Adaptive Behavior and Economic Theory
Author(s): Robert E. Lucas, Jr.
Source: *The Journal of Business*, Vol. 59, No. 4, Part 2: The Behavioral Foundations of Economic Theory (Oct., 1986), pp. S401-S426
Published by: The University of Chicago Press
Stable URL: https://www.jstor.org/stable/2352771
Accessed: 23-05-2019 02:22 UTC

Robert E. Lucas, Jr.
University of Chicago

Adaptive Behavior and Economic Theory*

This essay uses a series of examples to illustrate the use of rationality and adaptation in economic theory. It is argued that these hypotheses are complementary and that stability theories based on adaptive behavior may help to narrow the class of empirically interesting equilibria in certain economic models. An experiment is proposed to test this idea.

I. Introduction

The relationship between psychological and economic views of behavior, once a subject of heavy dispute, is now understood in a very similar way by practitioners of both these disciplines and of our sister social sciences. In general terms, we view or model an individual as a collection of decision rules (rules that dictate the action to be taken in given situations) and a set of preferences used to evaluate the outcomes arising from particular situation-action combinations. These decision rules are continuously under review and revision; new decision rules are tried and tested against experience, and rules that produce desirable outcomes supplant those that do not. I use the term "adaptive" to refer to this trial-and-error process through which our modes of behavior are determined.

If one is interested in modeling particular decisions in any very explicit way, it is obviously necessary to think about rather narrow aspects of an individual's entire set of decision rules: his or her personality. Experimental psychology has traditionally focused on the adaptive process by which decision rules are replaced by others. In this tradition, the influence of the subject's (or, as an economist says, the agent's) preferences

* Jacob Frenkel, Robin Hogarth, and Melvin Reder provided extensive and very helpful criticism of the version given at the conference.

(*Journal of Business*, 1986, vol. 59, no. 4, pt. 2)

0021-9398/86/5904-0022$01.50

are kept simple by choosing outcomes that are easily ordered (rewards vs. punishments), and the focus is on the way that behavior is adapted over time toward securing better outcomes.

Economics has tended to focus on situations in which the agent can be expected to "know" or to have learned the consequences of different actions so that his observed choices reveal stable features of his underlying preferences. We use economic theory to calculate how certain variations in the situation are predicted to affect behavior, but these calculations obviously do not reflect or usefully model the adaptive process by which subjects have themselves arrived at the decision rules they use. Technically, I think of economics as studying decision rules that are steady states of some adaptive process, decision rules that are found to work over a range of situations and hence are no longer revised appreciably as more experience accumulates.

From this point of view, the question whether people are in general "rational" or "adaptive" does not seem to me worth arguing over. Which of these answers is most useful will depend on the situations in which we are trying to predict behavior and on the experiences the people in question have had with such situations. It would be useful, though, if we could say something in a general way about the characteristics of social science prediction problems where models emphasizing adaptive aspects of behavior are likely to be successful versus those where the nonadaptive or equilibrium models of economic theory are more promising.

I do not know any general framework for addressing questions of this kind, so I will use the case method and discuss a series of examples. I will begin in Section II with an example of a social science question—the control of inflation—that has been successfully solved by economic methods. I hope that this example will work against a tendency that often appears in methodological discussions to disdain existing theories in favor of the yet-to-be-constructed models of the future. Economics works surprisingly well, under some conditions, and I think progress is more likely to follow from an understanding of the factors that have contributed to past successes and from trying to build on them than from attempts to reconstruct economics from the ground up in the image of some other science.

The rest of the paper consists mainly of more examples, each of which is drawn from recent work that is in some sense on the methodological boundary between economics and psychology. In each case, my focus will be on what methods I think of as "psychological" may contribute to the solution of economic problems of the sort described in Section II. Section III reviews a series of individual choice experiments with pigeons reported by Battalio et al. Section IV discusses some market experiments using undergraduate subjects conducted by Smith. Section V reviews a theoretical example by Bray on the stability of rational expectations equilibria. Section VI states a problem in monetary theory, and Section VII proposes an experimental resolution of it. In the course of discussing this research, I will also be advertising in passing a number of studies related to these. There is a good deal of current research on this interdisciplinary boundary, and I think that it has much promise for clarifying and advancing progress on some traditionally "economic" issues. I will expand on this belief, briefly, in Section VIII.

II. The Quantity Theory of Money

I will take an example of a solved economic problem from macroeconomics, partly because that is my own area of expertise but also because the aggregative character of macroeconomic problems serves to emphasize the distance between much of economics and the concerns of individual psychology. The particular theory I will discuss is called the quantity theory of money.

Consider the "equation of exchange,"

$$Mv = Py. \tag{1}$$

Here P denotes an economy's general level of prices at a point in time (as measured, say, by the consumer price index) and y denotes the rate of real production (real GNP, say). On the left, M is the quantity of money in circulation (say, M1, currency held by the public plus all checking deposits). Then v, velocity, is defined in terms of these other three magnitudes so that (1) always holds. One way of stating the quantity theory of money (there are many, not all consistent, as is true of most interesting theories) is to add to (1) the assumptions that velocity is constant and that movements in real output y are not affected by movements in M. Of course, neither of these assumptions is likely to be literally true (that velocity is not constant can be checked by plotting the observable series $P_t y_t / M_t$ against time for any economy), but this is to be expected in any theoretical model. If this theory were true, however, the rate of inflation, $(1/P)(dP/dt)$, would satisfy, from (1) and the assumption of constant velocity

$$\frac{1}{P}\frac{dP}{dt} = \frac{1}{M}\frac{dM}{dt} - \frac{1}{y}\frac{dy}{dt}.$$

Moreover, $(1/y)(dy/dt)$ would not vary systematically with $(1/M) \times (dM/dt)$. This way of stating the theory suggests that, if observations on $(1/P)(dP/dt)$ are plotted against corresponding observations on $(1/M)(dM/dt)$, the points so plotted should lie along a line with slope one: a 45-degree line shifted down by the rate of income growth. Figure 1 exhibits such a plot. (The graphics are from Lucas [1980], but the numbers were taken from Robert Vogel's [1974] study.)

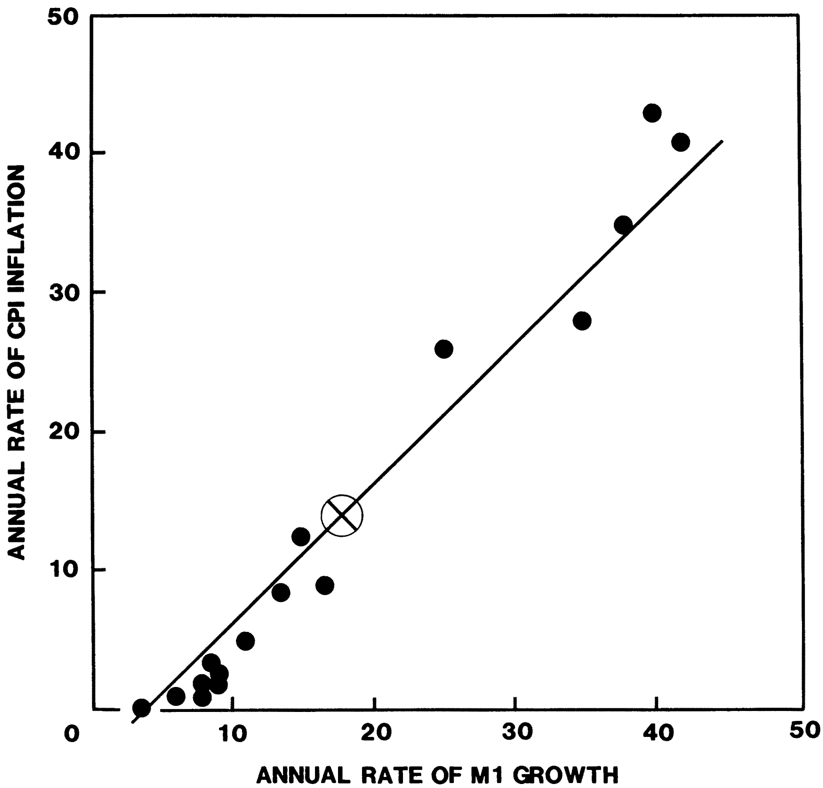

FIG. 1.—Sample averages from 16 Latin American countries, 1950–69

The X-coordinate for each point in figure 1 is the average rate of money growth for a single Latin American country over the period 1950–69, and the corresponding Y-coordinate is that country's average inflation rate for the same period. The line on the figure is drawn so as to pass through the average of all of these points, but its slope is 45 degrees, as predicted by the quantity theory. It is not fitted to the data. It is not easy to think of examples of nonvacuous social science theories that have recorded this kind of empirical success.

Not only does the quantity theory of money fit data, in the sense of figure 1, but it provides an operational answer to a problem of great social importance, the control of inflation. The rate of growth of money, in any country, can be controlled quite directly through government and central bank policy. That is, the location of a country on the X-axis of figure 1 is a matter of fairly simple policy choice. If, as figure 1 confirms, velocity and real income growth are largely independent of this choice, a society can thus indirectly dictate its long-run average inflation rate—its location on the Y-axis. This is what I meant in the introduction when I said that the problem of controlling inflation has been "successfully solved" in a scientific sense.[1]

So that we do not get carried away with this success, figure 2 (from my 1980 paper) also plots inflation rates against M1 growth rates, where in this case each observation is from a single quarter during the period 1955–75 for the United States. Where did the good fit go? To recover it, each observation in figure 2 can be replaced with a very long moving average of adjacent quarters' points, producing figure 3. Comparing figures 2 and 3 we can see that the use of averaged (over time) inflation and money growth rates was not incidental to the results displayed in figure 1. Without such averaging, the quantity theory (at least in the form that I have presented) does not provide a serviceable account of comovements in money and inflation.

What is the relationship between these three figures and economic theory? In particular, what role does the assumption that agents behave "rationally" play in equation (1) and in the assumptions I appended to this identity to obtain a nonvacuous model? This is not a simple question because each stage in the development of economic theory has produced its own collection of monetary theories, some consistent with the quantity theory and some not, so that it is possible to set out neither a set of agreed-on axioms on which monetary theory in general is "based" nor a set that is equivalent to the quantity theory in the form in which I have "tested" it. I think that this is a typical state of affairs in economics and not at all a deplorable one, but it does complicate the discussion.

It is certain that the quantity theory did not originate as an empirical generalization based on evidence such as that summarized in figure 1. When David Hume (1963) first enunciated the hypothesis, in 1742, the data needed to construct figure 1 were not collected for any economy. Nor did Hume deduce the theory from the axioms of utility theory, for the development of this useful equipment was more than a century in the future. Hume's argument was based on the idea that people hold money for the sole purpose of spending it on goods so that changes in money and prices in equal proportion "ought" to be pure units changes, affecting no one's decisions. The argument is tricky because a change in money does not automatically cause prices to move equiproportionally in any direct sense, so the proposition that individual behavior is invariant to units changes does not in itself give the result. One needs to argue, as Hume did, both that if prices move in proportion to the increase in money individuals will be willing to hold the increase and that, if prices do not move in proportion, money demand

1. Obviously, few societies have solved the problem of inflation in a political sense. I do not see this fact as qualifying my claim in the text, any more than I would view the current popularity of "creationism" as qualifying the scientific status of the theory of evolution.

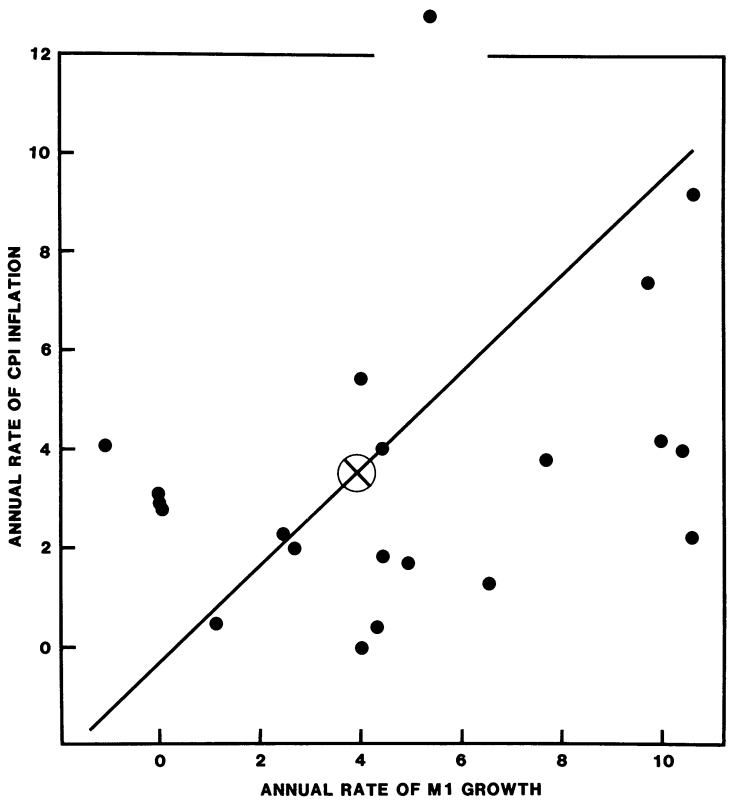

FIG. 2.—Original data for second quarters, 1955–75

cannot equal money supplied so that further price movements must follow. Nevertheless, there is a clear sense in which the theory rests on the hypothesis of individual rationality, at least in this limited sense that units that "ought" not to matter to "rational" people are assumed not to matter in fact.

In more recent times, a great variety of theories of monetary economies have been devised, in which the utility-maximizing behavior of agents and the environment in which they are assumed to interact have been made increasingly explicit. In the main, these theoretical developments have reinforced our understanding of and, I would say, belief in the quantity theory in roughly the form that I have stated it. But this theoretical work has also illuminated several distinct ways in which money can be "nonneutral," or ways in which changes in the rate of growth of money differ from pure units changes. Indeed, in Section VI, I will review one monetary model that has a continuum of equilibria, only one of which is quantity theoretic.

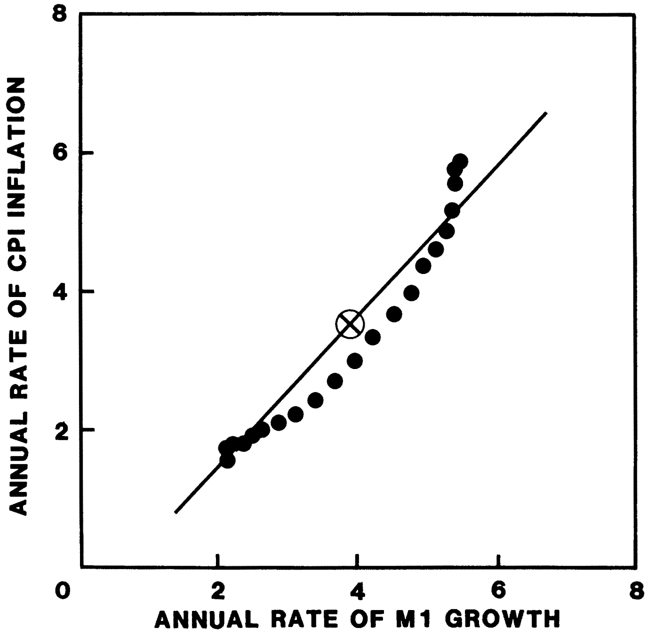

FIG. 3.—Smoothed data for second quarters, 1955–75

Given that Hume arrived at an empirically successful version of the quantity theory with relatively informal reasoning, what is the contribution of these more refined theoretical developments? I think there are several. First, by being explicit about agents' preferences, modern theory can illuminate the consequences of differing inflation rates on individual consumer welfare: equation (1) predicts that rapid money growth is inflationary but sheds no light on what is wrong with inflation. Second, theory that does not imply constant velocity under all circumstances suggests useful limits on the range of applicability of equation (1), as I have used it. Thus, the large movements in velocity observed during the onset of hyperinflations, predicted by more refined theory but not by (1), ought to reinforce our view that figure 1 is an important confirmation of the quantity theory: the theory works well when it "ought" to and fails when it "ought" to as well. The fact that we now view Newton's theory of gravity as a special case holding only in certain circumstances does not in any way compromise the usefulness of that theory. Third, and this is more an expression of hope than anything else, a main objective of more recent theorizing is to obtain models that are consistent with figures 1 and 3 and with figure 2. The simple form (1) of the quantity theory tells us nothing about the patterns (if there are any) shown in figure 2, nor does it suggest a line of attack on the problems of short-run monetary dynamics. I think the whole problem of the business cycle is hidden somewhere in this picture, but we do not yet have the monetary theory that can let us see if this conjecture is right or wrong.

Let me summarize. There are axiomatic developments of the quantity theory of money, but we do not believe that this theory is useful because it is built up from impeccable foundations. On the contrary, the more we understand the foundations, the more limits we see to the applicability of the theory. The empirical testing of the theory is critical precisely because we know that the axioms are abstractions, necessarily "false," so we need to know whether and under what range of circumstances these abstractions are adequate. Conversely, we learn very little about the axioms of utility theory from tests of its aggregative implications. Hume derived most of these implications before the theory of utility was discovered, and it seems likely to me (though we cannot know for sure in advance) that future evolution in utility theory will produce new statements of the quantity theory very similar to the ones we have now.

What is also clear is that adaptive elements of behavior have played no visible role in the theoretical development of the quantity theory or in its testing. If (as I have claimed) all behavior is adaptive, how can it be that important propositions about behavior can be obtained and applied without making any reference to this fact? Let us keep this question in mind.

III. An Individual Choice Experiment

The quantity theory of money provides a good illustration of the way that theoretical economic reasoning can lead to nonvacuous models that bear on questions of importance, but it is not a good context for isolating the contribution of rationality per se to economic modeling. Though all theoretical developments of the quantity theory assume some form of rational behavior, they require as well assumptions on the way agents interact and assumptions that certain theoretically possible feedback effects are small enough to be neglected. In this section and the next, I will draw on examples of research in which the separate effects of rationality and these other factors can be isolated and examined much more fruitfully.

There is an interesting and growing body of research, centered at Texas A&M University, involving experiments on individual choice with animals as subjects. The work is marked by its unusual mix of sophistication in the design and conduct of the experiments—methods drawn from psychological research—and in its use of the economic theory of choice. An excellent, recent example is provided by Battalio et al. (1981), who report results on commodity choice by pigeons. These experiments provided pigeons a choice between food and water in a controlled environment where "the price of each commodity was varied by altering the average time between deliveries of that commodity . . . while the income constraint was the total time available for the delivery of the two goods'' (p. 69). The pigeon exercised his choice with a single peck on a control key that determined whether his current income flow (out of a total stock of delivery time) was to go for food or for water. He could switch between these two goods at any time.

Battalio et al. are explicit about the background of the pigeons: ''Four male White Carneaux pigeons with no previous experimental history served as subjects. . . . All [subjects] had extensive training with the procedures prior to the start of the experiment. Experimental conditions [i.e. prices] were changed when inspection of graphs of the data indicated an absence of any significant drift in consumption patterns'' (p. 71).

The main focus of the experiments is the degree to which the choice behavior of these pigeons satisfied the weak axiom of revealed preference. The statistic used to test this hypothesis used an average over the last 5 days of choices at ''baseline'' (original) prices and over the last 10 days at a new set of prices. Total days at each set of prices varied (over subjects and over price vectors) from 5 to 56 days.

The authors report on the conformity of the choice patterns to which each pigeon converges (roughly) with the axiom of revealed preference. For this particular set of experiments, it is high (though the indifference maps appear quite different across individual subjects). They also report on the adjustment paths or learning curves the subjects exhibited following switches from one price vector to another, finding interesting differences between responses to increases and decreases in relative food prices. For some subjects, repeating the baseline price vector later on induces roughly the original choices; for others, it does not.

To evaluate the results of this particular experiment would require a much more detailed description than I have provided, as well as familiarity with closely related work with pigeons and other subjects. But my present interest is rather in what is taken for granted about the relationship of economic theory to observed behavior in this work, both in the design of the experiments and in the interpretation of the results.

The main hypothesis tested is derived from the economic theory of choice. The theory was designed to refer to human decision making, not that of pigeons, and is on a different level from curve-fitting generalizations from earlier experimental work (such as psychological learning curves). The theory delivers nonobvious predictions about behavior in interesting circumstances, and these predictions are fairly well confirmed. These are striking and exciting findings.

It is assumed from the outset, however, that the economic theory of choice will fall well short of a complete model of the decision-making process of pigeons. First, the predictions are limited in scope to the behavior of average performance over several trials, with the averaged trials selected so as to exclude an initial period of changing behavior patterns. It is clear from the reported results that had these restrictions not been imposed, the theory would have fit the data very badly. Second, the theory, even interpreted as a model of these limited aspects of behavior, does not attempt to deal with choice at the more fundamental physiological level. That is, pigeons used food and water to stay alive (they had no access to either, outside the "price systems" imposed by the experiment), and, given that pigeons have survived as a species to this date, we can infer that they are equipped with internal mechanisms that detect deficiencies in existing "stocks" of both and trigger actions to deal with them. The economic theory is in no sense derived from a description of these underlying mechanisms, nor is it an attempt to provide such a description. The economic theory being tested, then, is limited both in scope and in depth.

This application of the theory of choice presupposes, then, the existence of a broader and deeper theory. Though this theory is not spelled out, certain features of it are taken for granted in the design of the experiment. It is assumed, in the first place, that the subject's behavior is adaptive. His initial behavior will be influenced by his genetic makeup and his previous experience (which is why aspects of both are given explicitly in describing the experiment). This behavior will involve some erratic or "experimental" actions on the pigeons' part, as well as a continuing evaluation of outcomes. Further, it involves some presumption (on the subjects' part) of stationarity in the environment, so actions that yield good outcomes are repeated and those that yield bad ones are not, or are at least used less frequently.

The economic theory of choice is thus interpreted as a description of a kind of stationary "point" of this dynamic, adaptive process. The pigeons' demand functions are decision rules arrived at after a process of deliberate experimentation and assessment of outcomes. The behavior implied by these decision rules is "rational" in the sense that economists use that term. But not only is it consistent with adaptive, trial-and-error behavior; the experiment designed to discover this rationality assumes that, if it exists, it is the outcome of some (unspecified) adaptive process.

Would it be possible to reinterpret this entire process as "rational" in this sense, as the solution to some more complex maximum problem? I assume so. Every "point" must be at the top of some hill, in some "space."[2] But what would be the empirical reward from doing

2. The subjects, e.g., could be modeled as having a prior distribution over possible experimental setups and choose pecks on the control key so as to maximize the expected value (with respect to this distribution) of an objective that assigns value both to the immediate food-water reward and to the new information gained at each stage.

so? It is clear that the time path of a subject's behavior will depend on his "initial conditions"—on what he makes of the experimental scene when he is first introduced to it—and we have no way of knowing what these are. Battalio et al. sidestepped this problem by hoping that, whatever these initial conditions look like, their influence will disappear over time (that the underlying process is stable). By giving up on an empty theory of the entire process, they obtained in return a theory with real content about certain very important aspects of it.

On the second limitation of the theory of choice—its physiological shallowness—it seems to me that much more could be done. Battalio et al. have some interesting speculations about mechanisms that might underlie the observed asymmetry in subjects' responses to upward and downward movements in the relative food price. I think that these could be modeled and that such models would have additional testable implications and would also illuminate the economic theory—the indifference maps of pigeons over food-water combinations. Battalio et al. do not pursue this (perhaps this is on the agenda for future work), but, even if they had, I think it is likely that they would have been led to progressively more pigeon-specific models, models with less transferability to behavior of other subjects or to pigeons in other situations. Economists apply essentially the same model of choice to pigeon choices over food-water pairs as to, say, a corporation's choice over capital goods of differing durability. Insofar as much decisions can successfully be viewed in a unified way, it is not likely to be at the physiological level.

In summary, it is clear that the research on economic rationality in animal subjects rests on a maintained idea that behavior is determined by an adaptive process, with the economic theory of choice interpreted as applying to some kind of stationary point of this process. This is the way in which Battalio et al. interpret their results with pigeons, and it seems to me the only interpretation that is tenable. Moreover, it is inconceivable to me that this same general idea cannot be carried over, in some form, to interpreting the application of economic theory to human behavior in actual market situations. But this is getting slightly ahead of my story.

IV. Market Experiments with Human Subjects

Applications of economic theory to market or group behavior require assumptions about the mode of interaction among agents as well as about individual behavior. For example, a competitive equilibrium (the concept typically—though not necessarily—underlying quantity-theoretic models) assumes that each agent takes prices as given and that no trading occurs at non-market-clearing prices. Just as the assumption of individual rationality abstracts from the adaptive aspects we know are present in actual individual behavior, so does the assumption of competitive (or Nash) equilibrium abstract from adaptive aspects of group behavior. The consequences of this quite different abstraction can also be isolated and studied experimentally.

A large body of experimental results bearing on this aspect of the applicability of economic theory has been obtained by using human subjects in market systems. In this research, subjects are simply given a preference function by the experimenter, who pays them in dollars according to how well they succeed in maximizing their induced "utility." Hence no information about subjects' actual preferences over goods is obtained. On the other hand, subjects are left quite free as to how they choose to interact—to trade—with one another. The objective is to see under what conditions the predictions of competitive market equilibrium theory (given the artificially induced preferences of subjects) conform to the actual quantities exchanged and the prices at which these exchanges actually take place.

Vernon Smith's (1962) paper was the pioneering effort along this line.[3] Smith divided subjects into two groups, buyers and sellers, assigning to each buyer a "reservation price" giving the maximum amount he or she could pay for 1 unit of a good (no buyer could purchase more) and to each seller a minimum price at which he or she could sell the 1 unit of the good with which he was endowed. Each subject was rewarded in proportion to the difference between his maximum (minimum for sellers) price and the price at which he actually transacted. Subjects interacted during trading rounds lasting from 5 to 10 minutes, during which they were free to make public, verbal offers of any kind. On acceptance of an offer, the buyer-seller pair so matched withdrew from the market. On completion of a round, a new round, identical in structure and with identical preferences, opened. The process continued for two to six rounds.

"The most striking general characteristic of tests 1–3, 5–7, 9 and 10 is the remarkably strong tendency for exchange prices to approach the predicted equilibrium for each of these markets. As the exchange process is repeated through successive trading periods with the same conditions of supply and demand prevailing initially in each period, the variation in exchange prices tends to decline, and to cluster more closely around the equilibrium" (Smith 1962, p. 116). (In the exceptions, tests 4 and 8, the results exhibited a kind of bargaining power not

3. Since experimental methods have undergone considerable evolution since 1962, a more recent example might have been a more suitable basis for this discussion. But in comparing Smith (1982) to the main conclusions of Smith (1962), I am struck with the extent to which the main early findings of this research have stood up over time and over many replications.

predicted by competitive theory.) Smith obtained these results with about 10 subjects on each side of the market.

In 1874, Leon Walras (1954) had provided the first explicit theoretical description of a set of market "rules" under which it could usefully be asked whether and/or how a collection of economic agents could arrive at an equilibrium price—a price at which the quantity demanded by buyers of a good would equal the amount sellers wished to supply. Walras's scenario is centered on an auctioneer who initiates the process by announcing an arbitrary trial price. Agents then indicate how much they would be willing to buy or sell if this price should prevail. They have an incentive to answer this hypothetical question truthfully; if the announced price does in fact prevail, they are required to deliver or purchase whatever they said they would be willing to do. If the trial price equates demand and supply, it does prevail and trade is consummated. If it does not "clear the market," all bets are off and the auctioneer selects a new trial price. Under some quite reasonable conditions, this adaptive process (though it is only the auctioneer who does any adapting) converges to the market clearing price.

In Smith's experiment, as in subsequent experimental work, the market mechanics are not at all Walrasian: subjects set prices as they please, with no auctioneer to guide them. In Walras's auction, either price converges to the competitive equilibrium, or no trade occurs at all. In Smith's, trade can occur at any price that any buyer-seller pair can mutually agree on. Equilibrium prices in Smith's setting turned out to be stable (even when they differed from the competitive prediction), but patterns of convergence were not well described by Walras's adjustment hypothesis.

Walras's point of departure was the idea that an economic equilibrium is not an empirically interesting object unless one can imagine some way that a group of economic agents, with ordinary human mental equipment, might actually hit on it. The mechanism he proposed has the virtues of being concrete, of relying only on simple adaptive capacities, and of being, under a wide range of circumstances, stable. Smith's experimental setting retained these important features, but shifted the task of adaptation from the auctioneer to the same agents whose preferences determine the equilibrium, and permitted trades to be consummated whenever mutually agreeable, just as they are in actual free markets. In doing this, he reformulated the problem of stability of equilibria as a question (or set of questions) about the behavior of actual people—as a psychological question—as opposed to a question about an abstract and impersonal "market." His and subsequent experimental work has done much to illuminate this question, but in so doing it has left the standard theory of stability far behind.

V. Stability Theory: An Example

Recent work in stability theory has begun to examine situations in which convergence to equilibrium rests on adaptive behavior of individual agents (as opposed to the Walras auctioneer). An example given by Margaret Bray (1983) will serve to illustrate the main ideas.[4]

Bray's example concerns a sequence of spot markets where the market clearing price at date t, p_t, depends on the price that agents expect will prevail next period, date $t + 1$. (It is easy to think of models that would have this character. One example will be provided in the next section.) Call this expected price p^e_{t+1} (an expectation formed in t about an event in $t + 1$) and assume

$$p_t = a + bp^e_{t+1} + \epsilon_t. \tag{2}$$

Here $\{\epsilon_t\}$ is a sequence of independently and identically distributed normal "shocks" with mean zero. Following Muth (1961), call the price expectation

$$p^e = \frac{a}{1 - b} \tag{3}$$

rational because, if (for some reason) people always expect next period's price to be given by (3), the actual prices $\{p_t\}$ will be a sequence of independently and identically distributed normal random variables with mean $a/(1 - b)$ and the expectation that (3) will be confirmed, on average, by experience.

In (2), the current price p_t is market clearing, set presumably by some stable process such as Walras's or Smith's. Any adapting that takes place must be on individuals' common forecast p^e_{t+1} of next period's price. In particular, what if people begin with some price expectations that are not rational in the sense of (3), as they would if the situation were new to them, as in Smith's experiments? They will need to form some belief about p_1 in order to engage in trade at date $t = 0$. Thereafter, they will need to decide how to use their accumulating experience with actual prices $p_0, p_1, \ldots, p_t$ in forming an expectation about the next term in the sequence, p_{t+1}. But what actually happens

4. Bray (1982, 1983), Blume and Easley (1982), and others have studied convergence to rational expectations equilibria under various adaptive hypotheses. Townsend (1978) and others have examined the same general question from the point of view of Bayesian decision theory (see n. 2 above). In the latter approach, the entire path of approach to the rational expectations equilibrium is itself an equilibrium of a suitably specified game. In the former, it is not. The two approaches are complementary and both have their uses, but if the question is how or whether adaptive behavior on the part of "irrational" agents will lead to "rational" behavior over time, only the first is germane. I found the brief discussion in Blume, Bray, and Easley (1982) helpful in clarifying the relationship between these two ideas of stability. Another adaptive approach to this stability question is sketched in Lucas (1978). In that paper, agents' preferences over market goods are formed adaptively, as agents learn about the utility actually yielded by purchased goods.

will depend, in turn, on what people expect to happen: actual and expected prices are simultaneously determined.

Bray assumed that people simply use an average of past actual prices as a forecast of the next price,

$$p_{t+1}^{e} = \frac{1}{t}\sum_{i=0}^{t-1} p_i. \tag{4}$$

Under this hypothesis, both actual and expected prices are well-defined stochastic processes (given an initial price expectation p_1^e), and it is shown that, provided $|b| < 1$, $\{p_t^e\}$ converges over time, with probability one (over realizations of the shocks $\{\epsilon_t\}$) to the rational expectation (3), for all initial values of p_1^e. In this specific sense, then, the rational expectations equilibrium is stable, given adaptive behavior of the form (4).

What is the empirical content of this model (or of more realistic and complicated models that involve the same basic elements as this example)? Does one take the rational expectations equilibrium, (2) and (3), as the model's prediction about the behavior of actual prices, or the adaptive path, (2) and (4)? Except in the limit, the two are not the same. It does not seem to me that this question is usefully posed in the abstract.

In applications such as that described in Section II, one would clearly take only the rational expectations equilibrium itself as a serious (though possibly empirically unsuccessful) hypothesis. The initial dates of 1950 or 1955 are not $t = 0$ in any behavioral sense; they are just the points at which Vogel's or my data sets happened to start. Moreover, we have no way of knowing what agents' beliefs about future prices were in various countries in 1950 and no reason at all to imagine that these beliefs were the same across countries or across individual agents within a country. In any case, using any adaptive scheme amounts to the conjecture that we econometricians, using only aggregate data on variables like M1 and the consumer price index, can discover rents that were available to, say, Argentinians during 1950–69 but were invisible to Argentinians themselves, who we know were processing thousands of data points in addition to those in our data sets. In aggregative applications of this character, then, one would take the rational expectations equilibrium—the appropriate analogue to (2)-(3)—as the model to be tested and view the adaptive hypothesis (4) as being, at most, an adjunct to the theory that serves to lend it plausibility.

In applications such as those described in the last section, in contrast, in which a group of subjects is observed from the first date at which they are introduced to a particular economic situation and begin to operate within it, it seems clear that subjects could hit on the behavior (2)-(3) from the outset only by coincidence so unlikely as not to be an empirically serious possibility. One would test (2)-(3) only as a prediction about behavior after many trials, exactly as in Bray's theory, or as in the experiments described in Sections III and IV. For predicting actual behavior from $t = 0$ on, an adaptive hypothesis like (4) would be a serious candidate for describing actual behavior. Even if it should work poorly (as I think it would, based on Smith's and others' experimental results), the general idea of averaging past experience on which it is based is a flexible one, and it seems likely that some scheme of this type could provide a good description of the adaptive behavior we do observe.

VI. A Problem in Monetary Theory

For many problems in applied economics, then, the fact that people behave adaptively is of little or no operational consequence: one assumes that people have long ago hit on decision rules suited to their situations—"rational" rules—and utilizes theories about these rules to predict behavior. But this is certainly not true of all problems of interest. Even so well-established a theory as the quantity theory of money, reviewed in Section II, is subject to difficulties that I do not believe can be resolved on purely "economic" grounds. These difficulties involve the multiplicity of perfect foresight or rational expectations equilibrium paths in a Samuelson-type overlapping-generations model of a monetary economy. I will use an example to state the theoretical issue in this section. In the next, I will describe an experiment that I think is capable of fully resolving it.

The issue can be stated briefly, using a specific version of Paul Samuelson's (1958) model. The economy runs in discrete time, forever. Each period, N agents are born, each living for 2 periods, each endowed with 1 unit of a nonstorable consumption good in the first period of life and none in the second. An agent born in t has preferences $U(c_t^y, c_{t+1}^o)$ over consumption at t, c_t^y, and consumption at $t + 1$, c_{t+1}^o. Feasible allocations are nonnegative and satisfy $c_t^y + c_t^o = 1$, all t.

At $t = 0$, old agents each hold 1 unit of fiat money. Trade involves the young exchanging goods for the money held by the old. Letting q_t be the inverse of the price level (goods per unit of money), the decision problem of a young trader born in t is then

$$\max_m U(1 - q_t m, q_{t+1} m),$$

where m is the money balances he chooses to acquire in trade. In equilibrium, the first-order condition for this maximum problem must be satisfied (I assume increasing, concave U) at $m = 1$. Thus, one equilibrium condition is

$$U_1(1 - q_t, q_{t+1})q_t = U_2(1 - q_t, q_{t+1})q_{t+1}. \tag{5}$$

Nonnegativity adds another,

$$0 \leq q_t \leq 1. \tag{6}$$

Any solution to the implicit first-order difference equation (5) that satisfies (6) is a "perfect foresight" or "rational expectations" equilibrium.

Note that stationary solutions to (5) (sequences $\{q_t\}$ with q_t constant) correspond to the quantity theory of money. If the money supply is constant, so is the inverse q of the price level. If the money supply is initially doubled, the stationary equilibrium value of q is halved. Other solutions to (5) will not have these properties. Note also that (5) provides an example (though a nonlinear one) of an equilibrium condition of the form (2) postulated by Bray.

Until U is specified, this theory has a lot of possibilities. Since it is not my purpose here to explore all of these, let me specialize to the particular preferences

$$U(c_t^y, c_{t+1}^o) = (c_t^y)^{1/2} + 2(c_{t+1}^o), \tag{7}$$

so the equilibrium condition (5) becomes

$$q_{t+1} = \frac{1}{4}(1 - q_t)^{-1/2} q_t. \tag{8}$$

Figure 4 plots the right-hand side of (8) against the 45-degree line.

The stationary points of (8) are $q = 0$ and $q = 15/16$. The solution 15/16 is the quantity-theoretic equilibrium. The solution $q = 0$ describes a situation in which no agent has the faith that other traders will accept money at later dates, in which case money is valueless. Any solution to (8) with $q_0 > 15/16$ will violate (6) for some t. All solutions with $0 \leq q_0 \leq 15/16$ satisfy (6) for all t. All these solutions are perfectly legitimate equilibria. It is abundantly clear from much theoretical work that this multiplicity of equilibria does not in general disappear as this intergenerational model is complicated in various ways, provided all agents are assumed to be finitely lived. The simplicity of the example reveals the existence of a continuum of equilibria but it does not create them.

Indeed, if one thinks of trade in this economy as taking place in a sequence of spot markets, as seems necessary given its demographic structure, there are still many more sequences $\{q_t\}$ that may be interpreted as equilibrium prices. Suppose, for example, that at $t = 0$ a price $q_0 \in (0, 15/16)$ is established by young agents who believe, unanimously, that q_1 is given by (8). As established above, this is equilibrium behavior, no matter what q_0 value in this interval is hit on. Next period,

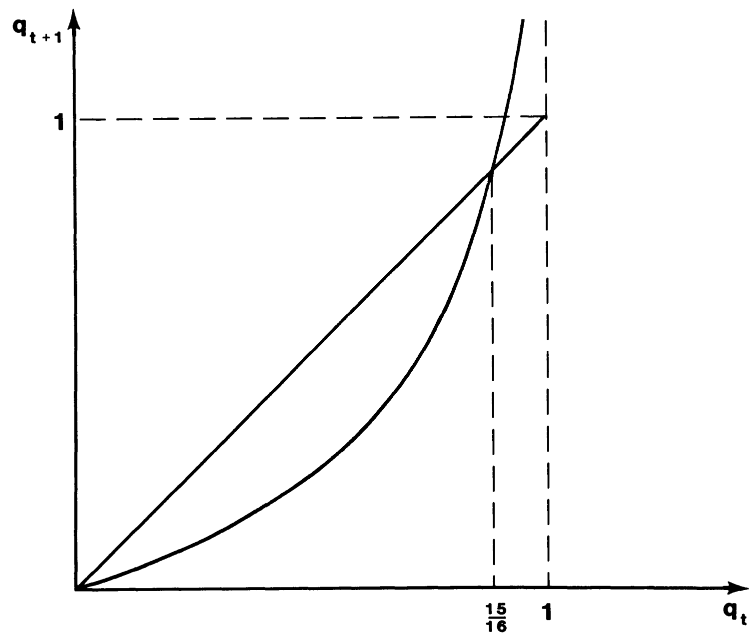

FIG. 4.—Price-level dynamics of eq. (8)

a new generation arrives, facing a situation that is exactly the same, in all respects, as that faced by the preceding generation. (Calendar time is clearly immaterial, as is history in the definition of a perfect foresight equilibrium.) Hence, for this new generation as well, any $q_1 \in (0, 15/16)$ is an equilibrium price. Continuing in this way, any sequence $\{q_t\}$, $0 \leqslant q_t \leqslant 15/16$, represents equilibrium behavior.

It is instructive to compare this theoretically confused situation with an otherwise identical economy with a finite life T and a given terminal price q_T of money. With such a given boundary condition, (8) has a unique solution $\{q_t\}$ for any $q_T \in [0, 1]$. If $q_T = 0$, this solution is $q_t = 0$, all t. If $q_T \in (0, 1]$, it is clear from figure 4 that an equilibrium $\{q_t\}$ must remain close to 15/16 most of the time, moving toward q_T appreciably only close to the terminal time. (With $q_T = .01$, e.g., and rounding prices to two places, $q_t = .94$ [$\cong$ 15/16] until $T - 5$, following the sequence .92, .79, .45, .14, .04, .01, home from this point on.)

In this finitely lived economy, then, the multiplicity of equilibria in the sense of solutions to (8) is entirely absent. So too, then, is the additional multiplicity arising from the irrelevance of calendar time. Calendar time does matter in the finite system since each generation is 1 period closer to the end than its predecessor and hence faces an objectively different solution.

The simplicity of the finite model contrasted to the apparently hopeless complexity of the infinite one has seemed to many to hold the promise that there is some purely mathematical way by which the paradoxes raised by the infinite-horizon case might be resolved. Cannot one simply let T go to infinity in the finite case and conclude that the limiting equilibrium behavior ($q = 15/16$, in our example) is the "right" equilibrium for the infinite case? Viewed as a purely economic question, the answer is no. The infinite horizon case offers genuine equilibrium possibilities that are not approximated by the limits of sequences of finite-horizon equilibria.[5]

The stability theory proposed by Bray does offer a resolution to the multiplicity problem arising in the infinitely lived economy. Thus write (8) as

$$q^e_{t+1} = \frac{1}{4}(1 - q_t)^{-1/2}q_t, \quad t = 0, 1, 2, \ldots, \tag{9}$$

where q^e_{t+1} is a point expectation formed at time t about the price in $t + 1$. This plays the role of (2). Then, as in (4), suppose q^e_t is formed adaptively; as a simple average of past, actual prices and the initial expectation q^e_t,

$$q^e_{t+1} = \frac{t}{t+1}q^e_t + \frac{1}{t+1}q_{t-1}, \quad t = 1, 2, \ldots. \tag{10}$$

Given an initial expectation q^e_1, the actual price q_0 is obtained from (9). Then the new forecast q^e_2 is given by (10), q_1 by (9), and so forth. It is easy to show diagrammatically (fig. 5) that the sequences $\{q_t\}$ and $\{q^e_t\}$ so generated satisfy

$$\lim_{t\to\infty} q_t = \lim_{t\to\infty} q^e_t = \frac{15}{16} \tag{11}$$

for all $q^e_1 \in (0, 1)$. (If $q^e_1 = 0$, which would imply that no one initially believes the money will be valued in the future, the solution is $q^e_t = q_t = 0$, all t.) That is, the system converges to the stationary rational expectations equilibrium.[6]

Figure 5 illustrates the proof of (11) for the case $0 < q^e_1 < 15/16$. Given q^e_{t+1} on the vertical axis, the actual price q_t on the horizontal axis can be read off the curve $1/4(1 - q)^{-1/2}q$. Since the curve is below the 45-degree line, $q_t > q^e_{t+1}$. Then the new forecast q^e_{t+2}, being an average of the old one and something larger, exceeds q^e_{t+1}, which implies that $q_{t+1} > q_t$, and so on. Bray's stability hypothesis comes close to running

5. McCallum (1983) has proposed as a "methodological principle" that equilibria with a "minimal set of state variables" be preferred. This principle, in the present context, would select the stationary equilibrium, but it is unclear what the behavioral rationale for this principle is. I think that the experiment proposed in Sec. VII is, however, very much in the spirit of McCallum's argument.

6. Proposition 2 in Wallace (1980, p. 56) is almost identical to the proposition illustrated in fig. 5.

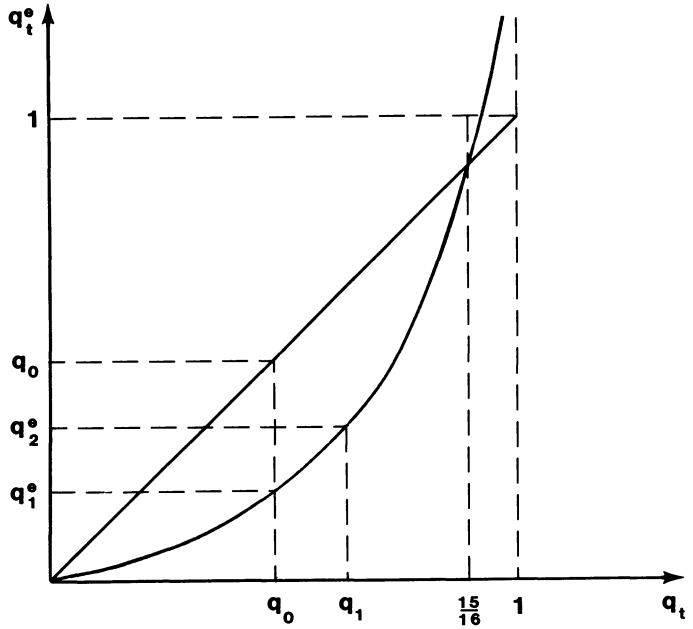

FIG. 5.—Price-level dynamics of eqq. (9) and (10)

the unstable difference equation (8) backward, converting it into a stable one.

The kind of adaptive behavior captured in this simple model seems to me to be a plausible conjecture as to how people might actually behave. Since it singles out the stationary equilibrium 15/16 as the one that would be converged to it seems to give this equilibrium a special substantive interest not shared by the infinity of other equilibria. But as a purely theoretical argument, this stability example does not seem to me to settle anything. The adaptive behavior it assumes is not based on any economic principle: (8) exhausts the implications of competition and "rationality," and it does not single out any one equilibrium path. Figure 5 is roughly consistent with what we know psychologically about the way people tend to behave in new situations, but so would be innumerable other adaptive schemes, different from the above, that one could have as easily worked through. In any case, the fact that one can produce an adaptive scheme that singles out a particular equilibrium does not rule out the possibility of producing other adaptive schemes that single out other equilibria or even suggest that this would be difficult to do.

The most that can be offered by the kind of stability argument just given seems to me to be the suggestion of the kind of experiment that might genuinely single out a particular equilibrium as being of more substantive interest than the others. The issue involves a question concerning how collections of people behave in a specific situation. Economic theory does not resolve the question. One can imagine other principles that would, but this cannot rule out the possibility that still other principles might resolve it quite differently. It is hard to see what can advance the discussion short of assembling a collection of people, putting them in the situation of interest, and observing what they do.

VII. A Proposed Experiment

The problem involved in convincing a collection of experimental subjects that they are in an infinite-horizon environment seems to me insurmountable. (Even to spell out what this means is not easy.) The central issues—whether people initially behave adaptively and, if so, what form this adaptive behavior takes—are as easily stated in the finite horizon case as in the infinite one. This observation (by Nancy Stokey) leads to the following proposed design.

Take N to be a number large enough to assure roughly competitive behavior in static experimental situations (say, 8 or 10), and let the number of subjects be $3N$, divided into three groups of N each. At $t = 0$, one group will be "old," each endowed with 100 white chips, each representing 0.01 units of fiat money (one "cent"). A second group will be endowed with 100 blue chips each, each representing 0.01 units of goods. The third group is in waiting; they will play the role of the young at $t = 1$. These groups are to rotate through ages 0, 1, and "nonexistent" for the duration of the experiment and are so informed.[7]

All subjects are informed that, at the end of each period as a young agent, each is to turn in all blue chips retained, and each is scored in proportion to the square root of this amount. At the end of each period as an old agent, each is to turn in all blue and white chips, and scores are given in proportion to twice the number of blue chips. The total pay each subject is to receive over the life of the experiment is simply the sum of the rewards so acquired in all of his successive "lives." Any white chips returned by subjects at "death" are redistributed as equally as possible to the newly "old" generation (acquired at price of zero) as of the beginning of the next period.

Subjects are convincingly informed that the experiment will last exactly T periods and that, at the conclusion of period T, the group playing the role of the "young" at that time will receive a fixed number q_T of blue chips for each white chip held at the end of that period. No other information as to the intrinsic value of white chips is given any subject at any time.

Subjects will be left free, in each period, to exchange white chips for

7. Sunder (1985) is currently conducting experiments of this general structure.

blue (or not to do so) on any terms they choose, on an entirely individual basis, exactly as in the Smith experiment described in Section IV. A period ends when no pair of subjects wishes to engage in further exchange. Thus blue chips and white will be traded, in general, in a variety of quantities and prices. Each period, the ratio $\hat{q}_t$ of total blue chips surrendered to total white chips surrended in exchange is recorded, and the history $(\hat{q}_0, \hat{q}_1, \ldots, \hat{q}_{t-1})$ is clearly displayed to all traders in period t. The sequence $(\hat{q}_0, \hat{q}_1, \ldots, \hat{q}_{T-1})$ is regarded as the outcome of experiment (T, q_T). The series of experiments here proposed involves varying T over various values (e.g., $T = 10, 50, 100$) and varying q_T over the interval [0, 1] at discrete values most definitely including the end points.

An immediate benefit from the discipline imposed by the attempt to set out an operational experimental counterpart to the theoretical economy of Section VI is that one is led to take that theory seriously as an aid in thinking about how actual people might really behave in the given situation. One is led to ask, In what respects is equation (8) a serious model of human behavior, and in what respects is it not?

Surely it is unlikely that any sizable group of subjects would unanimously realize that their collective behavior "should" be described by (8), whatever this means. Even if they did, it is less likely still that all would solve (8) backward from the given terminal condition (T, q_T) to the "correct" value of q_0. Yet if this turn of events is wildly implausible, how much less plausible is it that similar subjects, situated in an infinite stage version of this same economy, should unanimously hit on the identical, wholly arbitrary value of q_0 that sets them and (somehow) their successors off on one of this economy's unstable equilibrium paths? Insofar as "perfect foresight" or "rational expectations" equilibria are useful social-scientific constructs, it must be in some other sense than this.

The sense in which these constructs are useful is, I think, something like this. The subjects in the experiment just described will have to take some kind of guess as to what white chips acquired today in trade will be worth tomorrow. Without knowing a good deal more about these people than that they are "rational," it seems obviously impossible to predict in any reliable way what these guesses will be. Even rational people assess new situations in the light of their own experience and without knowing much about these experiences; how can one predict their assessments? Yet unless all subjects are convinced that white chips are forever valueless (and I do not believe this can be brought about even by *telling* them that $q_T = 0$ if T is as large as 10 or 15), trade will occur and some positive value for money will be established.[8] If

8. This conjecture is not entirely without foundation, since the experiments proposed above are similar in many respects to experiments with repeated Prisoner's Dilemma

so, and if q_t is much off the value 15/16, subjects will see that available rewards have been passed over in the past and will adapt their behavior in the direction of claiming these rents in the future.

These conjectures are very much in the spirit of Bray's stability model, and, indeed, it would be interesting to see if the simple average forecasting rule that she assumed performs well as a description of price formation in early rounds of trading. My guess is that it would not—subjects in Smith's and subsequent experiments seem to behave more erratically, perhaps because they are themselves experimenting a little, than would be consistent with averaging alone—but I would also guess that whatever price formation patterns are observed would have stability properties identical to those derived by Bray.

VIII. Conclusions

This paper has been an inquiry into the role, actual and potential, of adaptive elements in empirically oriented economic theory. Rather than attempt a general characterization of this role, I have proceeded by the method of cases, using examples of specific social science research that seemed to be capable of shedding light on this general issue. I will conclude by sketching some generalities that these cases suggest.

I began with an example of empirical success in economics: the quantity theory of money. The example is not typical, for it involves the use of theoretical reasoning to arrive at testable propositions that subsequently, and in ways the originators of the theory could not have foreseen, enjoyed striking empirical confirmation. The nature of the theoretical reasoning involved is quite varied. Some element of "rationality" is involved in all versions, in the sense that units that "ought" not to matter to people are assumed in fact not to matter. More recent models involve formal utility theory in a much more explicit way than did the original versions, as well as explicit notions of market equilibrium. It is interesting that more refined theory has not been found to vindicate or provide a "foundation" for the testable versions of the theory. On the contrary, it has suggested a number of qualifications or possible deviations between theory and observation. We do not "believe in" the theory because it is built up from impeccable axioms about more fundamental aspects of behavior. It is also interesting to note that the theory succeeds empirically only on data

games, as reported, e.g., by Axelrod (1981). In these games, there is a unique Nash equilibrium that can be calculated by "backward induction" from a simple terminal condition, yet subjects often pursue nonequilibrium cooperative strategies (the analogue to exchanging goods for ultimately valueless money in the model of Sec. VI).

that are heavily time averaged. The theory has virtually no ability to account for month-to-month comovements in prices and money.

The experimental work with pigeons and other animal subjects permits an examination of the role of "rationality" at a level that obviously cannot be carried out with aggregative data on entire economies. In this work, the relationship between adaptive and rational behavior is clear: the presumption that behavior is adaptive is built into the experimental design and the interpretation of the results. Working out the implications of utility theory involves calculations (by the experimenters) that we do not believe have counterparts in the mental processes of pigeons.

Experimental work with human subjects in market situations involves adaptive behavior as well, but of an entirely different character. Here subjects are using experience not to trace out their own preferences but to determine how other "players" are likely to react to their own moves. Again, it is possible for the experimenter to calculate certain features of the outcome of this adaptive process theoretically, but these calculations are not a description or a model of the adaptive process itself.

This experimental work bears on the question of the stability of economic equilibria, but the results suggest (as Smith observed in his original paper) processes very different from those treated in received stability theory. More recently, theorists such as Margaret Bray have begun to develop a stability theory that seems to correspond much more closely to the adaptive behavior documented by Smith and other experimentalists.

The models studied by Smith and Bray have unique equilibria. Their results, experimental and theoretical, have the effect of making us feel more comfortable with the predictions of certain theoretical models but do not lead to modifications or improvements in the predictions of these models (though I think they have the potential for doing so). In Section VI, I introduced a well-known example from monetary theory in which there is a continuum of equilibria so that the theory is virtually vacuous. Bray's stability theory selects exactly one of these as stable. In Section VII, I proposed an experiment that would, I believe, select out this same equilibrium as the stable one (although the process might well differ from that proposed by Bray).

Recent theoretical work is making it increasingly clear that the multiplicity of equilibria illustrated in Section VI can arise in a wide variety of situations involving sequential trading, in competitive as well as finite-agent games. All but a few of these equilibria are, I believe, behaviorally uninteresting: They do not describe behavior that collections of adaptively behaving people would ever hit on. I think an appropriate stability theory can be useful in weeding out these uninteresting equilibria, an important application of the Correspondence Principle that Samuelson (1947) proposed long ago. But to be useful, stability theory must be more than simply a fancy way of saying that one does not want to think about certain equilibria. I prefer to view it as an experimentally testable hypothesis, as a special instance of the adaptive laws that we believe govern all human behavior.

Each "point" in figure 1 represents the behavior over a period of 20 years of all the individual households and business firms in a single Latin American country. To a sociologist or an anthropologist, these 16 countries exhibit an enormous variety of quite different cultures. To a political scientist, they cover a range from liberal democracy through military dictatorship. To a psychologist, they consist of millions of individual personalities, with most of those alive at the end of the period not yet born at its beginning. To an economist, they are 16 points lying (more or less) on a theoretically predicted 45-degree line.

To observe that economics is based on a superficial view of individual and social behavior does not, in these circumstances, seem to me to be much of an insight. I think it is exactly this superficiality that gives economics much of the power that it has: its ability to predict human behavior without knowing very much about the makeup and the lives of the people whose behavior we are trying to understand.[9] Yet an ability such as this necessarily has its limits, and I have spent most of this essay on cases that seem to me to lie close to these limits, for this is where they can best be seen and, perhaps, transcended.

References

Axelrod, R. 1981. The emergence of cooperation among egoists. *American Political Science Review* 75:306–18.

Battalio, R. C.; Kagel, J. H.; Rachlin, H.; and Green, L. 1981. Commodity choice behavior with pigeons as subjects. *Journal of Political Economy* 89:67–91.

Blume, L. E.; Bray, M. M.; and Easley, D. 1982. Introduction to the stability of rational expectations equilibrium. *Journal of Economic Theory* 26:313–17.

Blume, L. E., and Easley, D. 1982. Learning to be rational. *Journal of Economic Theory* 26:340–51.

Bray, M. 1982. Learning, estimation and the stability of rational expectations. *Journal of Economic Theory* 26:318–39.

Bray, M. 1983. Convergence to rational expectations equilibrium. In Roman Frydman and Edmund S. Phelps (eds.), *Individual Forecasting and Aggregate Outcomes*. Cambridge: Cambridge University Press.

Hume, D. 1963. *Essays Moral, Political and Literary*. London: Oxford University Press.

Lucas, R. E., Jr. 1978. Asset prices in an exchange economy. *Econometrica* 46:1429–46.

Lucas, R. E., Jr. 1980. Two illustrations of the quantity theory of money. *American Economic Review* 70:1005–14.

Muth, J. F. 1961. Rational expectations and the theory of price movements. *Econometrica* 29:315–35.

McCallum, B. T. 1983. On non-uniqueness in rational expectations models: An attempt at perspective. *Journal of Monetary Economica* 11:139–68.

9. What I am here calling "superficiality" in economic theory is described more fully and more neutrally in Simon (1969).

Samuelson, P. A. 1947. *Foundations of Economic Analysis*. Cambridge, Mass.: Harvard University Press.
Samuelson, P. A. 1958. An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy* 66:467–82.
Simon, H. A. 1969. *The Science of the Artifical*. Cambridge, Mass.: MIT Press.
Smith, V. L. 1962. An experimental study of competitive market behavior. *Journal of Political Economy* 70:111–37.
Smith, V. L. 1982. Microeconomic systems as an experimental science. *American Economic Review* 72:923–55.
Sunder, S. 1985. Unpublished notes. University of Minnesota.
Townsend, R. M. 1978. Market anticipations, rational expectations, and Bayesian analysis. *International Economics Review* 19:481–94.
Vogel, R. C. 1974. The dynamics of inflation in Latin America, 1950–1969. *American Economics Review* 64:102–14.
Wallace, N. 1980. The overlapping generations model of fiat money. In J. H. Karaken and N. Wallace (eds.), *Models of Monetary Economies*. Minneapolis: Federal Reserve Bank of Minneapolis.
Walras, L. 1954. *Elements of Pure Economics*. Translated by William Jaffe. Homewood, Ill.: Irwin.