

**Asset Pricing Under Endogenous Expectations  
in an Artificial Stock Market**

by

W. Brian Arthur, John H. Holland, Blake LeBaron, Richard Palmer, and Paul Tayler \*

Dec 12, 1996

\* All authors are affiliated with the Santa Fe Institute, where Arthur is Citibank Professor. In addition, Holland is Professor of Computer Science and Engineering, University of Michigan, Ann Arbor; LeBaron is Associate Professor of Economics, University of Wisconsin; Palmer is Professor of Physics, Duke University; and Tayler is with the Dept. of Computer Science, Brunel University, London.

## Asset Pricing Under Endogenous Expectations in an Artificial Stock Market

### Abstract

We propose a theory of asset pricing based on heterogeneous agents who continually adapt their expectations to the market that these expectations aggregatively create. And we explore the implications of this theory computationally using our Santa Fe artificial stock market.

Asset markets, we argue, have a recursive nature in that agents' expectations are formed on the basis of their anticipations of other agents' expectations, which precludes expectations being formed by deductive means. Instead traders continually hypothesize—continually explore—expectational models, buy or sell on the basis of those that perform best, and confirm or discard these according to their performance. Thus individual beliefs or expectations become endogenous to the market, and constantly compete within an ecology of others' beliefs or expectations. The ecology of beliefs co-evolves over time.

Computer experiments with this endogenous-expectations market explain one of the more striking puzzles in finance: that market traders often believe in such concepts as technical trading, "market psychology," and bandwagon effects, while academic theorists believe in market efficiency and a lack of speculative opportunities. Both views, we show, are correct, but within different regimes. Within a regime where investors explore alternative expectational models at a low rate, the market settles into the rational-expectations equilibrium of the efficient-market literature. Within a regime where the rate of exploration of alternative expectations is higher, the market self-organizes into a complex pattern. It acquires a rich psychology, technical trading emerges, temporary bubbles and crashes occur, and asset prices and trading volume show statistical features—in particular, GARCH behavior—characteristic of actual market data.

### Acknowledgments

We are grateful to Kenneth Arrow, Larry Blume, Buz Brock, John Casti, Steven Durlauf, David Easley, David Lane, Ramon Marimon, Tom Sargent, and Martin Shubik for discussions of the arguments in this paper, and of the design of the artificial market. All errors are our own.

# Asset Pricing Under Endogenous Expectations in an Artificial Stock Market

by

W. Brian Arthur, John H. Holland, Blake LeBaron, Richard Palmer, and Paul Tayler

## Introduction

Academic theorists and market traders tend to view financial markets in strikingly different ways. Standard (efficient-market) financial theory assumes identical investors who share rational expectations of an asset's future price, and who instantaneously and rationally discount all market information into this price.<sup>1</sup> It follows that no opportunities are left open for consistent speculative profit, that technical trading (using patterns in past prices to forecast future ones) cannot be profitable except by luck, that temporary price overreactions—bubbles and crashes—reflect rational changes in assets' valuations rather than sudden shifts in investor sentiment. It follows too that trading volume is low or zero, and that indices of trading volume and price volatility are not serially correlated in any way. The market, in this standard theoretical view, is rational, mechanistic, and efficient. Traders, by contrast, often see markets as offering speculative opportunities. Many believe that technical trading is profitable<sup>2</sup>, that something definable as a “market psychology” exists, and that herd effects unrelated to market news can cause bubbles and crashes. Some traders and financial writers even see the market itself as possessing its own moods and personality, sometimes describing the market as “nervous” or “sluggish” or “jittery.” The market in this view is psychological, organic, and imperfectly efficient. From the academic viewpoint traders with such beliefs—embarrassingly the very agents assumed rational by the theory—are irrational and superstitious. From the traders’ viewpoint, the standard academic theory is unrealistic and not borne out by their own perceptions.<sup>3</sup>

While few academics would be willing to assert that the market has a personality or experiences moods, the standard economic view has in recent years begun to change. The crash of 1987 damaged economists’ beliefs that sudden price changes reflect rational adjustments to news in the market: several studies failed to find significant correlation between the crash and market information issued at the time

---

<sup>1</sup> For the classic statement see Lucas (1978), or Diba and Grossman (1988).

<sup>2</sup> For evidence see Frankel and Froot (1990).

<sup>3</sup> To quote one of the most successful traders, George Soros (1994): “this [efficient market theory] interpretation of the way financial markets operate is severely distorted. ... It may seem strange that a patently false theory should gain such widespread acceptance.”

(e.g. Cutler *et al.* 1989). Trading volume and price volatility in real markets are large—not zero or small, respectively, as the standard theory would predict (Shiller, 1981, 1989; Leroy and Porter, 1981)—and both show significant autocorrelation (Bollerslev *et al.*, 1990; Goodhart and O’Hara, 1995). Stock returns also contain small, but significant serial correlations (Fama and French, 1988; Lo and Mackinlay, 1988; Summers, 1986; Poterba and Summers, 1988). Certain technical-trading rules produce statistically significant, if modest, long-run profits (Brock, Lakonishok, and LeBaron, 1991). And it has long been known that when investors apply full rationality to the market, they lack incentives both to trade and to gather information (Milgrom and Stokey, 1982; Grossman 1976; Grossman and Stiglitz, 1980). By now, enough statistical evidence has accumulated to question efficient-market theories and to show that the traders’ viewpoint cannot be entirely dismissed. As a result, the modern finance literature has been searching for alternative theories that can explain these market realities.

One promising modern alternative, the noise-trader approach, observes that when there are “noise traders” in the market—investors who possess expectations different from those of the rational-expectations traders—technical-trading strategies such as trend chasing may become rational. For example, if noise traders believe that an upswing in a stock’s price will persist, rational traders can exploit this by buying into the uptrend thereby exacerbating the trend. In this way positive-feedback trading strategies—and other technical-trading strategies—can be seen as rational, as long as there are non-rational traders in the market to prime these strategies (De Long *et al.* 1990a, 1990b, 1991; Shleifer and Summers, 1990). This “behavioral” noise-trader literature moves some way toward justifying the traders’ view. But it is built on two less-than-realistic assumptions: the existence of unintelligent noise traders who do not learn over time their forecasts are erroneous; and of rational players who possess, by some unspecified means, full knowledge of both the noise traders’ expectations and their own class’s. Neither assumption is likely to hold up in real markets. Suppose for a moment an actual market with minimally intelligent noise traders. Over time, in all likelihood, some would discover their errors and begin to formulate more intelligent (or at least different) expectations. This would change the market, which means that the perfectly intelligent players would need to readjust *their* expectations. But there is no reason these latter would know the new expectations of the noise-trader deviants; they would have to derive their expectations by some means such as guessing or observation of the market. As the rational players changed, the market would change again. And so the noise traders might again further deviate, forcing further readjustments for the rational traders. Actual noise-trader markets, assumed stationary in theory, would start to unravel; and the perfectly rational traders would be left at each turn guessing the changed expectations by observing the market.

Thus noise-trader theories, while they explain much, are not robust. But in questioning such theories we are led to an interesting sequence of thought. Suppose we were to assume “rational,” but non-identical, agents who do not find themselves in a market with rational expectations, or with publicly-known expectations. Suppose we allowed each agent continually to observe the market with an eye to discovering

profitable expectations. Suppose further we allowed each agent to adopt these when discovered and to discard the less profitable as time progressed. In this situation, agents' expectations would become endogenous—individually adapted to the current state of the market—and they would co-create the market they were designed to exploit. How would such a market work? How would it act to price assets? Would it converge to a rational-expectations equilibrium—or would it uphold the traders' viewpoint?

In this paper we propose a theory of asset pricing that assumes fully heterogeneous agents whose expectations continually adapt to the market these expectations aggregatively create. We argue that under heterogeneity, expectations have a recursive character: agents have to form their expectations from their anticipations of other agents' expectations, and this self-reference precludes expectations being formed by deductive means. So, in the absence of being able to deduce expectations, agents—no matter how rational—are forced to hypothesize them. Agents therefore continually form individual, hypothetical, expectational models or “theories of the market,” test these, and trade on the ones that predict best. From time to time they drop hypotheses that perform badly, and introduce new ones to test. Prices are driven endogenously by these induced expectations. Individuals' expectations therefore evolve and “compete” in a market formed by others' expectations. In other words, agents' expectations co-evolve in a world they co-create.

The natural question is whether these heterogeneous expectations co-evolve into homogeneous rational-expectations beliefs, upholding the efficient-market theory, or whether richer individual and collective behavior emerges, upholding the traders' viewpoint and explaining the empirical market phenomena mentioned above. We answer this not analytically—our model with its fully heterogeneous expectations it is too complicated to admit of analytical solutions—but computationally. To investigate price dynamics, investment strategies, and market statistics in our endogenous-expectations market, we perform carefully-controlled experiments within a computer-based market we have constructed, the SFI Artificial Stock Market.<sup>4</sup>

The picture of the market that results from our experiments, surprisingly, confirms both the efficient-market academic view and the traders' view. But each is valid under different circumstances—in different regimes. In both circumstances, we initiate our traders with heterogeneous beliefs clustered randomly in an interval near homogeneous rational expectations. We find that if our agents adapt their forecasts very slowly to new observations of the market's behavior, the market converges to a rational-expectations regimes. Here “mutant” expectations cannot get a profitable footing; and technical trading, bubbles, crashes, and autocorrelative behavior do not emerge. Trading volume remains low. The efficient-market theory prevails.

---

<sup>4</sup> For an earlier report on the SFI artificial stock market, see Palmer *et al.* (1994).

If, on the other hand, we allow the traders to adapt to new market observations at a more realistic rate, heterogeneous beliefs persist, and the market self-organizes into a complex regime. A rich “market psychology”—a rich set of expectations—becomes observable. Technical trading emerges as a profitable activity, and temporary bubbles and crashes occur from time to time. Trading volume is high, with times of quiescence alternating with times of intense market activity. The price time series shows persistence in volatility, the characteristic GARCH signature of price series from actual financial markets. And it shows persistence in trading volume. And over the period of our experiments, at least, individual behavior evolves continually and does not settle down. In this regime, the traders’ view is upheld.

In what follows, we discuss first the rationale for our endogenous-expectations approach to market behavior; and introduce the idea of collections of conditional expectational hypotheses or “predictors” to implement this. We next set up the computational model that will form the basic framework. We are then in a position to carry out and describe the computer experiments with the model. Two final sections discuss the results of the experiments, compare our findings with other modern approaches in the literature, and summarize our conclusions.

## 2. Why Inductive Reasoning?

Before proceeding, we show that once we introduce heterogeneity of agents, deductive reasoning on the part of agents fails. We argue that in the absence of deductive reasoning, agents must resort to *inductive* reasoning, which is both natural and realistic in financial markets.

### A. Forming Expectations by Deductive Reasoning: an Indeterminacy

We make our point about the indeterminacy of deductive logic on the part of agents using a simple arbitrage pricing model, avoiding technical details that will be spelled out later. (This pricing model is a special case of our model in Section 3, assuming risk coefficient  $\lambda$  arbitrarily close to 0, and gaussian expectational distributions.) Consider a market with a single security that provides a stochastic payoff or dividend sequence  $\{d_t\}$ , with a risk-free outside asset that pays a constant  $r$  units per period. Each agent  $i$  may form individual expectations of next period’s dividend and price,  $E_i[d_{t+1}|I_t]$  and  $E_i[p_{t+1}|I_t]$ , with conditional variance of these combined expectations,  $\sigma_{i,t}^2$ , given current market information  $I_t$ . Assuming perfect arbitrage, the market for the asset clears at the equilibrium price:

$$p_t = \beta \sum_j w_{j,t} (E_j[d_{t+1}|I_t] + E_j[p_{t+1}|I_t]) \quad (1)$$

In other words, the security's price  $p_t$  is bid to a value that reflects the current (weighted) average of individuals' market expectations, discounted by the factor  $\beta = 1/(1+r)$ , with weights  $w_{j,t} = (1/\sigma_{j,t}^2) / \sum_k 1/\sigma_{k,t}^2$ , the relative "confidence" placed in agent  $j$ 's forecast.

Now, assuming intelligent investors, the key question is how the individual dividend and price expectations  $E_i[d_{t+1}|I_t]$  and  $E_i[p_{t+1}|I_t]$  might be formed. The standard argument that such expectations can be formed rationally (*i.e.*, using deductive logic) goes as follows. Assume *homogeneous* investors who (i) use the available information  $I_t$  identically in forming their dividend expectations, and (ii) know that others use the same expectations. Assume further that the agents (iii) are perfectly rational (can make arbitrarily difficult logical inferences), (iv) know that price each time will be formed by arbitrage as in (1), and (v) that (iii) and (iv) are common knowledge. Then expectations of future dividends  $E_i[d_{t+k}|I_t]$  are by definition known, shared, and identical. And homogeneity allows us to drop the agent subscript and set the weights to  $1/N$ . It is then a standard exercise (see Diba and Grossman, 1988) to show that by setting up the arbitrage equation (1) for future times  $t+k$ , taking expectations across it, and substituting backward repeatedly for  $E[p_{t+k}|I_t]$ , agents can iteratively solve for the current price as<sup>5</sup>

$$p_t = \beta^k \sum_{k=1}^{\infty} E[d_{t+k}|I_t]. \quad (2)$$

If the dividend expectations are unbiased, dividend forecasts will be upheld on average by the market and so the price sequence will be in rational-expectations equilibrium. Thus the price fluctuates as the information  $\{I_t\}$  fluctuates over time, and it reflects "correct" or "fundamental" value, so that speculative profits are not consistently available. Of course, rational-expectations models in the literature are typically more elaborate than this. But the point so far is that if we are willing to adopt the above assumptions—which depend heavily on homogeneity—asset pricing becomes deductively determinate, in the sense that agents can, in principle at least, logically derive the current price.

Assume now more realistically that traders are intelligent but heterogeneous—each may differ from the others. Now, the available shared information  $I_t$  consists of past prices, past dividends, trading volumes, economic indicators, rumors, news, and the like. These are merely qualitative information plus data sequences, and *there may be many different, perfectly defensible statistical ways*, based on different assumptions and different error criteria to use them to predict future dividends (Arthur, 1992; Kurz, 1993). Thus there is no objectively laid-down expectational model that differing agents can coordinate upon, and so there is no objective means for one agent to know other agents' expectations of future dividends. This is

---

<sup>5</sup> The second, constant-exponential-growth solution is normally ruled out by an appropriate transversality condition.

sufficient to bring indeterminacy to the asset price in (1). But worse, the heterogeneous price expectations  $E_i[p_{t+1}|I_t]$  are also indeterminate. For suppose agent  $i$  attempts rationally to deduce this expectation, he may take expectations across the market clearing equation (1) for time  $t+1$ :

$$E_i[p_{t+1}|I_t] = \beta E_i \left[ \sum_j \left\{ w_{j,t+1} (E_j[d_{t+2}|I_t] + E_j[p_{t+2}|I_t]) \right\} | I_t \right] \quad (3)$$

This requires that agent  $i$ , in forming his expectation of price, take into account his expectations of *others*' expectations of dividends and price (and relative market weights) two periods hence. To eliminate in like manner the price expectation  $E_j[p_{t+2}|I_t]$  requires a further iteration. But this leads agents into taking into account their expectations of others' expectations of others' expectations of future dividends and prices at period  $t+3$ —literally, as in Keynes's (1936) phrase, taking into account “what average opinion expects the average opinion to be.”

Now, under homogeneity these expectations of others' expectations collapsed into single, shared, objectively determined expectations. Under heterogeneity, however, not only is there no objective means by which others' dividend expectations can be known, but attempts to eliminate the other unknowns, the price expectations, merely lead to the repeated iteration of subjective expectations of subjective expectations (or equivalently, subjective priors on others' subjective priors)—an infinite regress in subjectivity. Further, this regress may lead to instability: If investor  $i$  believes that others believe future prices will increase, he may revise his expectations to expect upward-moving prices. If he believes that others believe a reversion to lower values is likely, he may revise his expectations to expect a reversion. We can therefore easily imagine swings and swift transitions in investors' beliefs, based on little more than ephemera—hints and perceived hints of others' beliefs about others' beliefs.

Under heterogeneity then, deductive logic leads to expectations that are not determinable. Notice the argument here depends in no way on agents having limits to their reasoning powers. It merely says that given differences in agent expectations, there is no logical means by which to arrive at expectations. And so, perfect rationality in the market can not be well-defined. Infinitely intelligent agents cannot form expectations in a determinate way.

### *B. Forming Expectations by Inductive Reasoning*

If heterogeneous agents cannot deduce their expectations, how then do they form expectations? They may observe market data, they may contemplate the nature of the market and of their fellow investors. They may derive expectational models by sophisticated, subjective reasoning. But in the end all such models will

be—can only be—hypotheses. There is no objective way to verify them, except by observing their performance in practice. Thus agents, in facing the problem of choosing appropriate predictive models, face the same problem that statisticians face when choosing appropriate predictive models given a specific data set, but no objective means by which to choose a functional form. (Of course, the situation here is made more difficult by the fact that the expectational models investors choose affect the price sequence, so that our statisticians' very choices of model affect their data and so their choices of model.)

In what follows then, we assume that each agent acts as a market “statistician.”<sup>6</sup> Each continually creates multiple “market hypotheses”—subjective, expectational models—of what moves the market price and dividend. And each simultaneously tests several such models. Some of these will perform well in predicting market movements. These will gain the agent’s confidence and be retained and acted upon in buying and selling decisions. Others will perform badly. They will be dropped. Still others will be generated from time to time and tested for accuracy in the market. As it becomes clear which expectational models predict well, and as poorly predicting ones are replaced by better ones, the agent learns and adapts. This type of behavior—coming up with appropriate hypothetical models to act upon, strengthening confidence in those that are validated, and discarding those that are not—is called *inductive reasoning*.<sup>7</sup> It makes excellent sense where problems are ill-defined. It is, in micro-scale, the scientific method. Agents that act use inductive reasoning we will call inductively rational.<sup>8</sup>

Each inductively-rational agent generates multiple expectational models that “compete” for use within his or her mind, and survive or are changed on the basis of their predictive ability. The agents’ hypotheses and expectations adapt to the current pattern of prices and dividends; and the pattern of prices changes to reflect the current hypotheses and expectations of the agents. We see immediately that the market possesses a *psychology*. We define this as the collection of market hypotheses or (expectational models or mental beliefs) that are being acted upon at a given time.

If there were some attractor inherent in the price-and-expectation-formation process, this market psychology might converge to a stable unchanging set of heterogeneous (or homogeneous) beliefs. Such a set would be statistically validated, and would therefore constitute a rational-expectations equilibrium. We investigate whether the market converges to such an equilibrium below.

---

<sup>6</sup> The phrase is Tom Sargent’s (1993). Sargent argues similarly, within a macroeconomic context, that to form expectations agents need to act as market statisticians.

<sup>7</sup> For earlier versions of induction applied to asset pricing and to decision problems, see Arthur (1992) and (1994: the *El Farol* problem), and Sargent, *op. cit.* For accounts of inductive reasoning in the psychological and adaptation literature, see Holland *et al.* (1986), Rumelhart (1980), and Schank and Abelson (1977).

<sup>8</sup> In the sense that they use available market data to learn—and switch among—appropriate expectational models. Perfect inductive rationality, of course, is indeterminate. Learning agents can be arbitrarily intelligent, but without knowing other’s learning methods cannot tell a-priori that *their* learning methods are maximally efficient. They can only discover the efficacy of their methods by testing them against data.

### 3. A Market with Induced Expectations

#### A. The Model

We now set up a simple model of an asset market along the lines of Bray (1982) or Grossman and Stiglitz (1980). The model will be neoclassical in structure, but will depart from standard models by assuming heterogeneous agents who form their expectations inductively by the process outlined above.

Consider a market in which  $N$  heterogeneous agents decide on their desired asset composition between a risky stock paying a stochastic dividend, and a risk-free bond. These agents formulate their expectations separately, but are identical in other respects. They possess a constant absolute risk aversion (CARA) utility function,  $U(c) = -\exp(-\lambda c)$ . They communicate neither their expectations nor their buying or selling intentions to each other. Time is discrete and is indexed by  $t$ ; the horizon is indefinite. The risk-free bond is in infinite supply and pays a constant interest rate  $r$ . The stock is issued in  $N$  units, and pays a dividend,  $d_t$ , which follows a given exogenous stochastic process  $\{d_t\}$  not known to the agents.

The dividend process, thus far, is arbitrary. In the experiments we carry out below, we specialize it to an AR(1) process

$$d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + \varepsilon_t, \quad (4)$$

where  $\varepsilon_t$  is gaussian, i. i. d., and has zero mean, and variance  $\sigma_e^2$ .

Each agent attempts at each period to optimize his allocation between the risk-free asset and the stock. Assume for the moment that agent  $i$ 's predictions at time  $t$  of the next period's price and dividend are normally distributed with (conditional) mean and variance,  $E_{i,t}[p_{t+1} + d_{t+1}]$ , and  $\sigma_{i,t,p+d}^2$ . (We say presently how such expectations are arrived at.) It is well known that under CARA utility and gaussian distributions for forecasts, agent  $i$ 's demand,  $x_{i,t}$ , for holding shares of the risky asset is given by:

$$x_{i,t} = \frac{E_{i,t}(p_{t+1} + d_{t+1}) - p(1+r)}{\lambda \sigma_{i,t,p+d}^2} \quad (5)$$

where  $p_t$  is the price of the risky asset at  $t$ , and  $\lambda$  is the degree of relative risk aversion.

Total demand must equal the number of shares issued:

$$\sum_{i=1}^N x_{i,t} = N, \quad (6)$$

which closes the model and determines the clearing price  $p$ —the current market price—in (5) above.

It is useful to be clear on timing in the market. At the start of time period  $t$ , the current dividend  $d_t$  is posted, and observed by all agents. Agents then use this information and general information on the state of the market (which includes the historical dividend sequence  $\{ \dots d_{t-2}, d_{t-1}, d_t \}$  and price sequence  $\{ \dots p_{t-2}, p_{t-1} \}$ ) to form their expectations of the next period's price and dividend  $E_{i,t}(p_{t+1} + d_{t+1})$ . They then calculate their desired holdings and pass their demand parameters to the specialist who declares a price  $p_t$  that clears the market. At the start of the next period the new dividend  $d_{t+1}$  is revealed, and the accuracies of the predictors active at time  $t$  are updated. The sequence repeats.

### B. Modeling the Formation of Expectations

At this point we have a simple, neoclassical, two-asset market. We now break from tradition by allowing our agents to form their expectations individually and inductively. One obvious way to do this would be to posit a set of individual-agent expectational models which share the same functional form, and whose parameters are updated differently by each agent (by least squares say) over time, starting from different priors. We reject this in favor of a different approach that better reflects the process of induction outlined in Section 2 above. We assume each agent at any time possesses a multiplicity of linear forecasting models—hypotheses about the direction of the market, or “theories of the market”—and uses those that are both best suited to the current state of the market and have recently proved most reliable. Agents then learn, not by updating parameters, but by discovering which of their hypotheses “prove out” best, and by developing new ones from time to time, via the genetic algorithm. This structure will offer several desirable properties: It will avoid biases introduced by a fixed, shared functional form. It will allow the individuality of expectations to emerge over time (rather than be built in only to a-priori beliefs). And it will better mirror actual cognitive reasoning, in which different agents might well “cognize” different patterns and arrive at different forecasts from the same market data.

In the expectational part of the model, at each period, the time series of current and past prices and dividends are summarized by an array or information set of  $J$  market descriptors. And agents' subjective expectational models are represented by sets of *predictors*. Each predictor is a *condition/forecast* rule (similar to a Holland classifier which is a condition/action rule) that contains both a market condition that may at times be fulfilled by the current state of the market and a forecasting formula for next period's price and dividend. Each agent possesses  $M$  such individual predictors—holds  $M$  hypotheses of the market in mind simultaneously—and uses the most accurate of those that are *active* (matched by the current state of

the market). In this way, each agent therefore has the ability to “recognize” different sets of states of the market, and bring to bear appropriate forecasts, given these market patterns.

It may clarify matters to show briefly how we implement this expectational system on the computer. (Further details are in Appendix A.) Suppose we summarize the state of the market by  $J=13$  bits. The fifth bit might correspond to “the price has risen the last 3 periods,” and the tenth bit to “the price is larger than 16 times dividend divided by  $r$ ,” with 1 signaling the occurrence of the described state, and 0 its absence or non-occurrence. Now, the condition part of all predictors corresponds to these market descriptors, and thus also consists of a 13-bit array, each position of which is filled with a 0, or 1, or # (“don’t care”). A condition array matches or “recognizes” the current market state if all its 0’s and 1’s match the corresponding bits for the market state with the #’s matching either a 1 or a 0. Thus the condition (#####1#####0####) “recognizes” market states in which the price has risen in the last 3 periods. The condition (#####0#####0###) recognizes states where the current price is not larger than 16 times dividend divided by  $r$ . The forecasting part of each predictor is an array of parameters that triggers a corresponding forecasting expression. In our experiments, all forecasts use a linear combination of price and dividend,  $E(p_{t+1} + d_{t+1}) = a(p_t + d_t) + b$ . Each predictor then stores specific values of  $a$  and  $b$ . Therefore the full predictor (#####1#####0###) / (0.96, 0) can be interpreted as “if the price has risen in the last 3 periods and if the price is not larger than 16 times dividend divided by  $r$  then forecast next period’s price plus dividend as 96% of this period’s.” This predictor would recognize—would be activated by—the market state (0110100100011) but would not respond to the state (0110111011001).

Predictors that can recognize many states of the market have few 1’s and 0’s. Those more particularized have more 1’s and 0’s. In practice we include for each agent a default predictor consisting of all #’s. The genetic algorithm creates new predictors by “mutating” the values in the predictor array, or by “recombination”—combining part of one predictor array with the complementary part of another.

The expectational system then works at each time with each agent observing the current state of the market, and noticing which of his predictors match this state. He forecasts next period’s price and dividend by combining statistically the linear forecast of the  $H$  most accurate of these active predictors, and given this expectation and its variance, uses (5) to calculate desired stock holdings and generate an appropriate bid or offer. Once the market clears, the next period’s price and dividend are revealed and the accuracies of the active predictors are updated.

As noted above, learning in this expectational system takes place in two ways. It happens rapidly as agents learn which of their predictors are accurate and worth acting upon, and which should be ignored. And it happens on a slower time scale as the genetic algorithm from time to time discards non-performing predictors and creates new ones. Of course these new, untested predictors do not create disruptions—they will be acted upon only if they prove accurate. This avoids brittleness and provides what machine-learning theorists call “gracefulness” in the learning process.

We can now discern several advantages of this multi-bit, multi-predictor architecture. One is that this expectational architecture allows the market to have potentially different dynamics—a different character—under different states or circumstances. Because predictors are pattern-recognizing expectational models, and so can “recognize” these different states, agents can “remember” what happened before in given states and activate appropriate forecasts. This enables agents to make swift *gestalt*-like transitions in forecasting behavior should the market change.

Second, the design avoids bias from the choice of a particular functional form for expectations. Although the forecasting part of our predictors is linear, the multiplicity of predictors conditioned upon the many combinations of market conditions yield collectively at any time and for any agent a nonlinear forecasting expression in the form of a piece-wise linear, non-continuous forecasting function whose domain is the market state space, and whose accuracy is tuned to different regions of this space. (Forecasting is of course limited by the choice of the binary descriptors that represent market conditions.)

Third, learning is concentrated where it is needed. For example,  $J=12$  descriptors produces predictors that can distinguish over four thousand different states of the market. Yet, only a handful of these states might occur often. Predictor conditions that recognize states that do not occur often will be used less often, their accuracy will be updated less often and, other things equal, their precision will be lower. They are therefore less likely to survive in the competition among predictors. Predictors will therefore cluster in the more visited parts of the market state space, which is exactly what we want.

Finally, the descriptor bits can be organized into classes or information sets which summarize fundamentals, such as price-dividend ratios or technical-trading indicators, such as price trend movements. The design allows us to track exactly which information—which descriptor bits—the agents are using and ignoring, something of crucial importance if we want to test for the “emergence” of technical trading. This organization of the information also allows the possibility setting up different agent “types” who have access to different information sets. (In this paper, all agents see all market information equally.)

A neural net could also supply several of these desirable qualities. However, it would be less transparent that our predictor system, which we can easily monitor to observe which information agents are individually and collectively using at each time.

## 5. Computer Experiments: The Emergence of Two Market Regimes

### A. Experimental Design

We now explore computationally the behavior of our endogenous-expectations market in a series of experiments. We retain the same model parameters throughout these experiments, so that we can make comparisons of the market outcomes using the model under identical conditions with only controlled

changes. Each experiment is run for 250,000 periods to allow asymptotic behavior to emerge if it is present; and it is run 25 times under different random seeds to collect cross-sectional statistics.

We specialize the model described in the previous section by choosing parameter values, and, where necessary, functional forms. We use  $N = 25$  agents, who each have  $M = 100$  predictors, which are conditioned on  $J = 12$  market descriptors. The dividend follows the AR(1) process in (4), with autoregressive parameter  $\rho$  set to 0.95, yielding a process close to a random walk, yet persistent.

The 12 binary descriptors that summarize the state of the market are the following:

**1-6** Current price  $\times$  interest rate/dividend  $> 0.25, 0.5, 0.75, 0.875, 1.0, 1.125$

**7-10** Current price  $>$  5-period moving average of past prices (MA), 10-period MA, 100-period MA, 500-period MA

**11** Always on (1)

**12** Always off (0)

The first six binary descriptors—the first six bits—reflect the current price in relation to current dividend, and thus indicate whether the stock is above or below fundamental value at the current price. We will call these “fundamental” bits. Bits 7-10 are “technical trading” bits which indicate whether a trend in the price is under way. They will be ignored if useless, and acted upon if technical-analysis trend following emerges. The final two bits, constrained to be 0 or 1 at all times, serve as experimental controls. They convey no useful market information, but can tell us the degree to which agents act upon useless information at any time. We say a bit is “set” if it is 0 or 1, and predictors are selected randomly for recombination, other things equal, with slightly lower probabilities the higher their specificity—that is, the more they contain bits set (see Appendix A). This introduces a weak drift toward the all-# configuration, and ensures that the information represented by a particular bit is used only if agents find it genuinely useful in prediction. This market information design allows us to speak of “emergence.” For example, it can be said that technical trading has emerged if bits 7 to 10 become set significantly more often, statistically, than the control bits.

We assume that forecasts are formed by each predictor  $j$  storing values for the parameters  $a_j, b_j$  in the linear combination of price and dividend,  $E_j[p_{t+1} + d_{t+1}|I_t] = a_j(p_t + d_t) + b_j$ . Each predictor also stores a current estimate of its forecast variance. (See Appendix A.)

Before we conduct experiments, we run two diagnostic tests on our computer-based version of the model. In the first, we test to see whether the model can replicate the rational expectations equilibrium of standard theory. We do this by calculating analytically the homogeneous rational expectations equilibrium (h.r.e.e.) values for the forecasting parameters  $a$  and  $b$ , (see Appendix A), then running the computation with all predictors “clamped” to these calculated h.r.e.e. parameters. We find indeed that such predictions are upheld—that the model indeed reproduces the homogeneous rational expectations equilibrium—which

assures us that the computerized model, with its expectations, demand functions, aggregation, market clearing, and timing sequence, is working correctly. In the second test, we show the agents a given dividend sequence and a calculated h.r.e.e. price series that corresponds to it, and test whether they individually learn the correct forecasting parameters. They do, though with some variation due to the agents' continual exploration of expectational space, which assures us that our agents are learning properly.

### B. The Experiments

We now run two sets of fundamental experiments with the computerized model, corresponding respectively to slow and medium rates of exploration by agents of alternative expectations. The two sets give rise to two different *regimes*—two different sets of characteristic behaviors of the market. In the slow-learning-rate experiments, the genetic algorithm is invoked every 1,000 periods on average, predictors are crossed over with probability 0.3, and the predictors' accuracy-updating parameter  $\theta$  is set to 1/150. In the medium-exploration-rate experiments, the genetic algorithm is invoked every 250 periods on average, crossover occurs with probability 0.1, and the predictors' accuracy-updating parameter  $\theta$  is set to 1/75.<sup>9</sup> Otherwise, we keep the model parameters the same in both sets of experiments, and in both we start the agents with expectational parameters selected randomly from a uniform distribution of values centered upon the calculated homogeneous rational expectations ones.(see Appendix A). In the slow-exploration-rate experiments, no non-r.e.e. expectations can get a footing: the market enters an evolutionarily-stable, rational-expectations regime. In the medium-exploration-rate experiments, we find that the market enters a complex regime in which psychological behavior emerges, there are significant deviations from the r.e.e. benchmark, and statistical “signatures” of real financial markets are observed.

We now describe these two sets of experiments and the two regimes or phases of the market they induce.

**The Rational Expectations Regime.** As stated, in this set of experiments, agents continually explore in prediction space, but under low rates. The market price, in these experiments, converges rapidly to homogeneous rational expectations value adjusted for risk, even though the agents start with non rational expectations. In other words, homogeneous rational expectations are an attractor for a market with endogenous, inductive expectations.<sup>10</sup> This is not surprising. If some agents forecast differently than the h.r.e.e. value, then the fact that most other agents are using something close to the h.r.e.e. value, will return a market-clearing price that corrects these deviant expectations: There is a natural, if weak, attraction to

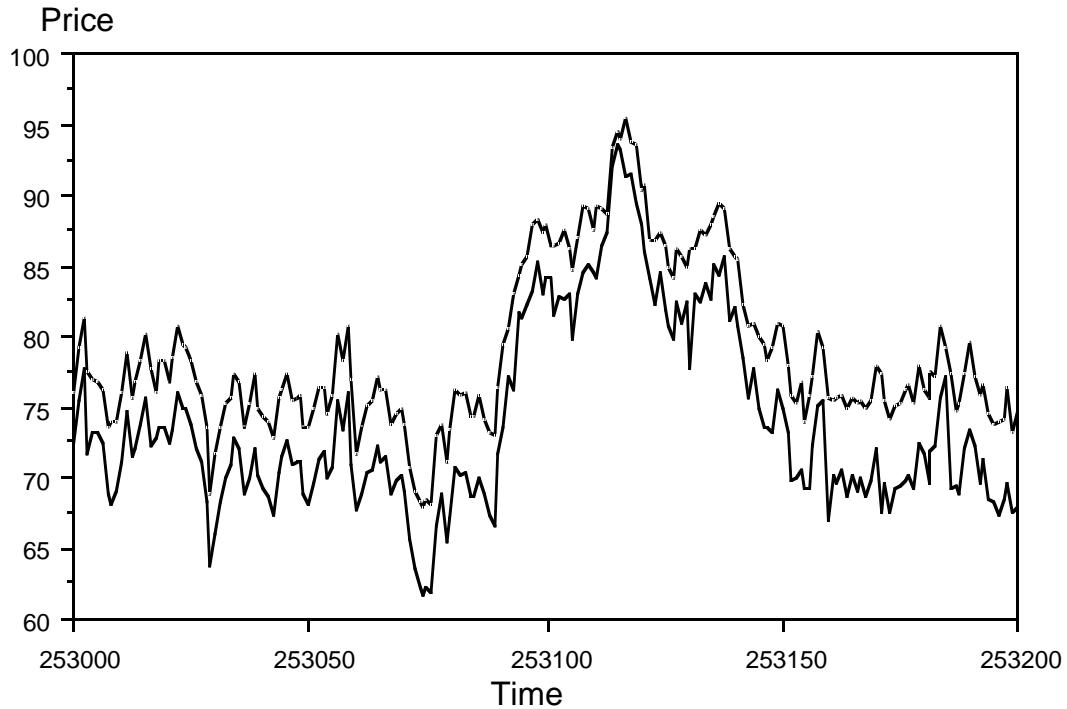
---

<sup>9</sup> At the time of writing, we have discovered that the two regimes emerge, and the results are materially the same, if we vary *only* the rate of invocation of the genetic algorithm.

<sup>10</sup> Within a simpler model, Blume and Easley (1982) prove analytically the evolutionary stability of r.e.e.

h.r.e.e. The equilibrium within this regime differs in two ways from the standard, theoretical, rational-expectations equilibrium. First, the equilibrium is neither assumed nor arrived at by deductive means. Our agents instead arrive inductively at a homogeneity that overlaps that of the homogeneous, theoretical rational expectations. Second, the equilibrium is a stochastic one. Agents continually explore alternatives, albeit at low rates. This testing of alternative explorations, small as it is, induces some “thermal noise” into the system. As we would expect, in this regime, agents’ holdings remain highly homogeneous, trading volume remains low (reflecting only variations in forecasts due to mutation and recombination) and bubbles, crashes, and technical trading do not emerge. We can say that in this regime the efficient-market theory and its implications are upheld.

**The Complex or Rich Psychological Regime.** We now allow a more realistic level of exploration in belief space. In these experiments, as we see in Fig. 1, the price series still appears to be nearly identical to the price in the rational-expectations regime. (It is lower because of risk attributable to the higher variance caused by increased exploration.)

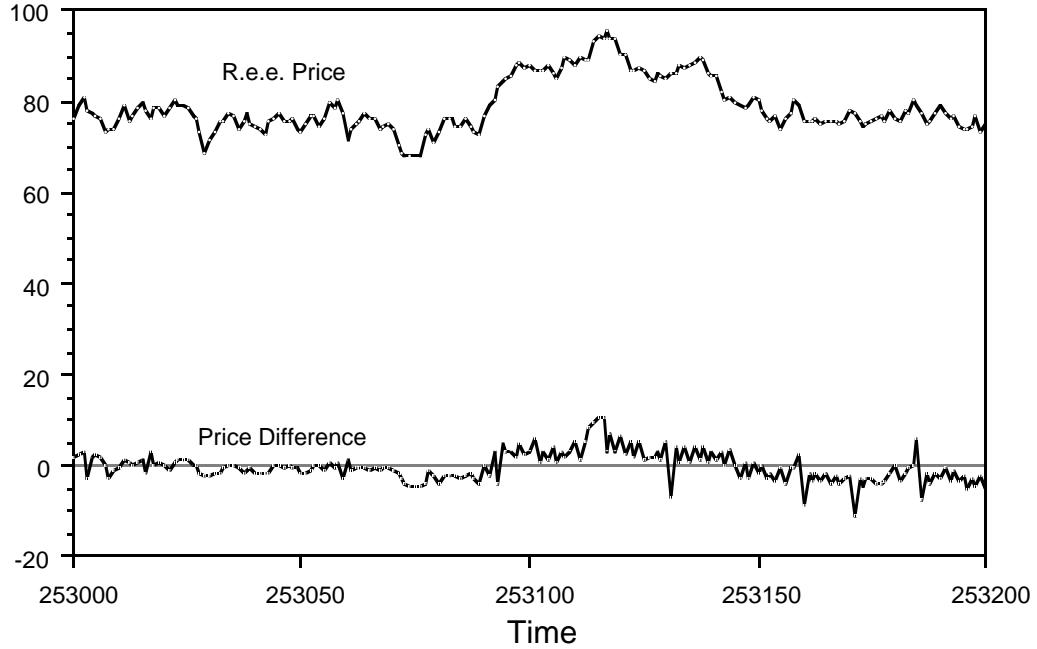


Figs 1a and 1b. Rational-expectations price vs. price in the rich psychological regime.

The two price series are generated on the same random dividend series. The upper is the homogeneous r.e.e. price, the lower is the price in the complex regime. The higher variance in the latter case causes the lower price through risk aversion.

On closer inspection of the results, however, we find that complex patterns have formed in the collection of

beliefs, and that the market displays characteristics that differ materially from those in the rational expectations regime. For example, when we magnify the difference between the two price series, we see systematic evidence of temporary price bubbles and crashes (Fig. 2). We call this new set of market behaviors the rich-psychological, or complex, regime.



Figs 2. Deviations of the price series in the complex regime from fundamental value.

The bottom graph shows the difference between the two price series in Fig 1. (with the complex series rescaled to match the r.e.e. one and the difference between the two doubled for ease of observation). The upper series is the h.r.e.e. price

This appearance of bubbles and crashes suggests that technical trading, in the form of buying or selling into trends, has emerged in the market. We can check this rigorously by examining the information the agents condition their forecasts upon. Figure 3 shows the number of technical trading bits that are used (are 1's or 0's) in the population of predictors as it evolves over time. In both sets of experiments, technical trading bits are initially seeded randomly in the predictor population. In the rational-expectations regime, however, technical trading bits provide no useful information and fall off as useless predictors are discarded. But in the complex regime, they bootstrap in the population, reaching a steady-state value by 150,000 periods. Technical trading, once it emerges, remains.<sup>11</sup>

---

<sup>11</sup> When we run these experiments informally to 1,000,000 periods, we see no signs that technical bits disappear.

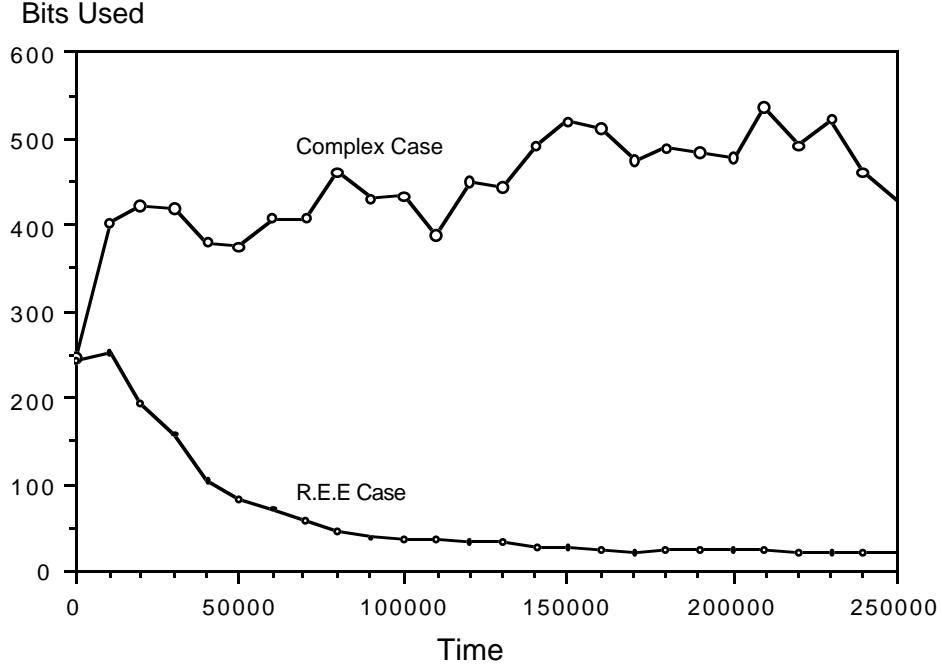


Figure 3. Number of technical-trading bits that become set as the market evolves, (median over 25 experiments in the two regimes).

Price statistics in the complex regime differ from those in the rational expectations regime, mainly in that kurtosis is evident in the complex case (Table 1) and that volume of shares traded (per 10,000 periods) is about 300% larger in the complex case, reflecting the degree to which the agents remain heterogeneous in their expectations as the market evolves. We note that fat tails and high volume are also characteristic of price data from actual financial markets.

	Mean	Std. Dev.	Skewness	Kurtosis	Vol. traded
R.e.e. Regime	0.000	2.1002	0.0131	0.0131	2,460.9
Complex Regime	0.000	2.1007	0.0204	0.3429	7,783.8

Table 1. Returns and volume statistics (medians) in the two regimes collected for 25 experiments after 250,000 periods.

How does technical trading emerge in psychologically-rich or complex regime? In this regime the “temperature” of exploration is high enough to offset to some degree expectations’ natural attraction to the r.e.e. And so, subsets of non-r.e.e. beliefs need not disappear rapidly. Instead they can become mutually reinforcing. Suppose, for example, predictors appear early on that by chance condition an upward price

forecast upon the market's showing a current rising trend. Then agents who hold such predictors are more likely to buy into the market on an uptrend, raising the price over what it might otherwise be, causing a slight upward bias that might be sufficient to lend validation to such rules and retain them in the market. A similar story holds for predictors that forecast reversion to fundamental value. Such predictors need to appear in sufficient density to validate each other and remain in the population of predictors. The situation here is analogous to that in theories of the origin of life, where there needs to be a certain density of mutually-reinforcing RNA units in the "soup" of monomers and polymers for such replicating units to gain a footing (Eigen and Schuster, 1979; Kauffman 1993). Thus technical analysis can emerge if trend-following (or mean-reversion) beliefs are by chance generated in the population, and if random perturbations in the dividend sequence activate them and subsequently validate them. From then on, they may take their place in the population of patterns recognized by the agents and become mutually sustainable. This emergence of structure from the mutual interaction of system subcomponents justifies our using the label "complex" for this regime.

What is critical to the appearance of subpopulations of mutually-reinforcing forecasts, in fact, is the presence of market information to condition upon. Market states act as "sunspot-like" signals that allow predictors to coordinate upon a direction they associate with that signal. (Of course, these are not classic sunspots that convey no real information.) Such coordination or mutuality can remain in the market once it establishes itself by chance. We can say the ability of market states to act as signals primes the mutuality that causes complex behavior. There is no need to assume a separate class of noise traders for this purpose. We can test this signaling conjecture in further experiments where we "turn off" the condition part of all predictors (by filling them with non-replaceable #'s). Now forecasts cannot differentiate among states of the market, and market states cannot act as signals. We find, consistent with our conjecture that signaling drives the observed patterns, that the complex regime does not emerge. As a further test of the significance of technical trading signals, we regress the current price on the previous period's plus the technical indicator (price  $> 500$  period moving average). In the rational-expectations regime, the technical indicator is of course not significant. In the complex regime, the trend indicator *is* significant (with t-value of 5.1 for the mean of the sample of 25 experiments), showing that the indicator does indeed carry useful market information. The corresponding test on actual financial data shows a similar result (Brock, Lakonishok, and LeBaron, 1991).

One of the striking characteristics of actual financial markets is that both their price volatility and trading volume show persistence or autocorrelation. And volatility and volume show significant cross-correlation. In other words, both volume and volatility remain high or low for periods of random length, and they are interrelated. Our inductive market also shows persistence in volatility or GARCH behavior in the complex regime, Fig. 4, (with the Chi-square statistic in the Engle GARCH Test significant at the 95%

level).<sup>12</sup> It also shows persistence in trading volume, Fig. 5., as well as significant cross-correlation between trading volume and volatility, Fig. 6. The figures include corresponding correlations for the often-used market standard, IBM stock. (Note that because our time period and actual market days do not necessarily match, we should expect no exact overlap. But qualitatively, persistence in our market and IBM's is similar.) These correlations are not explained by the standard model, where theoretically they are zero.

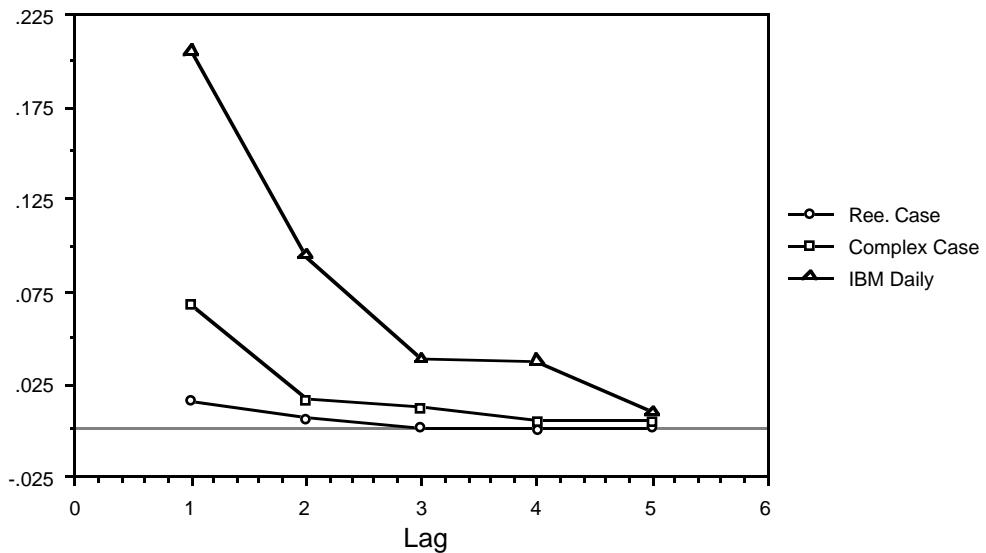


Fig. 4. Autocorrelation of volatility in rational-expectations and complex regimes, and in IBM daily returns

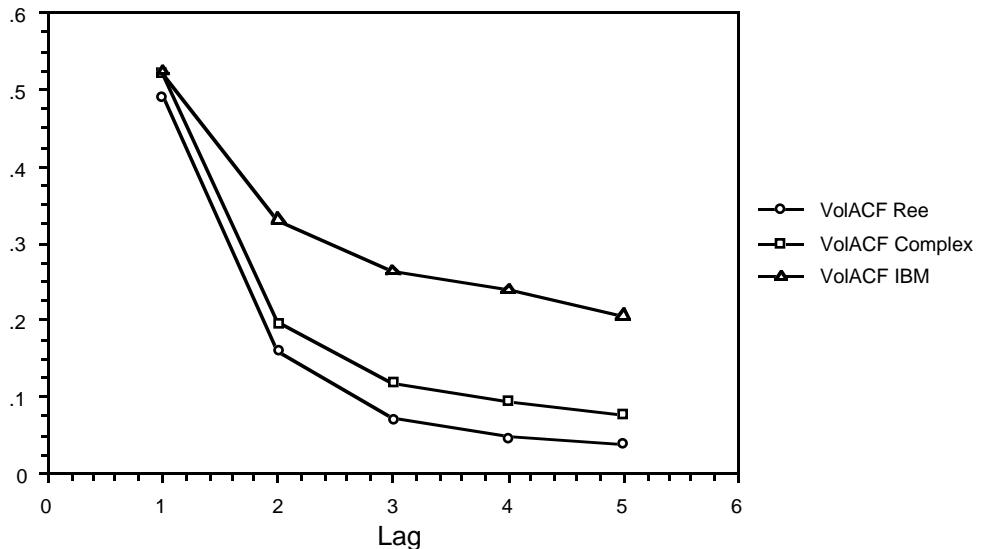


Fig. 5. Autocorrelation of trading volume in the rational-expectations and complex regimes, and in IBM daily returns

<sup>12</sup> Autocorrelated volatility is often fitted with a Generalized Autoregressive Conditional Heteroscedastic time series. Hence the GARCH label. See Bollerslev *et al.* 1990, and Goodhart and O'Hara, 1995.

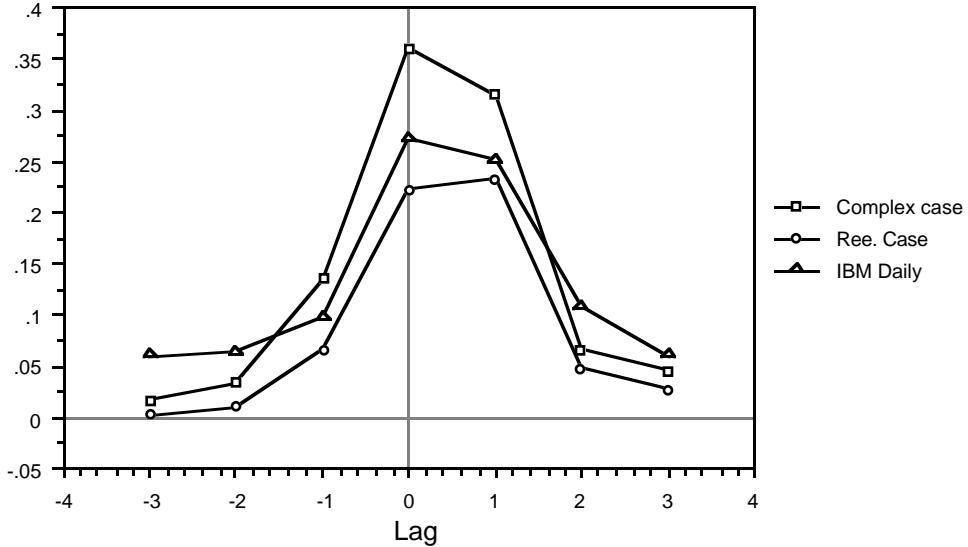


Figure 6. Crosscorrelation of trading volume with volatility, in the rational-expectations and complex regimes, and in IBM daily returns

Why financial markets—and our inductive market—show these empirical “signatures” remains an open question. We conjecture a simple evolutionary explanation. Both in real markets and in our artificial market, agents are constantly exploring and testing new expectations. Once in a while, randomly, more successful expectations will be discovered. Such expectations will change the market, and trigger further changes in expectations, so that small and large “avalanches” of change will cascade through the system. (Of course, on this very short time-lag scale, these avalanches occur not through the genetic algorithm but by agents changing their active predictors.) Changes then manifest in the form of increased volatility and increased volume. One way to test this conjecture is to see whether autocorrelations increase as the predictor accuracy-updating parameter  $\theta$  in (7) in Appendix A is increased. The larger  $\theta$  is, the faster individual agents “switch” among their predictors. Thus the more such switches should cascade. Experiments confirm that autocorrelations indeed increase with  $\theta$ . Such cascades of switching in time are absorbed by the market, and die away. Hence our evolutionary market exhibits periods of turbulence followed by periods of quiescence, as do actual markets.<sup>13</sup>

## 6. Discussion

To what extent is the existence of the complex regime an artifact of design assumptions made in our model? We find experimentally by varying both the model’s parameters and the expectational-learning

<sup>13</sup> For a discussion of volatility clustering in a different model, see Youssefmir and Huberman, 1995; and also Grannan and Swindle, 1994.

mechanism, that the complex regime and the qualitative phenomena associated with it are robust. These are not an artifact of some deficiency in the model.<sup>14</sup>

It might be objected that if some agents could discover a superior forecasting means to exploit the market, this might arbitrage complex patterns away, causing the market again to converge to rational expectations. We believe not. If a clever meta-expectational model was “out there” that might exploit others’ expectations, such a model would, by aggregation of others’ expectations, be a complicated non-linear function of current market information. To the degree that the piecewise linear form we have assumed covers the space of nonlinear expectational models conditioned on current market information, agents would indeed, via the genetic algorithm, pick up on an approximate form of this superior meta-model. The complex regime owes its existence then not to limitations of forecasting, but rather to the fact that in our endogenous-expectations model, market information can be used as signals, so that a much wider space of possibilities is open—in particular, the market can self-organize into mutually supporting subpopulations of predictors. (In fact, in a simpler, analytical model, with a small number of classes of trader whose beliefs adapt endogenously, Brock and Hommes (1996) find similar, rich, asset-price dynamics.) There is no reason these emergent subpopulations should be in stochastic equilibrium. Indeed agents may mutually adapt their expectations forever, so that the market explores its way through this large space, and is non-stationary. In some early exploratory experiments, we “froze” successful agents’ expectations, then reinjected these agents with their previously successful expectations much later. The reintroduced agents proved less successful than average, indicating that the market had evolved and was non-stationary.

It might be also objected that by our use of condition-bits in the predictors, we have built technical trading in to our model. And so it is no surprise that it appears in the complex regime. But actually, only the possibility of technical trading is built in, not its use. The use of market descriptors is selected against in the model. Thus market signals must be of value to be used, and technical trading emerges only because such market signals induce mutually supporting expectations that condition themselves on these market signals.

If the market has a well-defined psychology in our model, does it also experience “moods”? Obviously not. But notice we assume that agents entertain more than one market hypothesis. Thus we can imagine circumstances of a prolonged “bull-market” uptrend to a level well above fundamental value in which the market state activates predictors that indicate the uptrend will continue, and simultaneously other predictors that predict a rapid downward correction. Such combinations, which occur easily in both our market and actual markets, could well be described as “nervous.”

---

<sup>14</sup> One design choice might make a difference. We have evaluated the usefulness of expectational beliefs by their accuracy rather than by profit they produce. In practice these alternatives may produce different outcomes. For example, buying into a price rise on the basis of expectations may yield a different result if validated by profit instead of by accuracy of forecast, when “slippage” is present, that is, when traders on the other side of the market are hard to find. We believe, but have not proved, that the two criteria lead to the same qualitative results.







with random predictors: condition bits are set to 0 or 1 with probability 0.1, otherwise to  $\#$ . This avoids bias in choosing predictors at the outset, and allows intelligent behavior to bootstrap itself up as the artificial agents generate predictive models that perform better. For the bitstrings these procedures are standard genetic algorithm procedures for mutation and crossover (uniform crossover is used which chooses a bit at random from each of the two parents). The forecasting parameter vectors are mutated by adding random variables to each individual component. And they are crossed over component-wise, or by taking linear combinations of the two vectors, or by selecting one or the other complete vector. Each of these procedures is performed with equal probability. Cross-over on a predictor is performed with probability 0.3 or 0.1 in the r.e. and complex regimes, respectively. Individual bits are mutated with probability 0.03. New predictors are brought into the predictor set with variance set to the average of their parents. If a bit has been changed, the new predictor's variance is set to the average of that of all predictors. If this new variance is lower than the variance of the current default predictor less an absolute deviation, its variance is set to the median of the predictors' variance. This procedure gives new predictors a reasonable chance of becoming used.

*Market Clearing.* The price is adjusted each period by directly solving (5) and (6) for  $p$ , which entails passing agents' forecasting parameters to the clearing equation. In actual markets, of course, the price is adjusted by a specialist who may not have access to agents demand functions. But we note that actual specialists, either from experience or from their "books," have a keen feel for the demand function in their markets, and use little inventory to balance day-to-day demand. Alternatively, our market-clearing mechanism simulates an auction in which the specialist declares different prices and agents continually resubmit bids until a price is reached that clears the market.

*Calculation of the Homogeneous Rational Expectations Equilibrium.* We calculate the homogeneous r.e.e. for the case where the market price is a linear function of the dividend  $p_t = fd_t + g$  which corresponds to the structure of our forecasts. We can then calculate  $f$  and  $g$  from the market conditions at equilibrium. A homogenous equilibrium demands that all agents hold 1 share, so that, from (5)

$$E_t(p_{t+1} + d_{t+1}) - (1+r)p_t = \lambda\sigma_{p+d}^2 \quad (9)$$

From the dividend process (4) and the linear form for the price, we can calculate  $\sigma_{p+d}^2 = (1+f)\sigma_e^2$  and  $E_t(p_{t+1} + d_{t+1})$  as

$$E_t(p_{t+1} + d_{t+1}) = (1+f)[(1-\rho)\bar{d} + \rho d_t] + g.$$

Noting that the right side of (9) is constant, we can then solve for  $f$  and  $g$  as

$$f = \rho / (1 + r - \rho)$$

$$g = (1 + f) [(1 - \rho) \bar{d} - \lambda \sigma_e^2] / r$$

Therefore the expression:

$$E_t(p_{t+1} + d_{t+1}) = (1 + r)p_t + \frac{\lambda(2 + r)\sigma_e^2}{(1 + r - \rho)} \quad (10)$$

is the homogeneous r.e.e. forecast we seek.

## Appendix B: The S F I Artificial Stock Market

The Santa Fe Artificial Stock Market has existed since 1989 in various designs (see Palmer *et al.* 1994 for a description of an earlier version). Since then a number of other artificial markets have appeared: e.g. Beltratti and Margarita, 1992; Marengo and Tordjman, 1995; Rieck, 1994. The Santa Fe Market is a computer-based model that can be altered, experimented with, and studied in a rigorously controlled way. Most of the artificial market's features are malleable and can be changed to carry out different experiments. Thus the artificial market is a framework or template that can be specialized to focus on particular questions of interest in finance: for example, the effects of different agents having access to different information sets or predictive behaviors; or of a transaction tax on trading volume; or of different market-making mechanisms.

The framework allows other classes of utility functions, such as constant relative risk aversion. It allows a specialist or market maker, with temporary imbalances in fulfilled bids and offers, made up by changes in an inventory held by the specialist. It allows a number of alternative random processes for  $\{d_t\}$ . And it allows for the evolutionary selection of agents via wealth.

The market runs on a NeXTStep computational platform, but is currently being ported to the Swarm platform. For availability of code, and for further information, readers should contact Blake LeBaron or Richard Palmer.

## References

Arthur W.B. "Inductive Behavior and Bounded Rationality," *Amer. Econ. Review*, 84, May 1994.

Arthur W.B. "On Learning and Adaptation in the Economy," Santa Fe Institute Paper 92-07-038, 1992.

Arthur, W.B. "Complexity in Economic and Financial Markets," *Complexity*, 1, 1995.

Beltratti, A. and S. Margarita, "Simulating an Artificial Adaptive Stock Market," mimeo, Turin Univ. 1992.

Binmore, K. "Modeling Rational Players," *Economics and Philosophy*, 3, 179-214, 1987.

Blume L. and D. Easley, "Evolution and Market Behavior," *J. Economic Theory*, 58, 9-40, 1990.

Blume L. and D. Easley, "Learning to be Rational," *J. Economic Theory*, 26, 340-351, 1982.

Bollerslev, T., R. Y. Chou, N. Jayaraman, and K. F. Kroner, "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5-60, 1990.

Bossaerts, P. "The Econometrics of Learning in Financial Markets," *Econometric Theory*, 11, 151-189, 1995.

Brock, W. and C.H. Hommes, "Models of Complexity in Economics and Finance," Mimeo. Dept. of Econ, Univ. of Wisconsin, Madison.

Brock, W., J. Lakonishok, and B. LeBaron, "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." Santa Fe Institute paper 91-01-006, 1991.

Cutler, D.M. and J.M. Poterba, and L.H. Summers, "What Moves Stock Prices?" *Journal of Portfolio Management*, 4-12, 1989.

De Long, J.B., A. Shleifer, L.H. Summers, R.J. Waldmann, "Positive Feedback Strategies and Destabilizing Rational Speculation," *J. of Finance*, 45, 1990a.

De Long, J.B., A. Shleifer, L.H. Summers, R.J. Waldmann, "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98, 703-738, 1990b.

De Long, J.B., A. Shleifer, L.H. Summers, R.J. Waldmann, "The Survival of Noise Traders in Financial Markets," *J. of Business*, 64, 1-18, 1991.

Diba B.T. and H. I. Grossman, "The Theory of Rational Bubbles in Stock Prices," *Economic Journal*, 98, 746-754, 1988.

Eigen, M. and P. Schuster, *The Hypercycle: A Principle of Natural Self-Organization*. Springer, NY, 1979.

Frankel, J.A. and K.A. Froot, "Chartists, Fundamentalists, and Trading in the Foreign Exchange Market," *AEA Papers and Proceedings*, 80, 181-185, 1990.

Friedman, D. and Aoki, M. "Inefficient Information Aggregation as a Source of Asset Price Bubbles," *Bulletin of Economic Research*, 44, 251-79, 1992.

Goodhart C. A. E. and M. O'Hara, "High Frequency Data in Financial Markets: Issues and Applications," mss, London School of Economics, 1995.

Grannan E.R. and G.H. Swindle, "Contrarians and Volatility Clustering," SFI paper 1994.

Grossman, S.J. "On the Efficiency of Competitive Stock Markets where Traders have Diverse Information," *Journal of Finance*, 31, 573-85, 1976.

Grossman, S.J. and J. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393-408, 1980.

Holland J.H., K.J. Holyoak, R.E. Nisbett, P.R. Thagard, *Induction*, Cambridge, MA: MIT Press, 1986.

Kauffman, S. *The Origin of Order: Self-Organization and Selection in Evolution*, New York, Oxford Univ. Press, 1993.

Keynes J. M. *General Theory of Employment, Interest and Money*, London, Macmillan, 1936.

Kirman A., "Epidemics of Opinion and Speculative Bubbles in Financial Markets," in *Money and Financial Markets*, M. Taylor, ed., London, Macmillan, 1991.

Kleidon A.W. "Anomalies in Financial Economics: Blueprint for Change?"

Kurz, M., "On the Structure and Diversity of Rational Beliefs," *Economic Theory*, 1994.

Kurz, M. "Asset Prices with Rational Beliefs," Stanford Econ. Working Paper 96-003, 1995.

Leroy, S.F. and R.D. Porter, "Stock Price Volatility: Tests based on Implied Variance Bounds," *Econometrica*, 49, 97-113, 1981.

Lo, A. W. and C. MacKinlay, "Stock Prices Do Not Follow Random Walks: Evidence From a Simple Specification Test," *Review of Financial Studies*, 1, 41-66, 1988.

Lucas R.E. "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429-1445, 1978.

Marengo, L. and H. Tordjman, "Speculation, Heterogeneity and Learning: A model of Exchange Rate Dynamics," IIASA Working Paper WP-95-17, 1995.

Milgrom, P. and N. Stokey, "Information, Trade and common Knowledge," *J. Econ. Th.* 26, 17-27, 1982.

Nelson R. and S. Winter *An Evolutionary Theory of Economic Change*, Harvard: Bellknap) 1982.

Palmer, R.G., W. B. Arthur, J.H. Holland, B. LeBaron, P. Tayler, "Artificial Economic Life: A Simple Model of a Stockmarket. *Physica D*, 75, 264-274, 1994.

Rieck, C. "Evolutionary Simulation of asset Trading Strategies," in *Many-Agent Simulation and Artificial Life*, E. Hillenbrand and J. Stender (eds.) IOS Press, 1994.

Sargent, T.J., *Bounded Rationality in Macroeconomics*, New York: Oxford Univ. Press, 1993.

Shiller, R.J. "Do Stock Prices Move too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71: 421-36, 1981.

Shiller, R.J. *Market Volatility*, Cambridge, MA: MIT Press, 1989.

Shleifer A. and L. H. Summers, "The Noise Trader Approach to Finance," *J. of Economic Perspectives*, 4, 19-33, 1990.

Shubik, M. "Time and Money," SFI Paper 96-03-013, 1996.

Soros, G. *The Theory of Reflexivity*. Soros Fund Management, New York, NY. 1994.

Summers, L.H. "Does the Stock Market Rationally Reflect Fundamental Values?" *Journal of Finance*, 46, 591-601, 1986.

Thaler, R.H. *Advances in Behavioral Finance*, Russell Sage Foundation, New York, 1993.

Youssefmir, M., B. Huberman. "Clustered Volatility in Multiagent Dynamics," SFI Paper 95-05-051, 1995.