\{\tau \leq t : \sigma_\tau = s\}|$.

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 $\sigma^t = (\sigma_0, \ldots, \sigma_t)$. That is, $p'_t(\sigma) = p'(\{\sigma_0\} \times \cdots \times \{\sigma_t\} \times S \times S \times \cdots)$.

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* $c : \Sigma \to \prod_{t=0}^{\infty} \mathbf{R}_{++}$ 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, $c_t : H_t \to \mathbf{R}_{++}$. 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 presented 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 $u_i : \mathbf{R}_{++} \to \mathbf{R}$ on consumptions and a discount factor β_i strictly between 0 and 1. The utility of a consumption plan is

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

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 $e = \sum_i e^i$; then $e_t(\sigma) = \sum_i e_t^i(\sigma)$ denotes the aggregate endowment at date t on path σ.

A. 2. *There are numbers $\infty > F \geq f > 0$ such that for each trader i, all dates t and all paths σ, $f \leq \inf_{t,\sigma} e_t(\sigma) \leq \sup_{t,\sigma} e_t(\sigma) < F$.*

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 convenient, and entails no essential loss of generality.

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

If there is a possible state s which trader i believes to be impossible, the trader would trade away all claims to the commodity in any partial history in which s is ever reached. Every possible state will almost surely be realized at some date, so the trader will almost surely not survive. Hence there is no cost to discarding such traders at the outset.

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 $c^* = (c^{1*}, \ldots, c^{I*})$ is a Pareto optimal allocation of resources, then there is a non-negative vector of welfare weights $(\lambda^1, \ldots, \lambda^I) \neq \mathbf{0}$ such that c^* solves the problem

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

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

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

almost surely, and

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

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

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, the security's *current value price*, 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: $q_t^s(\sigma) = \tilde{q}_t^s(\sigma)/\sum_\nu \tilde{q}_t^\nu(\sigma)$.

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. Only 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, $e_t(\sigma) \equiv e$.*

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, $q_t^s(\sigma) = \rho'$.*

2. *On each path σ at each date t and for all $\epsilon > 0$ there is a $\delta > 0$ such that if $|c_t^i(\sigma) - e| < \delta$, then $||q_t(\sigma) - \rho^i|| < \epsilon$.*

This is to say that prices reflect beliefs. 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, $S = \{A, B\}$. In a small abuse of our notation, take the probability of state A at any date t to be ρ. Traders have logarithmic utility, and have identical discount factors, $0 < \beta < 1$. Trader i 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 $w_t^i(\sigma)$ 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 α_t^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:

$$w_{t+1}^i(A) = \frac{\beta\rho^i w_t^i}{q_t^A}$$
$$w_{t+1}^i(B) = \frac{\beta(1-\rho^i)w_t^i}{q_t^B}$$

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,

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

and so the asset prices at date t are

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

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,

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

and so the ratio of is wealth share to js evolves as follows:

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

This evolution is more readily analyzed in its log form:

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

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

To understand how the market can learn, consider a Bayesian whose prior beliefs about state evolution contain I iid models in her support, $\{\rho^1, \ldots, \rho^I\}$, and let r_t^i denote the probability she assigns to model i posterior to the first t 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. In this sense, *the market is a Bayesian.* We are not committing economic anthropomorphism; we simply note the identity of the equation describing the evolution of posterior beliefs for a hypothetical Bayesian learner and that describing the evolution of the wealth share distribution.

From these observation we can draw several conclusions. If some trader holds correct beliefs, then in the long run his wealth share will converge 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, but are not correct, 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 payoff functions 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:

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

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 $I_\rho(\rho^i) < \infty$.

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 $\lim_t c_t^i(\sigma) = 0$. She survives on path σ if $\limsup_t c_t^i(\sigma) > 0$. A survivor i is negligible on path σ if for all $0 < r < 1$, $\lim_{T\to\infty}(1/T)|\{t \leq T : c_t^i(\sigma) > re_t(\sigma)\}| = 0$. 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 have that

$$\frac{u_i'\big(c_t^i(\sigma)\big)}{u_1'\big(c_t^1(\sigma)\big)} = \kappa_i \left(\frac{\beta_1}{\beta_i}\right)^t \prod_{s \in S} \left(\frac{\rho_s^1}{\rho_s^i}\right)^{n_t^s(\sigma)} \tag{5}$$

Note that if two traders i and j have identical beliefs and discount factors, then conditional on the total amount the two traders consume, the division between the two

traders is non-stochastic. If the total amount they consume is $c_t^{ij}(\sigma)$ on path σ at date t, then (c_t^i, c_t^j) solves the problem

$$
\begin{aligned}
u_{ij}\big(c_t^{ij}(\sigma)\big) = \max_{(c^i, c^j)} \quad & \frac{\lambda_i u_i(c^i) + \lambda_j u_j(c^j)}{\lambda_i + \lambda_j} \\
\text{such that} \quad & c^i + c^j - c_t^{ij}(\sigma^t) \leq \mathbf{0} \\
& \forall t, \sigma, \text{ and } c^i, c^j \geq 0
\end{aligned}
\tag{6}
$$

because the common discount factors and beliefs 'cancel out'. Thus if in problem (1) traders i and j with weights λ_i and λ_j are replaced by the collective trader with weight $\lambda_i + \lambda_j$ and payoff function u_{ij}, a consumption allocation solves the disaggregated problem only if the consumption allocation formed by summing i's and j's consumption and leaving everything else unchanged solves the aggregated problem. Conversely, if an allocation solves the aggregated problem, then the consumption allocation constructing by disaggregating the ij-trader's consumption at each date-event pair by solving problem (6), and leaving all other consumption unchanged, solves the disaggregated problem. In summary, there is no loss of generality in assuming that the traders are "representative traders", each representing a class of traders with identical discount factors and beliefs but perhaps different payoff functions. In paticular, without loss of generality we can assume that each "trader" has a unique discount factor-belief pair (β_i, ρ^i).

It will sometimes be convenient to have equation (5) in its log form:

$$
\log \frac{u_i'\big(c_t^i(\sigma)\big)}{u_1'\big(c_t^1(\sigma)\big)} = \log \kappa_i + t \log \frac{\beta_1}{\beta_i} - \sum_s n_t^s(\sigma) \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right).
$$

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.

$$
\begin{aligned}
\log \frac{u_i'\big(c_t^i(\sigma)\big)}{u_1'\big(c_t^1(\sigma)\big)} = \log \kappa_i &+ t \log \frac{\beta_1}{\beta_i} - t \sum_s \rho_s \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right) + \\
& \sum_s \big(n_t^s(\sigma) - t\rho_s \big) \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right) \\
= \log \kappa_i &+ t\big(\log \beta_1 - I_\rho(\rho^1) \big) - t\big(\log \beta_i - I_\rho(\rho^i) \big) - \\
& \sum_s \big(n_t^s(\sigma) - t\rho_s \big) \left(\log \frac{\rho_s^i}{\rho_s} - \log \frac{\rho_s^1}{\rho_s} \right)
\end{aligned}
$$

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

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

Then

$$
\log \frac{u_i'\big(c_t^i(\sigma)\big)}{u_1'\big(c_t^1(\sigma)\big)} = \log \kappa_i + t(\mathsf{s}_1 - \mathsf{s}_i) - \sum_s \big(n_t^s(\sigma) - t\rho_s\big)\left(\log \rho_s^i - \log \rho_s^1\right) \tag{7}
$$

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* $\mathsf{s}_i < \max_j \mathsf{s}_j$, *then trader i vanishes.*

The analysis compares one trader, say trader 1, to other traders in the economy. We use equation (7) 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 (7).

Lemma 1. *If on a sample path σ,* $\log u_i'\big(c_t^i(\sigma)\big)/u_1'\big(c_t^1(\sigma)\big) \to -\infty$ *for some trader i, then* $\lim_t c_t^1(\sigma) = 0$. *On the other hand, if* $\limsup \min_i \log u_i'\big(c_t^i(\sigma)\big)/u_1'\big(c_t^1(\sigma)\big) > -\infty$, *then* $\limsup_t c_t^1(\sigma) > 0$.

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: Either the denominator converges to 0 or the numerator diverges to infinity. It must be the latter, because the denominator is bounded below by $u_i'(F) > 0$. Consequently, on any such path, $c_t^i(\sigma) \to 0$.

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

Proof of Theorem 2. We prove this theorem by examining equation (7). 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,

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

If s_1 is not maximal, there is an i for which $\mathsf{s}_1 - \mathsf{s}_i < -\epsilon < 0$. For almost all paths σ^t there is a T such that if $t > T$, then $u'_i(c_t^i(\sigma))/u'_1(c_t^1(\sigma)) < -\epsilon t$. According to lemma 1, trader 1 vanishes. $\qquad \square$

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. Lemma 1 also gives a lower bound rate at which traders vanish. Let s^* be the maximal survival index in the trader population.

Corollary 1. *If s_1 is not maximal, then $u'_1(c_t^1(\sigma^t)) \exp(\mathsf{s}_1 - \mathsf{s}^*)t > 0$ a.s.*

Proof. Replace $u'_i(c_t^i(\sigma^t))$ with the smaller $u'_i(F)$ and calculate. $\qquad \square$

If trader 1 is of the CRRA class, for instance, it follows that her consumption declines at an exponential rate, which depends upon the coefficient of relative risk aversion. This shows that although whether a trader vanishes or not is independent of the payoff function, the *rate* at which she vanishes is payoff function-dependent. More risk-averse traders vanish at a slower rate, and in fact, in the rate there is a trade-off between the survival index and the coefficient of relative risk aversion γ, since c_t^i is $O(\exp t(\mathsf{s}_1 - \mathsf{s}^*)/\gamma)$.

3.5   Market Equilibrium — Selection

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

Corollary 2. *If there is a unique trader i with minimal survival index among the trader population, then market prices converge to ρ^i almost surely.*

The limit market price ρ^i is not necessarily the price representing the most accurate beliefs in the market due to the tradeoff in the survival index between belief accuracy and patience. The limit price will be the best beliefs in the market if all traders have identical discount factors.

Proof. 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. □

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 (7) 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 (7). 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 (7) 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 $\mathtt{s}$.

More definitions are required. For a probability distribution θ on S, define the vector of log-probabilities: $\mathrm{lo}(\theta) = \left(\log(\theta(s)/\theta(\mathtt{s}))\right)_{s=1}^{\mathtt{s}-1}$. Let $\mathtt{Sur}$ 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 $\mathtt{Sur}$ is determined by how her beliefs, as represented by $\mathrm{lo}(\rho^i)$, are positioned relative to the beliefs of the other traders in $\mathtt{Sur}$. Denote by $C\{\mathrm{lo}(\rho^j)\}_{j\in\mathtt{Sur}}$ the closed convex cone generated by the log-probability vectors of the traders.

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

We are interested in markets with heterogeneous potential survivors.

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

1. *If* $i \in \mathtt{Sur}$, *then trader i survives.*

2. *Extremal traders are non-negligible and interior traders are negligible*

3. *If trader i is extremal, then $\lim_{T \to \infty}(1/T)|\{t \le T : c_t^i > re_t\}| > 0$ a.s.*

When $\mathbf{s} \le 3$, a maximal survival index is sufficient (as well as necessary) for survival. But how one survives depends upon ones position in the group of survivors. 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 in 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. *Interior traders vanish.*

2. *Extremal traders survive and are non-negligable.*

3. *If trader i is extremal, then $\lim_{T \to \infty}(1/T)|\{t \le T : c_t^i > re_t\}| > 0$ a.s.*

The discussion of survival possibilities is complicated by the possibility of heterogeneous discount factors. If discount factors are homogeneous, all traders with maximal survival index survive.

Corollary 3. *If for all i and j, $\beta_i = \beta_j$ then i survives iff $i \in \mathtt{Sur}$.*

Proofs of Theorems 3 4 and Corollary 3. Without loss of generality we can take the welfare weights to be 1 (by multiplying each payoff function by an appropriate positive weight). We already know that if a trader is not in $\mathtt{Sur}$, she vanishes, so suppose that trader 1 is in $\mathtt{Sur}$, and consider trader $i \ne 1$ who is also in $\mathtt{Sur}$. (Nothing additional is learned by studying traders who are not in $\mathtt{Sur}$). Since 1 and i are both in $\mathtt{Sur}$, they have identical survival indices, and so the right hand side of equation (7) becomes

$$\hat{z}_t^i(\sigma) = \sum_s \left(n_t^s(\sigma) - tp_s\right) \log \frac{\rho_s^1}{\rho_s^i}$$

We want to investigate if, infinitely often, for all such traders $i \neq 1$ in $\mathtt{Sur}$, $\hat{z}^i$ is arbitrarily large, arbitrarily small, or both. To do this we study the evolution of the system

$$(z_t^i)_{i \in \mathtt{Sur}, i \neq 1} = A \begin{pmatrix} w_{1t} \\ \vdots \\ w_{\mathtt{s}-1,t} \end{pmatrix}$$

where $a_{is} = \log \rho_s^1 / \rho_{\mathtt{s}}^1 - \log \rho_s^i / \rho_{\mathtt{s}}^i$ for $s = 1, \ldots, \mathtt{s} - 1$ and $i \in \mathtt{Sur}, i \neq 1$, and w_t is the mean-0 random walk such that $w_{st}(\sigma) = n_t^s(\sigma) - t\rho_s$ for $s = 1, \ldots, \mathtt{s} - 1$. We can neglect the traders not in The random walk $(n_t^s(\sigma) - t\rho_s)_{s=1}^{\mathtt{s}}$ is an $\mathtt{s} - 1$ dimensional random walk since the sum of all elements of the vector is 0, so we have chosen a representation which drops the $\mathtt{s}$ coordinate. Thus the vector $z_t(\sigma) = (z_t(\sigma)_{i \in \mathtt{Sur}, i \neq 1})$ is a random walk in $\mathbf{R}^{\#\mathtt{Sur}-1}$. For ease of reference, define $\mathrm{lo}(\rho^i) = (\log \rho_s^i / \rho_{\mathtt{s}}^i)_{s=1}^{\mathtt{s}-1}$.

We consider outcomes for trader 1 corresponding to two types of A matrix. (1) There is a vector (direction) x such that $Ax \gg 0$. (2) for each vector $x \neq 0$ there is a row vector a_i of A such that $a_i \cdot x < 0$. The remaining possibility is that $Ax \geq 0$, but is not strictly greater than 0. As we will see, cases (1) and (2) correspond to extremal and interior traders. The remaining possibility is that a trader's beliefs are on the boundary of the polytope $C\{\mathrm{lo}(\rho^j)\}_{j \in \mathtt{Sur}}$ but are not extreme points. That is, they lie in the relative interior of some facet of the polytope. We call such traders *boundary traders*. We have no results for them, for reasons that will be discussed below.

If case (1) holds, then there is an open cone C such that for all $x \in C$, $Ax \gg 0$. Whenever $w_t \in C$, $\log u_i'(c_t^i(\sigma^t)) / u_1'(c_t^1(\sigma^t))$ is positive for all $i \in \mathtt{Sur}, i \neq 1$. Furthermore, for all consumption shares $r < 1$ there is a bound $b(r)$ such that if, for all $i \in \mathtt{Sur}, i \neq 1$, $a_i \cdot w_t(\sigma) > b(r)$, then $c_t^1(\sigma) > re_t(\sigma)$. The set of such w values is the open cone less a compact set containing the origin. Such sets are recurrent, so $\Pr\{c_t^1(\sigma) > re_t(\sigma) \text{ i.o.}\} = 1$.

If case (2) holds, then for any direction of the walk, there is a trader $i \in \mathtt{Sur}, i \neq 1$ 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 null-recurrent for two dimensional walks, and transient for walks in dimension three or higher. So when $\mathtt{s} \leq 3$ (so that the dimension of the walk does not exceed 2), trader 1 is negligible; otherwise trader 1 vanishes.

Cases 1 and 2, respectively, correspond to extremal and interior traders. The next lemma establishes this, which completes the proof.

Lemma 2. *The inequality system $Ax \gg 0$ has a solution iff $\mathrm{lo}(\rho^1)$ is an extreme point of the convex hull of the $\mathrm{lo}_i(\rho^i) \in \mathtt{Sur}$, that is, iff trader 1 is extremal. If trader 1 is interior, then for all directions x there is an a_i such that $a_i x < 0$.*

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 $\sum_{i \geq 2} \lambda_i \mathrm{lo}(\rho^i) = \mathrm{lo}(\rho^1)$. In other words, a strictly positive direction x exists iff $\mathrm{lo}(\rho^1)$ is extremal in $C\{\mathrm{lo}(\rho^j)\}_{j=1}^I$.

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. □

This proves both theorems. To prove the corollary, suppose that discount factors are identical. Then $s_i = -I_\rho(\rho^i)$ and for some $\mathtt{s}$ and $i \in \mathtt{Sur}$, $-I_\rho(\rho^i) = \mathtt{s}$. This function is strictlyconcave in ρ^i, so the set of all μ such that $-I_\rho(\nu) \geq \mathtt{s}$ is strictly convex, and all traders with maximal survival index are extreme points of this convex set. □

The following figures demonstrate the geometry of Theorem 3 when $\mathtt{s} = 3$. 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-odds ratios of five beliefs in $\mathtt{Sur}$ when $\mathtt{s} = 3$. (The discount factors for all traders cannot be all the same.) The upper left figure displays the polytope $C\{\mathrm{lo}(\rho^j)\}_{j \in \mathtt{Sur}}$. Traders A, B and C are extremal, E is interior and D is boundary. The figure indicates there are many directions of increase for extremal trader A, only 1 direction for boundary trader D, and no directions of increase for interior trader E.

Figure 1: Five Beliefs in `Sur`.

The three remaining figures indicate the directions of increase. The cone on the upper 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. Far enough out in the cone, trader A consumes an arbitrarily large share of aggregate endowment. On the bottom left, any direction in the shaded half-space improves D relative to B and E, but these directions mostly improve either A or C relative to B and E as well, and in fact improve one or the other of them more, so far enough out either A or C is consuming most of the endowment. When the random walk crosses the line orthogonal to the boundary of the half space, A, C and D have equal marginal utilities, and far out on this line (thus minimizing the consumption of B and E) is the best D can do. Thus D's share is bounded from above. Whether D vanishes or not depends upon the recurrence properties of lower dimensional cones, about which we have little to say at this level of generality. Finally on the right, any direction is better for some trader relative to E. The best E does is when the walk is near the origin. But when the walk is in a neighborhood of the origin, E's share is still bounded, and when $s > 3$ such neighborhoods are in any event transient.

4.2 Market Equilibrium — Balancing

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

Corollary 4. *If multiple traders have maximimal survival index, then for all extremal traders i and all $\epsilon > 0$, $|q_t - \rho^i| < \epsilon$ infinitely often. If $\mathtt{s} > 3$ 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 surivors 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 $\mathtt{Sur}$ 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 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 surviver, 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.

Figure 2: Multiple survivors, $\mathbf{s} = 3$.

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 (best) analysts do. 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 both 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 themselves 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.

The utility of the survival index is that it separates the stochastic evolution of marginal utility ratios into a deterministic drift and and a mean-0 stochastic component. One contribution of this paper is to show that the stochastic component matters for survival and long run asset prices. This analysis extends to Markov environments, where a survival index, the same kind of decomposition, also exists. In more complicated environments, such as that created by Bayesian traders, we already know that the simple decomposition into a deterministic drift and a stochastic residue in insufficient. We showed in Blume and Easley (2006) that when discount factors among Bayesian traders are identical, all survival indices converge to the log of the common discount factor, and that a necessary condition for survival is that the rate of convergence be fastest. The rate turns out to be related to the dimension of the support of prior beliefs; more parameters to learn slow down learning. We have not yet gone beyond this. Among traders who learn the fastest, whose parameter spaces are of minimal dimension, there are still differences due to different prior beliefs which appear in a stochastic term which is converging to 0. If the convergence is slow enough, there is still room for the kind of effects we describe here to work, but we have not even attempted this analysis. In still more complex environments, the analog of the survival index is a series of conditional means which are compared by divergence rates. We can construct example enviroments where exactly the effects we describe here work, and so the analysis of Blume and Easley (2006) is (again) necessary but not sufficient. These examples, however, are very artificial, and this level of analysis is probably too abstract to say anything useful.

Perhaps an even more compelling issue 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 smarter in the long run than its traders; and so we are led to ask, how is the market's learning experience different than that of its traders? The leading example of section 3.1 is a first step towards answering this question.

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:*

$$Ax \gg 0 \tag{8}$$

$$\begin{aligned} yA &= 0, \\ y &\geq 0, \quad y \neq 0 \end{aligned} \tag{9}$$

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 $f_i(x) < 0$, $i = 1, \ldots, m$, $x \in K$ 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 ith row of the matrix A. The f_i are convex functions and W is a convex set. If (8) has no solution, then according to Fan et. al., there are non-negative scalars y_i not all 0 such that for all $x \in W$,

$$\sum_i y_i(-a_i x) \geq 0. \tag{10}$$

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 $\sum_i y_i a_i = 0$. $\qquad\square$

Notes

[1] "The End of Management", Time Magazine Bonus Section, August 2004, http://www.time.com/time/insidebiz/article/0,9171,1101040712-660965-1,00.html. For a popular account of prediction markets, see Surowiecki (2003).

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

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.