



Queen's Economics Department Working Paper No. 854

# On Learning and Adaptation in the Economy

W. Brian Arthur

Department of Economics  
Queen's University  
94 University Avenue  
Kingston, Ontario, Canada  
K7L 3N6

5-1992

**On Learning and Adaptation in the Economy**

by

**W. Brian Arthur**

Stanford University

May 25, 1992

Mackintosh Lecture  
Queen's University  
Discussion Paper #854

Morrison Professor of Population Studies and Economics,  
F. R. I., Stanford University, Stanford, CA 94305-6084.

## Abstract

The standard mode of theorizing assumed in economics is deductive—it assumes that human agents derive their conclusions by logical processes from complete, consistent and well-defined premises in a given problem. This works well in simple problems, but it breaks down beyond a “problem complexity boundary” where human computational abilities are exceeded or the assumptions of deductive rationality cannot be relied upon to hold.

The paper draws upon what is known in psychology to argue that beyond this problem complexity boundary humans continue to reason well, but by using *induction* rather than deduction. That is, in difficult or complex decision problems, humans transfer experience from other, similar problems they have faced before; they look for patterns and analogies that help them construct internal models of and hypotheses about the situation they are in; and they act more or less deductively on the basis of these. In doing so they constantly update these models and hypotheses by importing feedback—new observations—from their environment. Thus in dealing with problems of high complexity humans live in a world of learning and adaptation.

I illustrate these ideas by showing that the processes of pattern recognition, hypothesis formation and refutation over time are perfectly amenable to analysis; and by using them to explain supposedly “anomalous” behavior in financial markets.

## Acknowledgments

Much of the thinking in this paper has been heavily influenced by my association with John Holland at the Santa Fe Institute over the last several years. I thank him in particular. I am also grateful to Kenneth Arrow, Marcel Fafchamps, Frank Hahn, Richard Herrenstein, David Lane, Richard Palmer, David Rumlhart, and other Santa Fe and Stanford colleagues and students for discussions on the ideas in this paper. I also thank the faculty and students at the 1991-92 Mackintosh Lecture at Queen's University, Kingston, Ontario, and in particular Lorne Carmichael for useful suggestions.

## On Learning and Adaptation in the Economy<sup>†</sup>

W. Brian Arthur

After long neglect by economic theorists, learning and adaptation are suddenly "hot" as research subjects in economics. Yet in spite of this surge of interest, it is hard to find much justification for *why* economics should concern itself deeply with learning and adaptive behavior. The new literature convinces us that learning models are useful in spelling out processes of adjustment to the standard equilibria of economic theory; and it raises fascinating questions about what it means to recursively incorporate new information into decision behavior. But it also leaves us with a vague feeling that learning is somehow ancillary to economics—an addition of adjustment dynamics to the core theory, not fully necessary perhaps to theorizing in the field though interesting in its own right.

In this talk I want to argue that learning is not a mere addendum to the standard theory. On the contrary, when economics gets beyond dealing with very simple problems, satisfactory theorizing in economics cannot proceed without facing and resolving key issues that are at heart ones of learning, adaptation, and cognition. I want to show that thinking about decision making in fact raises fundamental questions about how we acquire and process knowledge—how we learn and adapt.

### The Deductive Metaphor

My argument will be that in dealing with problems that are complicated and not well-defined, economic agents are forced to reason somewhat differently from the way our standard notion of rationality pictures things. They are forced to use *inductive* rather than deductive means of reasoning. They must rely heavily on internal models of and hypotheses about the problems they are dealing with; and they constantly monitor, update, and revise these by importing feedback—new data—from their environment. Thus in dealing with problems of any realistic complica-

---

<sup>†</sup> This paper is based on the 1991-92 W. A. Mackintosh Lecture, at Queen's University, Kingston, Ontario, delivered November 18, 1991.

tion, economic agents unavoidably live in a world of learning and adaptation.

But I am getting ahead of the story. Let me begin from first principles by looking at the standard mode of theorizing. Practically all theoretical reasoning in modern economics starts with the assumption that agents possess "perfect rationality," that is that they obey certain axioms of reasonable and logical behavior. This means among other things that agents know what is in their self-interest, act in their self-interest, and are able to perform the calculations necessary to discriminate the implications of alternative decisions. This last requirement in turn implies that agents have the analytical capacity—the smarts, the brainpower—in a well-defined decision problem to figure out the optimal action. This standard mode of theorizing in economics assumes that agents derive their conclusions by logical processes from the givens of each problem—premises that are assumed complete, consistent, and well-defined. These conclusions follow necessarily and inexorably, *deduced* from the given premises. I will call this view of how decisions are made *the deductive metaphor*.<sup>1</sup>

Let us look hard at this mode of theorizing for a moment. Certainly it is highly plausible. Humans are indeed far-sighted and calculating. There is evidence, for example, that in forming their human capital, people do look ahead and deduce roughly the amount of schooling or training that would be economically appropriate for them, given their age and circumstances. And experiments using human subjects in simple economic situations often corroborate the behavior assumed in standard theorizing. In problems like these the deductive metaphor works well.<sup>2</sup>

But how far can we push it? Does the deductive metaphor still work when decisions become more complicated? Let us imagine a scale of varying problem-complexity (or problem-complicatedness, to be more precise). At the low end would be problems like human capital formation and simple games like Tic-tac-toe. Much higher would be problems like scheduling production lines, or games like Chess and Go. Let us examine the deductive mode at work in a problem of intermediate complexity—an industrial-organization problem I have taken from the literature.<sup>3</sup> Consider  $N$  firms deciding where to locate their products on a unit circle. For concreteness, we can suppose these to be 20 airlines deciding where to book take-off times or "slots" for flights (their product) from La Guardia airport, on a 24 hour clock. The airlines may have innate preferences about the placement of their slot times; but they will also be affected by any clustering of other airlines' flights near their takeoff-times which would eat into their demand. How, according

Shleifer and Lawrence H. Summers, 1990, "The Noise Trader Approach to Finance," in *Journal of Economic Perspectives*, 4, 19—33; also papers cited therein by Black; De Long *et al.*; Frankel and Froot; Shiller; and others. A different approach, one that allows for contagiousness of market opinion is due to Alan Kirman: "Ants, Rationality and Recruitment," mimeo, European University Institute; and "Epidemics of Opinion and Speculative Bubbles in Financial Markets," in *Money and Financial Markets*, M. Taylor, ed., London, Macmillan, 1991.

19. For the reasons given earlier, deductive rationality is not defined in this market. There may be smart and foolish traders, successful and unsuccessful traders, and traders using useful and inferior predictors; but in any well-defined sense, there are no "rational" and "irrational" traders.

20. This co-evolutionary setup is reminiscent of those in R. Nelson and S. Winter models (in *An Evolutionary Theory of Economic Change*, 1982, Harvard: Ballknop) except that here hypotheses rather than firms are co-evolving, and these do not die or mutate but instead are shunted in and out of the active population depending on their recent accuracy.

21. For example, sophisticated statistical analysis of the significance of a rule's predictions can not be perfectly carried out. The agent can not define the type of process generating the prices; and likely it is non-stationary.

22. For evidence see William Brock, Josef Lakonishok, and Blake Le Baron, 1991, "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." Santa Fe Institute paper 91-01-006.

23. For a very interesting computer experiment along these lines, exploring the co-evolution of strategies in the Iterated Prisoner's Dilemma, see Kristian Lindgren, "Evolutionary Phenomena in Simple Dynamics," in *Artificial Life II*, Santa Fe Institute Studies in the Sciences of Complexity, Vol X, C. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen, (eds.) Addison -Wesley, 1991. Lindgren finds among other things: coevolution of mutually supporting strategies, periods of stasis, large extinctions, and evolutionarily stable strategies.

10. The notion of transfer goes back to Edward L. Thorndike's *Educational Psychology*, in 1903. New York: Lemke and Buechner. For a variety of modern discussions on transfer see Gordon H Bower and Ernest R. Hilgard, *Theories of Learning*, Englewood Cliffs: Prentice Hall, 1981.
11. For evidence and discussion see *Thought and Choice in Chess*, Adriann De Groot, 1965, in the series *Psychological Studies*, 4, Mouton & Co., The Hague, Paris.
12. Julian Feldman, "Computer Simulation of Cognitive Processes," in Harold Borko (ed.), *Computer Applications in the Behavioral Sciences*, Prentice Hall, 1962. See also Feldman's (unpublished) doctoral dissertation, "An Analysis of Predictive Behavior in an Two-Choice Situation," Carnegie Institute of Technology, 1959. Feldman used the binary symbols "+," and "-"/" which I have here translated into "1" and "0." I thank Kenneth Arrow for drawing this experiment to my attention.
13. This is an actual rainfall series for the years 1949-1958 for the Forrest River Mission weather station in Western Australia. Source: *Rainfall Statistics, Australia*, p.xii. Canberra: Australian Govt. Publishing Service, 1977. The next year's rainfall turned out to be 10.09."
14. A similar argument has been made repeatedly by G.L.S. Shackle (see for example *Epistemics & Economics*, Cambridge: Cambridge University Press, 1972). And of course Frank Knight in *Risk, Uncertainty and Profit*, (London: LSE, 1921) makes the well-known distinction between risk, which can be objectively measured and *uncertainty*, which can not. See also the excellent discussion in G. Dosi and M. Egitidi, "Substantive and Procedural Uncertainty," *Journal of Evolutionary Economics*, 145-168, 1991.
15. See for example, Behzad T. Diba and Herschel I. Grossman, 1988, "The Theory of Rational Bubbles in Stock Prices," *Economic Journal*, 98, 746—754. See also Robert E. Lucas's classic "Asset Prices in an Exchange Economy," 1978, *Econometrica*, 46, 1429—1445. For a general introduction to asset pricing and financial markets, the reader might consult *The New Palgrave: Finance*, J. Eatwell, M Milgate, P. Newman (eds.) 1989, in particular the articles by Ross; Dybåg and Ross; Malkiel; and Connor.
16. *Asian Wall Street Journal*, March 7, 1992.
17. J.M. Keynes. *General Theory of Employment, Interest and Money*, London, Macmillan, 1936.
18. The so-called noise trader approach goes part way in the direction of this model. It allows the existence of heterogeneous traders: "rational," arbitraging investors, and "irrational" investors who respond to "pseudo-signals" and the advice of brokers and financial gurus. For example, see Andrei

to the deductive mode, is the solution—the placement of takeoff slots—arrived at?

The deductive solution starts by assuming that the airlines decide in a known order. Thus, given any arbitrary placement of the first 19 airlines' takeoff slots on the 24 hour clock, the 20th could choose its preferred takeoff time. But then, given any arbitrary placement of the first 18 airlines on the 24 hour clock, the 19th could also choose its preferred takeoff time, knowing how *its* decision would affect the placement of the 20th, choosing behind it. But then, given any arbitrary placement of the first 17 airlines on the 24 hour clock, the 18th could choose its preferred takeoff time, knowing how *its* decision would affect the placement of the 19th and in its turn the 20th. And so on, back to the first airline to decide. The solution can be *deduced*, dynamic-programming-style, from the givens of the problem—the potential profits of each airline, given the placement of the others.

There are several pleasing aspects to this solution. In exchange for a well-defined decision problem, it gives us back a well-defined solution. The solution method is not trivial. It invokes a logic that is clean, relentless, and consistent. The logic acts step by step on premises that are well-defined: and so the process by which a solution is arrived at itself becomes mathematics. In turn, in this deductive mode, economics itself becomes mathematics. The solution is built upon a flattering view of firms' behavior as rational, well-informed, and logical. And it rests comfortably with notions we have inherited from physics of entities following maximizing principles to the good of the whole. Above all, the solution is orderly and precise.

The solution has other pleasing properties. It is self-consistent in that if other agents behave according to the deductive solution, then I as the representative agent should behave according to the deductive solution. It is self-enforcing in that if other agents behave according to the deductive solution, then it would not be in my interest—I would be less than rational—to deviate from it. And if implemented perfectly, it confirms the deductions that went into it; it is a (perfect foresight) *rational expectations equilibrium*.

As an internal system of consistent logic, the solution here is a success. But to be an accurate picture of reality it must fulfill several requirements. It requires a correct definition of the problem on the part of each agent, which in this case stipulates knowledge not just of his own locational preferences given the placement of other firms, but the locational preferences of other firms given all possible

placements of other firms. It requires further that it is known that the problem is correctly defined by the other firms—that is, that the conditions and definition of the problem are full common knowledge.

The solution requires that each firm is “rational”—that it has the ability to see how to arrive at the optimal placement and the ability to compute it. These are not trivial requirements. Further, it requires more subtly that each firm know that the others are completely rational, that the others know the others are completely rational, that the others know the others know, and so on. Rationality itself, in other words, must be common knowledge. And failing perfect rationality, there are no instructions for computing placement of time slots—“out of equilibrium” the problem is not well-defined. We can see this if we suppose all requirements hold except that firm 3 believes that firm 17 is not deductively rational (and the other firms know of firm 3’s belief). Because firm 3 does not believe in firm 17’s ability to compute and arrive at the prescribed solution, it modifies its positioning from the prescribed one. But then other firms, even if they know firm 17 is deductively logical, must modify *their* behavior to reflect firm 3’s modification. This in turn causes firms to further need to modify. The deductive solution begins to unravel as more and more firms modify *their* behavior to adjust to departures from the deductive prescription. But there is no guide to how these modifications might take place; “non-deductively rational” decision behavior is not specified in this problem and the solution breaks down.

The solution also requires uniqueness of each optimal placement as the firms work backward in the deductive logic. If there were multiple solutions so that some firms down the line were indifferent between two or more time placements, their resolutions of these ties may affect the optimal placement of earlier firms. The problem, without additional assumptions, again becomes ill-defined.<sup>4</sup>

Thus the solution constructed under the deductive metaphor is indeed beautiful, orderly, tidy. But it requires: (a) full knowledge of the problem; (b) perfect ability to compute the solution; (c) a unique solution; and (d) common knowledge that other agents are operating under (a) and (b). Thus the solution is highly brittle. If any link in this network of requirements fails, it buckles and collapses. Of the deductive metaphor in action here we might rightly say, *c’est magnifique, mais ce ne pas la guerre*.

Of course, if some of these requirements are not fulfilled in problems at this

## Notes

1. Webster’s Third New International Dictionary (1976, unabridged) defines deduction as “deriving by logical process, or drawing a conclusion necessarily from given premises, or inference from evidence in which the conclusion follows necessarily from the premises.” Induction on the other hand is “reasoning from a part to a whole, from particulars to generals, or from the individual to the universal.”
2. For a thoughtful discussion of the evidence on whether human agents conform to deductive logic in sequential choice problems, see John Rust’s “Do People behave according to Bellman’s Principle of Optimality?” Hoover Institution Working Paper, E-92-10, 1992.
3. The example is based on Edward Prescott and Michael Visscher’s influential and well-argued paper “Sequential Location among Firms with Foresight,” *Bell Journal of Economics*, 378–393.
4. We might further allow that the firms could differentially price their products; this would not tax the theory greatly, but would further diminish any chances it would correspond to the real world.
5. To use a more old-fashioned term, what I am asking here in a very particular way is how we might deal with “bounded rationality” in the economy.
6. Erwin Schrödinger, 1956, *Mind and Matter*, reprinted together with *What Is Life*, Cambridge University Press, 1967.
7. These ideas go back to the work of Bartlett and of Piaget in the 1930s. For modern discussions see David Rumelhart, “Schemata: the Building Blocks of Cognition,” in R. Spiro, B. Bruce, and W. Brewer (eds.), *Theoretical Issues in Reading Comprehension*. Hillsdale, N.J.: Lawrence Erlbaum; and R. Schank and R.P. Abelson, *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, N.J.: Lawrence Erlbaum.
8. See Section 2.1 and Appendix 12A in *Induction* by John H. Holland, Keith J. Holyoak, Richard E. Nisbett and Paul R. Thagard, MIT Press, 1986.
9. On analogy, see Keith J. Holyoak, 1984, “Analogical Thinking and Human Intelligence, in R.J. Sternberg (ed.), *Advances in the Psychology of Human Intelligence*, vol. 2. Hillsdale, New Jersey: Lawrence Erlbaum.

generic algorithm. We find that early in the experiment the price settles to random noise about fundamental value. But after some time, mutually reinforcing trend-following or technical-analysis-like rules begin to appear in the predictor population. These begin perhaps as simple extrapolation predictors, but because they validate each other they are able to gain a footing in the population of active predictors. Eventually a slowly changing "ecology" of hypothesis-predictors becomes established, with self-reinforcing technical trading rules very much part of the system.<sup>23</sup>

We find no evidence that market behavior ever settles down; the population of predictors continually coevolves. One way to test this is to take agents out of the system and inject them in again later on. If market behavior is stationary they should be able to do as well in the future as they are doing today. But we find that when we "freeze" a successful agent's predictors early on and inject the agent into the system much later, the formerly successful agent is now a dinosaur. His predictions are unadapted and perform poorly. The system has changed. From our vantage point looking in, the market—the "only game in town" on our computer—looks much the same. But internally it coevolves and changes and transforms. It never settles.

### Conclusion

Economists have long been uneasy with the assumption of perfect, deductive rationality in decision contexts that are complicated and potentially ill-defined. Yet it has not been clear what to put in its place. From the reasoning given above, I believe that what humans actually use in these contexts is *inductive* reasoning: they recognize patterns; construct representations and internal models based on these; use these as working hypotheses, possibly carrying out a good deal of deduction based on them; and strengthen or replace them as they receive feedback from their environment. If this is true, learning and adaptation are no longer addenda to the core theory; they become central to it in problems of high complexity.

I have also argued that inductive reasoning is perfectly amenable to analysis. The model of Feldman's experiment shows that we can study analytically the dynamics of pattern recognition, hypothesis formation and refutation. Finally, I believe that inductive reasoning goes quite some way toward explaining supposedly "anomalous" behavior in financial markets.

intermediate level of problem complexity, we may not want to toss away the deductive solution. Agents might use some behavioral process not based on deductive logic that brings them to the deductive solution. They would then indeed behave "as if" they were fulfilling the deductive solution. But the question of whether actual behavior, in some unconscious way, by some unforeseen process, using some unstated reasoning, mimics purely logical behavior is one to be settled empirically from problem to problem, not assumed.

Let us move farther up the scale of complexity and confront the deductive metaphor with the game of Chess. In Chess the rules, preferences of the players (to win, lose, or draw), and starting positions are known and certain. It is a well-defined two-person zero-sum game; therefore it possesses a solution in minimax strategies. This solution could in principle, much as before, be arrived at by deducing backward from all possible end-games to the opening move. But here we are well beyond being able to fulfill the requirement that the player can compute the solution. Because of the combinatorial complexity of the game (it has been reckoned that there are about  $10^{120}$  possible moves), no player, no analyst, and no computer to date has come anywhere remotely close to "solving" Chess. In the deductive mode of analysis then, there is a theory of chess. But the theory is no longer an operational one. It merely an existence proof together with a design for a solution algorithm, both ghostly objects with no operational meaning in terms of delivering hard moves or strategies. And of course it is no guide to real, human Chess play. On our scale of problem complexity, Chess appears to be well beyond the ability of humans to fulfill the requirements of the deductive metaphor.

### The Complexity Boundary

Where does this leave us? I believe we can say that the deductive metaphor works well for problems in economics where the givens of the problem are known and available, and the computations required are not demanding. But there appears to be what I will call a *problem complexity boundary* beyond which arriving at the deductive solution and calculating it are unlikely or impossible for human agents; and beyond which other agents cannot be relied upon to carry out their part of the deducing process. Beyond this boundary where the requirements fail, there are often no instructions on how to proceed logically; so that away from the prescribed solution—"out of equilibrium"—the problem becomes ill-defined and the deductive



metaphor can not operate at all. Exactly where this boundary lies is of course fuzzy. In standard games, it would be somewhere in between Tic-tac-toe and Checkers. Humans *can* deduce, and they *can* calculate. But the level at which they can do this reliably and accurately, I believe, is surprisingly modest.

It may seem that beyond the complexity boundary, where the requirements of logic fail to be met and problems become ill-defined, decision making becomes impossible. Yet this is not so in practice. Humans make decisions quite comfortably in situations that are not well-defined. In fact when they do this they are barely aware that the problem they are dealing with is not well-defined. Notice I am *not* saying that whereas in simple problems humans act optimally in complicated ones they act less than optimally. Where a problem is ill-defined, optimizing behavior is also ill-defined; and so "acting less than optimally" is not defined. What I am saying is that beyond the problem complexity boundary either human rationality cannot be relied upon to meet the demands of the deductive metaphor or rational behavior from a logical point of view is not well-defined. Yet humans do make decisions beyond this boundary: ones that carry consequences of greater or less effect.

This raises two problems in economics: Beyond this problem complexity boundary how *do* human agents arrive at decisions? And how might we carry out proper theorizing in economics that recognizes the realities of human decision making beyond the problem-complexity barrier? It might be useful here to make a distinction between *deductive rigor*, exact logical theorizing based on given, but possibly unrealistic assumptions; and *behavioral rigor*, the construction of economic theory based on the precise, observed actualities of human behavior. What I am asking then is how we might construct models and theories in economics that are behaviorally rigorous in this sense of accurately reflecting *real* human decision-making in problems of complication—problems that might even as far as deductive logic is concerned be ill-defined.<sup>5</sup>

### How We Know

Some time ago, Herbert Simon introduced the idea of *satisficing* as a way to describe less-than-logical behavior in certain management decision contexts. It is tempting to propose that in our case, beyond the problem complexity boundary, humans satisfice—they simply attempt in a vague way to achieve some reachable outcome under circumstances that are not amenable to full rational action. But in

likely be mixed.

Fourth, under some circumstances we might indeed be able to ascribe psychological moods such as "depressedness" or "nervousness" or "jitteriness" to such a market. For example, assume price has indeed recently gone high and that many traders are monitoring both trend-following rules and reversion-to-fundamental-value rules. Trend-following predicts continued high profits from buying. Reversion predicts the possibility of a downturn that could validate other sell-predictors and trigger a crash—a major loss. On the basis of past behavior, both predictors could show similar forecasting accuracy. Our model of the market of course has no emotions. But it would not take too much imagination here to see that in this case in individual agents, something that could be described as "fear" and "greed" might be more or less balanced at high, disquieting levels; so that the market could certainly be described as "nervous."

I have used a highly simplified picture of the market to show that a more realistically based theory of the market might explain features the deductive theory can not. In a more elaborate model we might want to allow for sources of actual randomness in underlying price introduced by quantum-like effects in the bid-and-ask process or by random liquidity needs. We would want to allow for non-unit trades. And we would want to allow agents to form new predictors and discard old ones from their set of monitored ones. This would further add to the "openness" of the dynamics of price behavior—the tendency of the market to be always changing, always exploring new territory.

Of course the points I have made above about the likely emergence of technical trading rules and a mixed population of active predictors are conjectures. They may not be difficult to settle by analyzing what classes of predictors are mutually supportable or evolutionarily stable. The broader questions of what might emerge if the market started from scratch and of whether market behavior might never settle down may not be easily analyzed. In fact, full examination of the implications of the type of model I have just outlined may require controlled experimentation on the computer rather than standard mathematical analysis.

John Holland, Richard Palmer, Paul Taylor and I have in fact set up a computer experiment at the Santa Fe Institute along the lines of this model, with 100 computerized agents each using 60 predictors. Instead of a fixed set of hypotheses, we allow new hypotheses-predictors to be generated from time to time using Holland's

his constant motivation to trade is well founded or irrational. While is easy to see a subjective motive for investing here; it is not possible for the agent to settle the desirability of trading by any logical test.<sup>21</sup>

Second, *all* trading rules here forecast on the basis of patterns in past prices; therefore if some emerged that produced sustained, genuine profits they would be examples of profitable technical trading in action. Technical trading would then be a valid outcome in this market. Can this happen? Likely it can with trend-following rules. These have the property that may be self-reinforcing or self-validating in the population of active predictors: If a subset of investors is using trend-following predictors, and the market price has risen recently, they will buy; and this will enhance the probability that the price will indeed go up; thereby enhancing the accuracy of such predictors. (In the real market, moving average rules predict the direction of prices correctly about 52% of the time.<sup>22</sup>) Such self-reinforcing rules would therefore tend to linger in the active set and become a mutually supporting or symbiotic subpopulation in the “ecology” of active predictors. Of course, if such rules profited over the long term, they would need a substrate of predictors that lost on average. This might automatically be provided by agents who start with and monitor less successful predictors. (Or it might be provided in the real world by agents buying or selling in the market for liquidity or hedging reasons or by a constant flow of new entrants to the market.) Of course if the market started with no such rules active, they may face a threshold problem in gaining a footing in that their accuracy depends on the number of other such predictors active. But I don’t believe this is a major problem. Trend following is extrapolation; and extrapolation is often validated by chance, so that “technical” predictors could easily emerge as a robust and healthy subpopulation.

Third, the population of active predictors would likely see a balance between simple, robust rules or predictors, and elaborately thought-out, conditional ones. Elaborate predictors based upon highly special conditions (“the market will go up if conditions A, B, C or V, but *not* R, G, and S hold”) may show great accuracy when invoked, but being conditional they will pertain only to very particular price circumstances that occur rarely in the price history. There will few recurrences of such circumstances to validate them, and so their high correctness will be offset by their rarer chance to be validated. Simple predictors may tend to show the opposite: lower accuracy but frequent validation. The predictor population that emerges will

actual fact, this does not appear to be true. Humans use characteristic and sophisticated methods of reasoning in contexts that are complicated or ill-defined. To get an idea of how these work we need to go back and think about some basics.

What we are interested in is how agents in the economy make decisions in complicated problems; how they act upon the information available to them; in other words, how they *use* what they *know*. But “knowing” and “learning” —the updating as it were of knowing—are not straightforward. They are subtle affairs. Cognition is not perfectly understood by any scientific discipline. But what is understood is crucial for us to think about, and it is worth reviewing here.

Superficially, because we tend to think in fixed, laid-down categories, we tend to think of knowledge as data or facts, organized into these fixed categories, the information set. In this superficial model of knowing, an agent might notice that the Dow Index is running at 2955. He stores this fact in the Dow category in his information set; therefore he *knows* this. Tomorrow, he finds the Dow is at 2978. He revises his information set. He updates, *ergo* he has learned. In this model, “knowing” means being cognizant of data that are boxed within a category. Learning means revising these data: updating values, quantities, and parameters according to new information from outside and modifying behavior as a result.

This is fine for many purposes as an approximation. But we need to go deeper. The world—certainly that part of it that has to do with economic decision-making—does not come to us pre-packaged in our minds in the form of fixed categories or conceptual models that we can use to forecast, analyze, and act upon. Our knowledge, together with the concepts, notions and models we use to represent it, is *created*, put together over time by us and by others in society as a whole. As Schrödinger put it in *Mind and Matter*:<sup>6</sup>

The world is a *construct* of our sensations, perceptions, memories. It is convenient to regard it as existing objectively on its own. But it certainly does not become manifest by its mere existence. Its becoming manifest is conditional on very special goings-on in very special parts of this very world, namely on certain events that happen in a brain.

What we do then in acquiring knowledge, is not just notice, say, that the Dow

Index is at 2955, or that inflation is running at 4.3%. We need first to construct the concept of a "Dow Index" or of "inflation" by organizing certain of our perceptions into categories, creating conceptual labels for these categories, and acting upon these categories if they are both useful and agreed upon by others. We appear to be good at doing this even at the finest levels, good at creating categories that summarize useful aspects of our perceptions and that allow minute distinctions. In a sense that is what we use language itself for. But we do this at higher levels too. We construct conceptual frameworks (schemas and frames in the psychology literature), "chunking" our representations into higher and more general structures in which information and useful observations can be stored and used to guide action. Thus we constantly build and rebuild *representations* of our world, from time to time constructing new ones as we find them useful and discarding old ones as they become obsolete.<sup>7</sup>

Acquiring knowledge in the form of new sensory perceptions may well indeed be a matter of dropping new data into well-defined categories or representations. But just as often, it can involve a search over our representations to see which one comes closest—which one can organize the new perception, make sense of it, house it for later use. We see this most clearly in research, where the difficult part of the work is to find a representation—a conceptual framework—that helps us make sense of the problem we are grappling with. We are hardly aware we do this. But occasionally, looking at a framework used by someone in the past or taken for granted in another culture—Marxism, say—the representation becomes all too obvious. It looks strange and quaint. Of course, this searching over the space of representations or concepts happens on a less grand scale, subtly and constantly, in all the tiny activities of life. Acquiring knowledge is then not simply accruing "facts." A major part of it lies in finding categories—often approximate ones—in which to organize and store perceptions. It follows that if two people have been exposed to different experiences in the past, with resulting differences in the stock of conceptual representations they have formed, they may act upon the same data differently. The economic world, given this sort of learning, is path-dependent.

This is only one aspect of knowing and learning. Given such representations, we indeed appear to be moderately good at looking ahead and anticipating consequences; we reason and deduce and analyze in a way economics is comfortable with, coming to new conclusions when faced with new data—learning,

continuum in  $T$ -space; and so the function  $g$  is likely to be extremely complicated. All each agent can do is form models of his own from estimations over past prices—a highly limited data set.

Thus, even in this simple deterministic model, price movements are complicated and unpredictable. To the individual agents they may appear to a large extent random, with at best some hope of correct prediction some of the time. We could say here that the market is at an equilibrium—one that consists of an extremely complicated, dynamic trajectory. But there is no reason that this trajectory should cycle or show any simple behavior. In fact there is no reason to suppose that the price trajectory will ever repeat. Except in the most trivial of cases, it may appear as if it is always "exploring" new territory.

Notice that the market here is "rational" in the wide sense that investors are free to use the ultimate in deductive and inductive methods in trying to determine how price dynamics work.<sup>19</sup> But the price dynamics themselves depend on their efforts to do this. Predictors or hypotheses are acted upon as long as they are accurate and profitable; if not they are temporarily dropped until perhaps by chance they become accurate again. To borrow a metaphor from John Holland, we have an "ecology"—not of species—but of hypotheses or predictors. The fitness or profitability of using a given predictor depends on the composition of the population of other predictors. Predictors are therefore co-dependent and the set of active predictors co-evolves. Some classes of predictors may be competitors, others may be co-adaptive—mutually supportive. And the overall composition of the population of active predictors will change over time.<sup>20</sup>

I want to make several points with this inductively based market model.

First, it gives an interesting answer to why investors should want to trade in the market at all. In standard models, traders' expectations of profits are always zero; therefore they have no motivation to speculate. In this model agents invest on the basis of predictors that performed well on past data and that appear to be good bets for future profits. Therefore they have a constant, subjective incentive to invest. Of course such motivation would be ruled out if in actuality it were spurious, based on a pattern-predictivity that did not in fact exist. But as in Feldman's experiment there is no objective way for an agent to settle this. There is no way to settle whether a predictor's accuracy in the past is "real" or mere "chance"—the result of fitting "noise" in the past data better than other methods. Therefore he can not tell whether

There is no reason to suppose that an agent monitors only one predictor in the market. Each predictor is a hypothesis and the agent may want to hold several in mind simultaneously and monitor each for their current predictive accuracy. I will assume that each agent:

(f) monitors at most  $M$  predictors; and updates the forecasting accuracy of these when the new price is revealed. (In general, we might allow agents to change their set of monitored predictors; but here will suppose the set for each is fixed)

(g) forms a composite predictor from the most accurate subset  $M'$  of his  $M$  predictors, (say, the best three, or best one, or all predictors) by weighing these according to their accuracy in the standard way

(h) buys, sells, or holds, as this composite predictor directs.

We are now back in a world quite similar to the one of the Feldman experiment.

What are the dynamics of prices in this market? Define the overall set of investors' hypotheses or predictors as

$$H = \{H_1^1, \dots, H_M^1; H_1^2, \dots, H_M^2; \dots; H_1^K, \dots, H_M^K\}.$$

where  $H_{ij}$  is the  $i$ th predictor that the  $j$ th agent chooses to monitor.

Now, each agent monitors his predictors' accuracies over the last  $N$  periods, using data  $W$  periods back to forecast. Thus any *price history*  $P_t(\cdot) = \{P_t, P_{t-1}, \dots, P_{t-T}\}$ , where  $T = W + N$  assigns unambiguous accuracy values to each hypothesis-predictor. This in turn determines the set of "active" predictors—the ones that are used—and each agent's buying or selling behavior in the market. Taking all agents collectively, this determines excess demand for the stock, and in turn the price adjustment. Thus, given the set of monitored hypotheses of all agents, there is a well-defined mapping  $g_H$  from the current price history into the next period's price:

$$P_{t+1} = g_H(P_t, P_{t-1}, \dots, P_{t-T})$$

In this model of market behavior, given the set of pattern-hypothesis-predictors in use, price moves through the  $T$ -dimensional space  $R_T^+$  deterministically. Why then can not each agent reconstruct this mapping, and predict correctly at each time? Well, there are many agents, acting on many hypotheses, defined over a price

in other words. Somewhat more subtly, we appear to form *internal models* or *mental models* of the situations we are dealing with, models or representations that we can use to play out and "rerun" scenarios in our minds. (This is a mechanism that John Holland and his associates have repeatedly stressed.<sup>8</sup>) By mentally simulating what-if questions in this way, rehearsing them and thinking about the imagined outcomes, not necessarily logically, we come to new conclusions. Psychologists call this latent learning.

We appear to be superb at seeing patterns, recognizing patterns, and matching patterns. It is by mentally processing patterns that I "recognize" a face as a friend, or a certain painting I have never seen before as a Vermeer. And we appear to be able to hold patterns in our minds temporarily as hypotheses—"I think possibly it is a Vermeer"—rejecting or building confidence in the hypothesis as more evidence becomes available. It is by something akin to pattern recognition we appear to be good at seeing analogies and structural correspondences.<sup>9</sup>

We appear to do much of our decision making by finding problems we have dealt with in the past that are similar to the one we now dealing with. Often these are problems whose resolutions we have deemed satisfactory, in some way. In the language of psychology, we *transfer* experience from one problem to another.<sup>10</sup>

We appear to have a great capacity to generalize: to reason from the particular to the universal. We do this not by some primal ability but by using many of the capacities I have outlined above—in particular those for analogy, transfer and pattern recognition. If I have sufficient previous experience, I need see only a corner of a painting to declare it a Vermeer. I see the part, recognize the pattern, and fill in the blanks, so to speak.

We can see some of these modes of knowing and using knowledge if we think of how actual human behavior would operate in airline-time-slot model and in Chess. How would airlines actually decide? Likely they would not deduce backward, but would formulate hypotheses about how those that follow them in the booking of slots might react given the decisions they might make. They would seek out new knowledge and update these hypotheses before booking began as they gained more information on their competitors' needs and positions. They would act on past experience with their competitors in such circumstances and transfer it to the present situation. They might make internal models, and partially deduce forward on the basis of these. They would allow for many outcomes and seek decisions

robust against a range of these. Similarly in Chess, competitors would watch each other, formulate hypotheses about each others' motives and intended strategies; and update these as the play unfolded. They would categorize play: "He's using a Caro-Kann defense." They would search over previous games to form analogies: "This looks a bit like the isolated pawn problem in the 1936 Botvinnik-Vidmar Nottingham game." They would also search over representations to organize the data received: "My pawn formation is getting stronger." They may discern patterns in their opponent's play in past games, and exploit these to advantage. They would also of course carry out local deductions—analyzing the possible implications of moves several moves deep. But overall, they would operate in a world of pattern, hypothesis, analogy, and representation.<sup>11</sup>

To summarize, we certainly reason, analyze and deduce. But to do so at all we must first have constructed categories, representations, and models. Further, in ill-defined problems, or in complicated ones like Chess, pure deductive reasoning either is not possible or is beyond our abilities and it is applied to only part of the decision process. And so we reason by other, different means. We formulate internal models; we search for and use analogies; we recognize patterns; we transfer experience from other, similar situations. We use these methods to fill the gaps in our understanding so to speak, to allow us to infer from part-information to the whole, to extrapolate from the particular to the general. In other words, in "knowing" and "learning"—in gathering and gaining understanding—we operate heavily in the *inductive mode*. In any particular problem, of course, we combine the deduction we can carry out with a considerable amount of induction.

Granted these aspects of cognition and the formation of understanding, it is tempting to push them aside as messy, subtle, not easily mathematized and therefore not amenable to theorizing in economics—in a word, non-scientific. But notice that science itself—the process of understanding and coming to terms with the natural universe—can not proceed deductively. Science proceeds by formulating representations, categories, frameworks; by creating analogies; by noticing patterns; by formulating hypotheses, strengthening and refuting them; and by transferring understanding from one discipline to another. In the broadest sense, science—our most rational of endeavors—is largely inductive.

The inductive world has a close connection with learning and adaptation. Deduction—the deriving of a necessary conclusion from given premises by a logi-

discounted value of the dividend stream)

(c) Agents can buy or sell one unit of stock at each time, or can "hold"—do nothing

(d) Given buying and selling orders, all trades in this market are fulfilled by a "specialist" who declares a new price for the next period, and this new price is a deterministic function of excess demand and today's price

(e) Agents can not directly communicate their buying or selling intentions. In fact, the only observable variables in this market are current and past prices in a moving window that goes  $W$  periods back, that is  $\{p_t: t-W < \tau < t\}$

Assumptions (b) and (c) are hardly realistic; but we can complicate them later if needed.

Now, each agent's basic problem is to profit in the next period, and the only way he can do this is by predicting the direction of the market. But given these assumptions, as I said earlier, constructing a predictive model—even a stochastic one—is not a well-defined problem. There is no deductive method by which to proceed. Like our Chess players, all our agents can use is the available public data (here the price history) to form their own subjective, private conjectures about the predictions and intentions and subjective, private conjectures of their fellow investors given the current market situation. Given these conjectures about what is currently moving the market, any prediction that results from them, whether derived empirically or analytically, is predicated on the currently available price data. Thus any agent's predictive model is a hypothesis/predictor as defined earlier, a mapping or function from the pattern formed by the current price data set into a forecast of the next period's price. But this time each such function is from a continuum into a continuum; and there are infinitely many of these to choose from.

Of course, these hypotheses or predictors or models-of-the-market may be chosen on the basis of very elaborate reasoning about the choices, intentions, intelligence, and predictors of other investors. Or they may be adhered to as simple rules of thumb: "If the average of the last 5 days' price is higher than the average of the last 50 days, the price will go up. Conversely the price will go down." "If the price goes above six times fundamental level, a reversion downward will follow." We could call this last type of predictor a *conditional* predictor. It is only invoked and monitored on a subspace of the price-history space, in this example that where current price is greater than six times fundamental price.

follow (1). They may not be rational. They may not understand the basic arbitrage of the market. They may be acting on strange or crude or superior information. Then (1) would not be a good guide as to how price is formed.

Even if I know that other investors *are* perfectly rational and possess the same information as I do, and that these facts are common knowledge, notice that (1) implicitly invokes the common-expectations assumption that there is an objective method—a unique, implicitly agreed, prescriptive way—to use the information  $I_t$  to form the expectation  $E[P_{t+1} + d_{t+1} | I_t]$ . But this is patently unreasonable. The available shared information  $I_t$  typically consists of past prices, trading volumes, economic indicators, and the like. These are just sequences of numbers; and once again there may be many different and perfectly defensible statistical ways based on many different assumptions and error criteria to use them to estimate or forecast tomorrow's price. I might *assume* that everyone uses the same method to form expectations. But I have no objective way to figure what this method is. Nor have other investors.

And so, exactly as previously, the deductive solution begins to break down. Not knowing how other investors arrive at their expectations, I cannot form mine in a well-defined way. And so I must make assumptions—subjective ones—about how other investors form expectations and behave. I must begin to try to “psyche out” the market, in particular other investors’ beliefs and expectations. I must try, in Keynes’s words, to figure out “what average opinion expects average opinion to be.”<sup>17</sup> But this lands us immediately in a world of subjective beliefs, and beliefs about subjective beliefs. The deductive behavior of (1) no longer has meaning, and we are propelled into a world of induction.

#### *A Stock Market Model*

How would prices behave if we were to suppose that beliefs were formed in an inductive way? Let me briefly sketch out a model that is highly simplified, but will nevertheless serve to illustrate several points I want to make.<sup>18</sup>

I will assume:

- (a) A single stock is traded by  $K$  investors or economic agents, each of which is interested in maximizing the value of his portfolio next period
- (b) The stock pays a deterministic dividend  $d$  every period and there is an outside interest rate  $r$ . (Thus the “fundamental value” of the stock is  $d/r$ , the

cal process—is a largely static affair. Unless the problem changes, the solution calls for no updating. Induction, on the other hand—reasoning from part knowledge to a whole solution by using subjective assumptions or hypotheses—can always be improved by updating the assumptions or hypotheses. Induction calls for a dynamic approach. It calls for as much feedback on and refinement of the solution and the ideas invoked to arrive at it as it can get. In other words, it calls for learning.

At this point let me say a word about the term “rational.” We use “rational” in economics to describe only behavior that would conform to the logical, deductive decision mode outlined above. If we use “rational” in its wider sense to connote intelligent, sensible, savvy, or reasonable behavior, then it coincides with our economists’ perfectly-rational-deductive mode of behavior only in fairly simple problems where the deductive requirements are fulfilled. Beyond the complexity boundary where these requirements fail and problems become ill-defined, real rational behavior—intelligent behavior—is forced to use subjective assumptions and temporary hypotheses to reason from part information in the problem to a whole solution. It requires that inductive reasoning take over. Of course, both deductive and inductive modes are normally present, but the deductive one is much more limited in the complexity it can cope with, and when it becomes overloaded it recedes in importance to give way to the inductive mode.

Coming to grips with this wider definition of “rationality” in economic decision making therefore means coming to grips with inductive behavior in complicated economic decision making. We therefore need to look at inductive reasoning more closely, and to explore the style of theorizing that might capture it works.

#### **Hypotheses and Refutations**

I have now given a lengthy answer to the question: how do human decision makers proceed in the context of complicated or ill-defined problems? In most problems of complication, they proceed in a fairly predictable and simple sequence: they look for patterns; construct representations and internal models based on these; use these as working hypotheses; carry out deductions based on these hypotheses; and strengthen or replace these models or hypotheses as they receive feedback from their environment. Interestingly, humans appear to be able to hold in mind several internal models or working hypotheses at once, making quick switches among the ones they act upon as fresh data validate or negate them.



In the rest of this talk I want to show this inductive process in action using two example problems as illustration.

The first example derives from a classic experiment carried out by the psychologist Julian Feldman in 1959 that shows pattern recognition, hypothesis formation and inductive behavior at work in a particularly pristine form.<sup>12</sup> Feldman had human subjects predict which of two events, the appearance of a "1" symbol or the appearance of a "0" symbol, would occur next in a sequence of two hundred trials where the experimenter could control whether 1 or 0 appeared next. Subjects were not given any information on the process the experimenter was using to generate the next symbol. At each step, the subject announced his prediction and was then told which event occurred. He was asked to "think aloud" and give the reason for his choice at each step.

Feldman found that each subject was quick to spot patterns in the sequence of 1's and 0's and to form hypotheses on the process generating the sequence. Subjects would form a hypothesis—"You're following with a 0 and two 1's," or "You've begun a progression of two 1's and two 0's"—for some time, and stick to it as long as it predicted well. They allowed for exceptions fairly liberally: "You just gave me the 1 to throw me off." But if the pattern-hypothesis performed badly over several predictions they would change or drop it in favor of a different one.

The interesting thing is that the sequence of 1's and 0's in each case that Feldman used was perfectly *random*. Yet each subject could "see" patterns to act upon, albeit different ones as the experiment progressed.

In a way it seems ridiculous—overdone somehow—that our minds see patterns where none exist. But actually, this is not absurd. In most contexts we face as human beings, patterns *do* exist and the ability to discern them, conceptualize them, hypothesize about them, and act on them, has considerable evolutionary value. Evolution does not look for an optimal response, and it does not set much store by logic. It merely brings to the fore behaviors that work in our environment better than others. Pattern recognition is one of them.

#### *Temporarily Fulfilled Expectations*

Feldman's experiment appears to be something of a curiosity. But I believe it is important in showing how humans continually formulate and discard hypotheses to act upon, in situations where new data are constantly arriving. And so I want to

information that could be used to forecast tomorrow's; therefore methods such as technical trading rules that use past price patterns to forecast tomorrow's price cannot in the long-run make money; and therefore temporary surges and crashes are simply adjustments to new information.

Traders, on the other hand, tend to believe that to some degree the market is forecastable from past prices; most of them take technical trading rules or "chartism" seriously; and they see surges and crashes often as the result of herd behavior rather than adjustments to new information. They talk about the "mood of the market;" they worry about its "nervousness" or feel buoyed by its "confidence." Traders, in fact, see the market almost as if it had a personality of its own, a complex psychology that can be read and understood and profited from as acquaintance ship with it deepens. "The dollar," writes *The Asian Wall Street Journal* of the previous day's foreign exchange trading, "rallied on good news and shrugged off bad news."<sup>16</sup>

From the point of view of economic theory, the traders' view seems to be unduly anthropomorphic, a reading into the market of an "aliveness," a personality, that simply does not exist. It seems to be a nonsense. But are these notions so absurd? From the viewpoint of inductive behavior I believe they are not. Let us examine this question by looking first at whether the deductive theory of the market is reasonable.

The deductive solution comes in several flavors and degrees of rigor. Models may be complicated. But at the heart of most reasoning about the market pricing of an asset is a behavioral equation of the type:

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

We can interpret this as an arbitrage condition. Today's price,  $p_t$ , adjusts so that it reflects the expectation of tomorrow's price plus dividend,  $p_{t+1} + d_{t+1}$ , given discounting at rate  $\beta$  and today's information  $I_t$ . If this were not so, there would be an opportunity to make a profit and investors would be less than rational if they forewent this opportunity. But in pursuing it, they would bid today's price into line with their expectation of tomorrow's gains. Notice this assumes all investors are alike, have the same information  $I_t$ , and form expectations from  $I_t$  in the same way.

This seems reasonable and persuasive. But now suppose that  $I_t$ , as a reference investor, suspect that some significant portion of the investor pool does not exactly

Myself, I certainly do not believe there is a well-defined answer, given this information. Even if we had a longer data set going back 200 years. I know of no prescriptive logical forecasting procedure would be more defensible than many others. The common-expectations assumption that all agents share the same knowledge of the process generating next period's variables and use it in the same way is only valid when agents have knowledge of the actual mechanisms generating the series. This is reasonable maybe in well-defined cases like roulette or coin-tossing but not in the case of rainfall or in many economic situations. Absent such a mechanism, we are forced to admit that the formation of expectations is not well-defined: perfectly intelligent agents might use very different methods in extrapolating the series. Some may posit an underlying model generating the series. Some may look for outside data on other variables and try to incorporate these. Still others may use the given data alone, and fit a time-series, or do a linear extrapolation, or fit a non-linear model like a neural net. Even those using the same forecasting method may have different error criteria.

Thus, usually in economic problems it is *not* reasonable to suppose that the probabilistic process generating variables of interest is known and shared by all—that the common-expectations assumption is valid.<sup>14</sup> On the contrary, it is usually unknown, and reasonable men and women may use very different subjective methods—different pattern-recognition devices, different hypothesis-predictors—in forming expectations. This may seem a finicky point; but it is not. A high proportion of problems in economics involve expectations and the standard analysis of most of these invokes the common-expectations assumption to arrive at a deductive theory. Proceeding without it can lead to very different results, as I will now show.

### Financial Markets

The striking thing about financial markets is that the standard theory is completely at odds with practitioners' views of the market. In the standard theory of financial markets, in particular of the dynamics of stock prices, very few outcomes are possible. The price of a stock settles either to "fundamental value"—the discounted worth of its dividend stream—with perhaps some noise or perturbations, or to exponential increase driven by rising expectations.<sup>15</sup> In this theory the market is efficient in the sense that today's price discounts in all

show that this sort of behavior is not nebulous: it is perfectly amenable to analysis.

To make a model of Feldman's experiment, let us suppose the experiment continues indefinitely; and that each subject possesses what I will call "pattern-hypothesis-predictors" that he monitors at each step—rules or functions that map the immediate past pattern of digits at each time into a 1 or a 0 which serves as a forecast of the next binary symbol or digit. These rules are hypotheses in the sense that the subject holds them in his mind as conjectures on how the digits are generated. For example, the hypothesis that the sequence is alternating would be a function that "recognizes" the last observed digit as 0 or 1 and maps it into 1 or 0 respectively.

Of course many such predictors or hypotheses could be formed. For concreteness, let us suppose each subject is capable of looking back at past patterns at most four digits. Each predictor then maps all possible four-bit sequences into 1 or 0. There are 16 different four-bit sequences. Therefore there are  $2^{16}$  or 65,536 possible rules or functions—predictors—for mapping the 16 possible patterns into a 1 or 0. This is a large number. We will suppose the subject monitors or holds in mind some small, possibly very small, subset of these hypotheses-predictors,  $H_1, H_2, \dots, H_N$ , and assess their predictive accuracy at each step. (This subset is fixed through out the experiment.) Let us further suppose that the subject uses or acts upon the predictor that has been correct most often in the last ten trials to predict the next symbol in the sequence. I will call this the *active* predictor or hypothesis. (We can also suppose that he numbers his predictors, and in the case of ties he chooses as active the one with the lowest index-number.)

Let us now focus on the process of how predictors or hypotheses from the given set become active and used, or discarded back into the group that is merely passively monitored over time. Now, because we are assessing predictors on their accuracy in the last ten trials and the accuracy ten trials back requires input from the four digits preceding this, the last 14 symbols tell us unambiguously which hypothesis-predictor is the most accurate and therefore the one used. So we will be able to look at the dynamics of the predictors if we look first at the dynamics of the sequence taken 14 digits at a time. Consider then the binary sequence taken 14 digits at a time. This lives, so to speak, in the space of all fourteen-bit binary patterns. A typical point or state in this space might be 01101100001010, where the digits are arranged with the rightmost the most recent. At the next time, we lose the oldest, leftmost digit, and gain a new one on the right. So our point can transit into



11011000010100 or 11011000010101. In this way each point, or state, has two out-neighbors or successor states, the two states it can transit into. By the same token, it has two in-neighbors, the two fourteen-bit states or binary patterns that could have preceded it. There are  $2^{14}$  or 16,384 binary patterns that are 14 binary symbols long, and so there are 16,384 points in this space. We can imagine them connected by directed lines signifying each point's two in- or out-neighbors. Now, under Feldman's setup, the next binary digit is completely random; so at each point, the process forks to one of the two out-neighbors with equal probability. The state of the system therefore performs a random walk in our connected, binary 14-space.

It is now very easy to determine the dynamics by which hypotheses are taken up and discarded. All we need do is label each point in this space with the most accurate predictor at that point. Denote as  $R_i$  the set of points where  $H_i$  is the most accurate predictor. This then divides the space into regions  $R_1, R_2, \dots, R_N$  corresponding to the predictors  $H_1, H_2, \dots, H_N$ . As the process walks from point to point, it also walks through the regions corresponding to the most accurate, or active, or dominant hypotheses labeling those points.

Notice that the active hypothesis will not change every time. Suppose we are at a point  $p$  in our binary space, where predictor,  $H_3$  say, predicts the last 10 digits of the sequence the most accurately. At the successor to  $p$ ,  $H_3$  may still remain the most accurate even if its latest prediction is false. (The smaller the number of monitored predictors, the more likely it is that a point's successors are dominated by the same predictor.) Thus, as the state of the system moves farther from  $p$  it may stay in the region of dominance of  $H_3$ . Eventually, of course, errors to  $H_3$ 's predictions will build up, and the state of the system will enter the region of dominance of another predictor. In this way a human subject will linger with the most recent, accurate predictor believing possibly in its efficacy; but will drop it when it no longer predicts well, in favor of a better one. The subject will transit from one hypothesis-predictor to another, spending a random amount of time with each one before dropping it, from time to time making abrupt changes in perception as it were, and over time covering the space of all hypotheses he is monitoring. We could speak of this system as being one of *temporarily fulfilled expectations*—beliefs or hypotheses or predictors that are fulfilled, albeit here by a perfectly random process, but which give way to different beliefs or hypotheses or predictors when they cease to be fulfilled.

In this experiment of course there is no real connection between the predicted outcome and the actual one. In real life there may be. To allow for this we can suppose there is some exogenous mechanism that generates the sequence; and for some reason it arranges the probabilities of the next symbol matching the predicted one to be .55 instead of .50. That is, there is a small tendency toward self-fulfillment of each predictor or hypothesis. What will now happen? Once again, the subject will transit through the predictor regions, but the increased probability of transiting to a forecasted successor state means that sojourns in each region will be *longer* under this partial self-fulfillment. To the degree that prediction actually means something then, the expected time spent with each hypothesis will increase. By the same token of course, if predictions are self-negating, the expected sojourns with each hypothesis will decrease.

The Feldman experiment illustrates inductive reasoning at work. It illustrates the temporary acceptance of, and discarding of plausible working hypotheses that enable action. By the type of analysis I have outlined here, inductive behavior can be analyzed.

### The Common-Expectations Assumption

I want to give a second example of inductive reasoning at work, this time in a more important arena, financial markets. Here I will again be looking heavily at the formation of expectations and beliefs; but before I do this I want to comment on how this is handled in the deductive literature.

In problems involving expectations in economics, we ask that agents predict future variables  $z_{t+1}$  from their past values  $\{z_t, z_{t-1}, \dots, z_1\}$  or some other information set  $I_t$ . And in such cases, we typically start an economic analysis with something like: "Assume homogeneous agents who share full knowledge of the probabilistic process that generates  $z_{t+1}$  given information  $I_t$ ."

But where do agents get such knowledge of the probabilistic process that generates  $z_{t+1}$ ? A typical agricultural decision problem might involve predictions of next season's rainfall given knowledge of rainfall in the past. Rainfall in the last ten years might have been 8.04," 8.89," 5.69," 2.88," 5.46," 6.55," 8.44," 5.96," 10.68," 4.30" (where the series is moving forward in time). How might agents forecast next year's rain from this sequence?<sup>13</sup> The reader might take a serious look at these numbers and reflect for a moment on how to answer this.