Complexity in Economic and Financial Markets [§]

by W. Brian Arthur

Citibank Professor, Santa Fe Institute, Santa Fe; and Morrison Professor of

Economics and Population Studies, Stanford University, Stanford, California .

Abstract

Actions taken by economic decision makers are typically predicated upon hypotheses or predictions about future states of a world that is itself in part the consequence of these hypotheses or predictions. When we attempt to model how such predictions might be generated we become stymied: the predictions some economic agents might form depend on the predictions they believe others might form; and the predictions these might form depend upon the predictions *they* believe the original group might form. Predictions or expectations can then become self-referential and deductively indeterminate. This indeterminacy pervades economics and game theory.

This paper argues that in such situations agents predict not deductively, but inductively. They form subjective expectations or hypotheses about what determines the world they face. And these expectations are formulated, used, tested, and possibly changed, in a world that forms from others' subjective expectations. This yields individual expectations trying to prove themselves against others' expectations. The result is an ecology of co-evolving, possibly ever-changing expectations. The resulting dynamics often can be analyzed only by computation.

Inductive expectation formation is illustrated in an artificial computer-based stock market. Coevolution of expectations explains phenomena seen in real markets that appear as anomalies to standard finance theory.

Acknowledgments

Complexity in Economic and Financial Markets

by W. Brian Arthur

One way to look at the economy, the standard way in fact, is to view it in physical terms as a collection of activities, technologies, and needs, all interacting though a market system peopled by decision-making agents such as firms, banks, consumers, and investors. A very different way—the one I want to explore here—would be to view the economy in psychological terms: as a collection of beliefs, anticipations, expectations, and interpretations; with decision-making and strategizing and action-taking predicated upon these beliefs and expectations. Of course, the two views are related. Activities follow from beliefs and expectations. And beliefs and expectations are mediated and sculpted by the physical economy they find themselves in.

Why might a psychological or cognitive view of the economy be useful? Economic agents make their choices based upon their current beliefs or hypotheses (I will use these terms along with the jargon terms *expectations* or *predictions*) about future prices, or future interest rates, or competitors' future moves, or the future character of their world. And these choices, when aggregated, in turn shape the prices, interest rates, market strategies, or world these agents face. The beliefs or hypotheses that agents form in the real economy are largely individual and subjective. They are often private. And they are constantly tested in a world that forms from their and others' actions—a world that is ultimately formed from their and other agents' subjective beliefs. Thus at a sub-level, we can think of the economy ultimately as a vast collection of beliefs or hypotheses, constantly being formulated, acted upon, changed and discarded; all interacting and competing and evolving and coevolving; forming an ocean of ever-changing, predictive models-of-the-world. This view is useful, I believe, because it forces us to think about how beliefs create economic behavior—and how economic outcomes create beliefs. And it leads to different insights. Beyond the simplest problems in economics, this ecological view of the economy becomes inevitable; and it leads to a world of complexity.

The standard way to handle predictive beliefs in economic analysis is to assume identical agents who possess perfect rationality and arrive at shared, logical conclusions or expectations

about the situation they face. When these expectations induce actions that aggregatively create a world that validates them as predictions, they are in equilibrium and are called *rational expectations.* Rational expectations are useful in demonstrating logical equilibrium outcomes and analyzing their consequences. But in the real world they break down easily. If some agents lack the computing power to deduce the posited outcome; or if some arrive logically at different conclusions from the same data (as they might in a pattern recognition problem); or if there is more than one rational expectations equilibrium with no means to coordinate which is chosen; then some agents may deviate in their expectations. And if some deviate, the world that is created may change, so that others should logically predict something different and deviate too. And so rational expectations can unravel easily. Unless there are special circumstances, they are not robust.

There is a game in economics that illustrates this unraveling of rational expectations beautifully. It is the Guessing Game, where *N* players choose a number between 0 and 100, and the winner is the one closest to 2/3 of the average guess (see Nagel [16]). Obviously here, beliefs of what constitutes a good guess depend on one's view of others' beliefs of what constitutes a good guess. Now, uniform predictions of zero would constitute rational expectations; they would be self-validating in that if agents expected other agents to choose zero they should also choose zero. Therefore expectations that everyone will choose zero would be in equilibrium. And no other real number, if chosen by all, would constitute an expectational equilibrium. But does that mean that zero will necessarily be chosen? If I, as a reference player, suspect that some players—or even one player—may choose non-zero, then logically I ought to choose non-zero. And if I believe that others believe that someone may choose non-zero, I will deduce that they too will choose non-zero. Thus beliefs that some may choose non-zero lead others to expect non-zero and choose non-zero. The game leads to a self-referential sequence of "If they choose *x*, I and others should choose *y*. But if I and others choose *y*, they will have to choose *z*." There is no closure here, and ultimately beliefs or expectations in this game are deductively indeterminate, no matter how logical or rational the agents are.[1]

Consider as a second example my Bar Problem (Arthur [2]). One hundred people must decide independently each week whether to show up at their favorite bar (*El Farol* in Santa Fe, say). The rule is that if a person predicts that more that 60 (say) will attend, he will avoid

---

[1] For a readable account of logical unravelling in Guessing Games with actual human subjects see Nagel [16].

the crowds and stay home; if he predicts fewer than 60 he will go. This seems innocuous; but it destroys the possibility of long-run shared expectations. If all believe *few* will go, than *all* will go, thus invalidating these expectations. And if all believe *many* will no, *no one* will go, invalidating those expectations. Predictions of how many will attend depend on others' predictions, and others' predictions of others' predictions. Once again there is no rational means to arrive at deduced *a-priori* predictions.

These two problems are of course toy problems, concocted like the famous Prisoner's Dilemma to make a point. But they illustrate a foundational difficulty in economics. Where forming expectations means predicting an aggregate outcome that is formed in part from others' expectations, expectation formation can become self-referential. The problem of logically forming expectations then becomes ill-defined, and rational deduction finds itself with no bottom ground to stand upon. This indeterminacy of expectation-formation is by no means a rarity or anomaly within the real economy. On the contrary, it pervades all of economics and game theory.

What do we as humans do in cases like this? If the problem were a one-off one, the answer would depend largely on what is currently in our minds. But if the situation were repeated, as humans, we tend with high reliability to follow a well-known logical procedure. As the situation is replayed regularly, we look for patterns; and we use these to construct temporary expectational models or hypotheses to work with. We carry out localized deductions based on these hypotheses and act on them. As feedback from the repeated plays comes in, we strengthen or weaken our beliefs in our current hypotheses, discarding some when they cease to perform, and replacing them as needed with new ones. In other words, to use Tom Sargent's phrase [17], we act as statisticians, using and testing and discarding simple expectational models to fill the gaps in our understanding. Such behavior is perfectly logical; in fact on a micro-scale it is the scientific method. But as logic it is not deductive, but *inductive*. [2]

Financial Markets

---

[2] For a full account of induction, see Holland *et al.* [10]; and for earlier versions of induction applied to asset pricing, to decision problems and to the Bar problem, see Arthur [1] and [2]. Sargent [17] applies an inductive approach to problems in the macro economy.

To see that expectation formation matters, and that circularity in the formation of expectations is not merely an academic worry, let us look in some detail at a major arena of real economic activity, financial markets, and within them the formation of stock market prices.

One of the major puzzles in finance is that academic theorists, by and large, see markets differently from the way traders or practitioners do. The academic view sees investors as perfectly rational, from which it follows that markets are efficient in the sense that all usable information is discounted into current prices. Thus speculative profits, via technical trading (using past prices to forecast future prices) or other means, are not obtainable. Temporary bubbles and crashes cannot occur, other than as adjustments to market news. (The information that might lead to them would be discounted into the price instantaneously.) And such notions as a "market psychology" are absurdities. The theory lends useful insight into the workings of financial markets. But it is by no means shared by traders themselves. They, in contrast, tend to see markets as offering speculative opportunities. Many believe that technical trading is profitable, and that "market psychology" and herd effects unrelated to market news can cause bubbles and crashes. Some traders and financial writers even see the market itself as possessing "moods," sometimes describing it as "nervous" or "sluggish" or "jittery."

Which side is correct? Empirically this is hard to settle. Markets do turn out to be reasonably efficient; but statistics show that trading volume and price volatility in real markets are a great deal higher than the standard theory predicts (Shiller [18]). Statistical tests also show that technical trading can produce consistent, if modest, long-run profits (Brock *et al.* [5]). And the crash of 1987 showed dramatically to both traders and economists that sudden price changes do not always reflect rational adjustments to news in the market (Cutler *et al.* [6]).

Let us begin our examination of financial markets by looking at the standard theory. Full-dress models here are complicated (for example, Lucas [15]), but at their core is a simple logic (see Diba and Grossman [8]) which I will lay out, ignoring extraneous technicalities. In essence, this logic is going to say there is a natural (or fundamental) valuation at any time of a stock, and *given that all investors share the expectational model that produces this valuation*, it will be realized and will therefore be in rational expectations equilibrium. The difficulty will come with the stipulation "given that all investors share the expectational model that produces this valuation." Once we look at *how* investors might come to share this

expectation, or how they might form a good expectational model, the solution will unravel and we will have to turn to inductive reasoning to resolve stock pricing behavior.

So let us now go into the standard argument in some detail. Assume that at any time t, N investors can purchase shares in a single security, a stock say, and this provides unknown-in-advance stochastic payoffs or dividends $\{d_t\}$, exogeneously determined. Investors place the remainder of their wealth in a risk-free asset, a T-bill say, that pays a constant r units per period. Market information I_t (economic indicators, rumors, historical price data, Reuters news, and the like) is available at time t, which may suggest the direction the dividend may move. On the basis of this information, the theory has each investor deciding at each time how much wealth to allocate between the stock and the T-bill. What interests us is not this allocation, but how the stock price forms and moves in the resulting market.

The standard reasoning now runs as follows. Assume *homogeneous* investors who use the available information I_t identically in forming unbiased expectations $\{ E[d_{t+k}|I_t] \}$ about future dividends at times $t+k$. We want to show that they can reach deductive conclusions about the market price. Now suppose the stock is bid for each day in an auction among these identical investors. It is a simple exercise [8] to show that its price will settle at the level:

$$p_t = \beta \big(E[p_{t+1}|I_t] + E[d_{t+1}|I_t] \big) \qquad (1)$$

where $E[p_{t+1}|I_t]$ is the investors' shared expectation of tomorrow's price. This tells us that a share of the stock today is worth tomorrow's price plus any intervening dividend, discounted by the factor $\beta = 1/(1+r)$, which reflects that a dollar tomorrow is worth β dollars today if invested overnight in the T-bill. In other words, today's stock price p_t, is arbitraged to a value that reflects its current expected worth or *fundamental value*, given the shared beliefs derived from the information I_t.

But today's price is still not fully determined; it depends upon our investors' expectation of tomorrow's price $E[p_{t+1}|I_t]$. How do our infinitely rational investors form this expectation? They may take mathematical expectations across the above arbitrage equation as it will play out tomorrow (*ie.* with time index $t+1$),

$$E[p_{t+1}|I_t] = \beta \big(E[p_{t+2}|I_t] + E[d_{t+2}|I_t] \big) \qquad (2)$$

But this process leaves the expectation of the price in two days' time, $E[p_{t+2}|I_t]$, unknown. In fact, by repeating this operation for all times from $t+2$ onward, substituting future price expectations backward and eventually into (1), our rational investors can eliminate all price expectations, and settle the current price simply as

$$p_t = \sum_{k=1}^{\infty} \beta^k E[d_{t+k}|I_t] \, . \qquad (3)$$

These deduced beliefs that the stock's price are the discounted summation of the values $E[d_{t+k}|I_t]$ (the expected valuation of the future dividend stream) are unbiased, and will on average be upheld. They therefore constitute a rational expectations equilibrium. Such beliefs will of course fluctuate day by day as the information I_t changes, and so the daily price moves randomly, yielding the random-walk character noted in the popular Wall Street literature.

We now have a rational-expectations theory. But it is valid only under certain conditions in fine print, so to speak. The solution requires: (i) that identical, known, unbiased expectations of future dividends $E[d_{t+k}|I_t]$ are given; (ii) that the agents are perfectly rational; (iii) that they know that the price at each time will be formed by arbitrage as in (1); and (iv) that (i) (ii) and (iii) are known and accepted by all.

In passing let us note the implications of this reasoning that were mentioned earlier. Because we assume that dividend expectations are unbiased, prices at each time on average reflect "correct" or fundamental value so speculative profits are not available except by luck. Thus any information contained in past prices say, or in news, cannot generate profits; it would already be discounted in the price. Therefore the market is efficient and technical trading (using past prices to forecast future ones) cannot generate consistent profits. Price crashes (and temporary bubbles) are ruled out as well: sudden bad news is instantaneously fully reflected in the price, and can therefore not cause a snowball effect. These ideas—whether correct or not—are buried deep in our half-conscious notions of how markets operate.

Let us now redo the pricing exercise, but this time more realistically not taking condition (i) for granted. In fact, we will assume our *N* investors are heterogeneous: each possibly

differs from the others; and we will ask how intelligent price and dividend expectations might be formed.

As mentioned, the available shared information I_t consists of past prices, past dividends, trading volumes, economic indicators, rumors, news, and the like. But these are merely qualitative information plus sequences of data points, and *there may be many different, perfectly defensible statistical ways* based on different assumptions and different error criteria to use them to estimate future dividends [1] [14]. The information I_t is a Rorschach inkblot, as it were, and so there is no objective means for one investor to know other investors' expectations of future dividends. Neither is there an objectively laid-down expectational model that differing investors can coordinate upon. We immediately enter the dilemma pointed out in the Guessing Game above. Given that I do not know others' expectations, and that I know they will form them on the basis of *their* expectations of others' expectations, how can I form mine?

Let us follow the story logically, and watch it unravel. Assuming individual expectations $E_i[d_{t+1}|I_t]$ and $E_i[p_{t+1}|I_t]$ (now indexed for investor $i= 1,2,...,N$) the same argument as before yields the arbitrage equation:

$$p_t = \beta \sum_i \frac{1}{N}\left(E_i[p_{t+1}|I_t] + E_i[d_{t+1}|I_t]\right) \qquad (1')$$

Once again, this contains expectations of tomorrow's price. But how can agent i form this expectation? He must either act subjectively, or alternatively try to deduce it by applying the pricing equation $(1')$ to period $t+1$:

$$E_i[p_{t+1}|I_t] = \beta E_i\left[\sum_j \frac{1}{N}\left(E_j[p_{t+2}|I_t] + E_j[d_{t+2}|I_t]\right)\right] \qquad (2')$$

This calls for agent i's expectations of *everyone's* expectations of the dividend and price at time $t+2$. Substituting for *this* price expectation in turn yields

$$E_i[p_{t+1}|I_t] = \beta E_i\left[\sum_j \frac{1}{N}\left(E_j[d_{t+2}|I_t] + \beta E_j\left[\sum_k \frac{1}{N}\left(E_k[p_{t+3}|I_t] + E_k[d_{t+3}|I_t]\right)\right]\right)\right] \qquad (4)$$

This leads agent *i* into taking into account his expectations of others' expectations of others' expectations of future dividends and prices. Keynes famously remarked upon this, when he pointed out that valuing stocks called for taking into account "what average opinion expects the average opinion to be" [13]. Expectation formation by deduction is now facing severe difficulties. There is no objective means by which others' dividend expectations can be known, and the repeated iteration of subjective expectations of subjective expectations merely widens the uncertainty and never comes to closure.

Worse, the situation here leads to instability. As we see in (2′), investor *i*'s expectation of the price at *t*+1 depends on his ideas of other investors' expectations of the price at *t*+2. Thus the sense he makes of the Rorschach pattern of market information I_t is influenced by the sense he *believes* others may make of the same pattern. If he believes that others believe the price will increase, he will revise his expectations to anticipate upward-moving prices (in practice helping validate such beliefs). If he believes others believe a reversion to lower values is likely, he will revise his expectations downward. All we need to have self-reinforcing suspicions, hopes, and apprehensions rippling through the subjective formation of expectations (as they do in real markets) is to allow that I_t contains hints—and imagined hints—of others' intentions. Under heterogeneity, deductive logic leads not just to indeterminate expectations but also to unstable ones.

Notice the difficulty depends in no way on investors being limited in their reasoning powers. It merely says that given differences in agent expectations, deductive expectation formation becomes indeterminate, and so even perfectly rational investors cannot form expectations in a determinate way. So our investors can be arbitrarily intelligent, yet be unable to implement deductive rationality.

Life in an Ocean of Expectations

We have seen that deductive reasoning unravels if we take expectation formation seriously. How then might our inductive approach fare? And given the innate complication of dealing with not just one expectational model but an ocean of beliefs, how might its implications be studied?

To study asset pricing under inductive reasoning, John Holland, Blake LeBaron, Richard Palmer, Paul Tayler and I have created a artificial stock market on the computer, peopled by investors who are individual, artificially-intelligent computer programs that can reason inductively [3]. In our market-within-the-machine, each artificial investor acts as a market statistician, and continually creates multiple "market hypotheses"—subjective, expectational models—of what moves the market price and dividend.[3]

These expectational models are simple mathematical expressions based on past dividends and prices and conditioned to "recognize" particular states of the market. For example, (in words) a model might be "If today's price is higher than its average in the last 100 days, predict that tomorrow's price will be 3% higher than today's. Each agent possesses many such models, and can therefore have different predictive hypotheses that "recognize" different states of the market. Each acts on his most accurate model or hypothesis that applies to the current market state. Some expectational models perform well in predicting price and dividend movements; these are retained and acted upon in buying and selling decisions. Others perform badly and are dropped. The agents use the genetic algorithm to produce new forecasting models from time to time (by recombining and mutating those of theirs that work well), as well as to replace poorly performing ones. The agents learn both from discovering useful new expectational models or beliefs, and from finding the ones within their current set that perform well.

The market price of course forms endogenously from the bids and offers of the agents, and thus ultimately from their beliefs. And so the computerized market is its own self-contained, simple, artificial financial world.

In the experiments we carry out with this computerized market, we typically use 100 artificial investors each with 60 expectational models. Thus there are at any time 6,000 expectational models. This is hardly an "ocean" of expectations, but it is still a sizable tidal pool. To see how these beliefs might feed upon each other and co-evolve in this pool, consider expectational models that predict an upward trend. If such expectations appeared by random generation in several agents' sets of expectational models, and if prices trended upward by chance driven by an upward dividend, these trend expectations would become validated and would cause buying behavior, furthering the trend and possibly producing a temporary price bubble. Thus trend expectations in sufficient density in the population of

---

[3] For a non-computational evolutionary approach to asset pricing see Blume and Easley [4]. They assume arbitrary but fixed expectations and investment strategies that compete.

expectations are mutually reinforcing.[4] Suppose the price showed sinusoidal swings of period K unrelated to dividend. Then expectations that tomorrow's price at time $t+1$ equals that at $t-K+1$ would become validated. But use of such expectations would cause buying at the bottom of the cycle and selling at the top; and this would arbitrage the cycle away and render such cyclical expectations inaccurate. Cyclical expectations in combination would therefore be mutually negating, and would disappear from the pool of beliefs. Expectations that price follows fundamental value would be valid as long as they were held by others. But if price were trending upward above fundamental value, and others believed the price rise would continue, this might temporarily lead to the continuation of the trend, so that fundamentalist beliefs might be rendered inaccurate.

From this we see that expectations—like species—can come and go in the ocean of beliefs; they can mutually reinforce or mutually negate each other; and they form a co-evolving ecology of beliefs. The situation is similar to that of Lindgren's beautiful study of co-evolving strategies in the Prisoner's Dilemma [12]. In our actual market a huge variety of beliefs can in principle be generated. How then does the ecology organize itself? How do fundamentalist beliefs fare; and do technical trading beliefs gain a footing? Does the outcome in our experimental market validate standard financial theory or the traders' view?

Our findings are still preliminary. But it appears that the academic view and the traders' view are both upheld, but under different conditions, or in different regimes.

If we start our traders off with identical, fundamental-value expectations (by setting the parameters of all their expectational models to reproduce prices that validate these expectations), we find that deviating, non-fundamentalist expectations cannot get a footing. If a large majority of investors believe the fundamentalist model, the resulting prices will validate it, and deviant predictions that arise by mutation in the population of expectational models will be rendered inaccurate. Thus in our market, the homogeneous rational expectations equilibrium of the standard literature is evolutionarily stable: it cannot be invaded by small numbers of deviating expectations. The standard finance theory, under these special circumstances is upheld.

---

[4] In an interesting set of papers, De Long, Summers and others [7, 19] show analytically that certain expectations can be self-fulfilling, so that herd effects or positive feedbacks can be demonstrated. Typically they have two groups of homogeneous traders with fixed expectational models.

If, however, we start our traders off *in an exactly identical market*, except that their initial expectations are randomly distributed uniformly about the fundamentalist ones, we find that trend-following beliefs that appear by chance have enough density to become self reinforcing in the ecology of beliefs. The situation here is reminiscent of that in certain theories of the origin of life. If a sufficient density of small nucleotide molecules appears in the soup or ocean of molecules, they can become mutually reinforcing and "take off" (Eigen and Schuster [9]). If expectations appear that are mutually reinforcing, they too can "take off." Now, in our artificial financial world, these trend-following expectations should cause temporary price bubbles and crashes to occur, and these are indeed observable in the price series. Thus we can restate all this by saying that, in this initially-heterogeneous-expectations regime, using past prices to forecast future ones—technical trading—becomes an emergent property.

We also find that the market possesses a non-trivial *psychology*. We can define this as the collection of market hypotheses (or expectational models or mental beliefs) that are being acted upon at a given time. We believe (and yet need to validate statistically) that in this regime no stationary equilibrium is reached. There are indications that if we "freeze" successful agents early on and inject them back into the system much later, they do no better than average: the market has evolved and moved on as new strategies are discovered from time to time that exploit the earlier winning strategies and these no longer succeed. Of course, we find no evidence of market "moods." But from time to time we do see agents being torn between beliefs that a price trend will continue upward and beliefs that a downturn is imminent—hints of conjoined "greed" and "fear," the source of nervousness in real markets.

We also observe in the data our inductive-expectations market generates a curious statistical phenomenon that is a signature of actual financial markets. This is GARCH (Generalized AutoRegressive Conditional Heteroscedastic) behavior. Essentially GARCH behavior means that there are periods of persistent high volatility (or high variance) in the price series followed randomly by periods of persistent low volatility. This makes no sense in the standard theory. But in our evolutionary market prices might continue in a stable pattern for quite some time, so there is low price volatility, until new expectations are discovered that exploit that pattern. Then there will be swift changes of *gestalt*, swift readjustments in expectations that change the market itself and cause avalanches of further change. Volatility will remain high until such change is absorbed. GARCH behavior is statistically measurable in

our price series, and resembles that shown in actual financial markets. The initially-homogeneous-expectations regime, by contrast, displays no GARCH signature.

To see these phenomena note that we have no need to invoke "irrationality" or "behaviorism" or "investor sentiment" (cf. [20]). They follow simply from the subjective-inductive nature of belief formation.

We can conclude that given sufficient homogeneity of (unbiased) beliefs, the standard equilibrium of the literature is upheld. The market in a sense in this regime is essentially "dead." As the dial of heterogeneity of initial beliefs is turned up, the market undergoes a phase transition and "comes to life." It develops a rich psychology and displays phenomena regarded as anomalies in the standard theory but observed in real markets. The inductive, ecology-of-expectations model we have outlined is by its definition an *adaptive nonlinear network* (Holland [11]). In its heterogeneous mode it displays complex, pattern-forming, non-stationary behavior. We could therefore rename the two regimes or phases *simple* and *complex*. We conjecture that actual financial markets live within the complex regime.

Conclusion

An economy of course, does indeed consist of technologies, actions, markets, financial institutions and factories—all real and tangible. But behind these, guiding them and being guided by them on a sub-particle level are beliefs: the subjective expectations, multiple hypotheses, and half-hoped anticipations held by real human beings. Beliefs can be mutually-reinforcing, or mutually competing. They can arise, get a footing, become prominent, fall back, and disappear. They can be generated privately by theoretical reasoning or by pattern recognition; they can be transmitted from one agent to another. They shape in aggregate the macro economy; they give rise to the movements of financial markets; they direct flows of capital internationally; they govern strategic behavior; and they govern investment. They are the DNA of the economy, and they are everywhere dense. When beliefs form an ocean of interacting, competing, arising and decaying entities, occasionally they simplify into a simple, homogeneous equilibrium set. More often they produce complex, ever-changing patterns. Within the most significant parts of the economy, interacting, non-equilibrium beliefs are unavoidable, and with these so is a world of complexity.

References

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

2. W. B. Arthur: Inductive Behavior and Bounded Rationality. Amer. Econ. Review, 84, 406-411, May 1994.

3. W. B. Arthur , J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler: An Artificial Stock Market. Forthcoming as Santa Fe Institute paper, 1995.

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

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

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

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

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

9. M. Eigen and Schuster, P. "The Hypercycle: A Principle of Natural Self-Organization." Naturwissenschaften 64, p.541, 1977.

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

11. J.H. Holland: The Global Economy as an Adaptive Process. In The Economy as an Evolving Complex System, P.W. Anderson, K.J. Arrow, D. Pines (eds.) 117-124, Addison-Wesley: Reading, Mass., 1988.

12. K. Lindgren: Evolutionary Phenomena in Simple Dynamics. In Artificial Life II, C.Langton et al. (eds.), 295-312, Addison-Wesley: Reading, Mass., 1991.

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

14. M. Kurz: On the Structure and Diversity of Rational Beliefs. Santa Fe Institute Paper 91-02-013, 1991.

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

16. R. Nagel. Reasoning and Learning in Guessing Games: An Experimental Investigation. Ph.D dissertation. Univ. of Bonn, 1994.

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

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

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

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