Evolution, Learning, and Economic Behavior*

REINHARD SELTEN

Department of Economics, University of Bonn, Adenauerallee 24-42, D-5300 Bonn 1, Germany

Received August 7, 1990

This is the entire text of the 1989 Nancy L. Schwartz Lecture delivered by the author at the J. L. Kellogg Graduate School of Management, Northwestern University, Evanston, Illinois. *Journal of Economic Literature* Classification Numbers: 011, 036, 215. © 1991 Academic Press, Inc.

It is doubtful whether I shall be able to meet the high standards set by previous speakers. I shall not prove deep theorems. I shall not present astonishing new results. Instead of this I shall try to catch your attention by a fictitious dialogue. I shall employ the help of imaginary discussants like the "Bayesian" or the "experimentalist." A "chairman" will determine who speaks next, but he shall also make his own remarks.

Chairman: I open the discussion with a question: What do we know about the structure of human economic behavior?

One of our participants, the Bayesian, has signalled his willingness to answer this question. I give the floor to the Bayesian.

Bayesian: As far as economic activities are concerned, it is justified to assume that man is a rational being. Since Savage (1954) simultaneously axiomatized utility and subjective probability we know what rational economic behavior is. Rational economic behavior is the maximization of subjectively expected utility.

Chairman: Among us is an economist. I would like to ask him whether this is the agreed upon opinion in economic theory.

Economist: Yes, to a large extent this is the agreed upon opinion. Most

* We thank Dean Donald P. Jacobs of the J. L. Kellogg Graduate School of Management for permission to publish this lecture.

of microeconomics takes Bayesianism for granted. However, there are exceptions. Some theorists have different views.

As an example let me mention. Allais. Since 1953 he insists that in the evaluation of risky decisions not only the expectation, but also the variance of utility must be taken into account. He has shown that his theory is in better agreement with observed behavior than Bayesianism is.

Chairman: One of our discussants is an experimentalist. His background is in experimental economics. I would like to ask the experimentalist for his opinion on Allais and Bayesian decision theory and their agreement with observed behavior.

Experimentalist: If I understand the opening question correctly, we are here to discuss *human* economic behavior, not the behavior of a mythical hero called "rational man," a mythical hero whose powers of computation and cogitation are unlimited. For this mythical hero it is easy to form consistent probability and preference judgments, but not for ordinary people like you and me.

People are not consistent. I would like to mention just one of many experimental results which show this. In a paper published by Tyszka (1983) he describes an experiment in which the subjects had to make choices from triples, say $\{A, B, C\}$ or $\{A, B, D\}$. Tyszka succeeded in constructing triples such that

95% of the subjects chose A from $\{A, B, C\}$ and

95% of the subjects chose B from $\{A, B, D\}$.

The choice between A and B is influenced by the irrelevant alternatives C and D. This should not be the case, if consistent preference judgments are formed.

The theory of Allais (1953) is in better agreement with behavior than Bayesianism is, but the agreement is only slightly better. The violations of ideal rationality observed in experiments are much more basic than the theory of Allais suggests. Rejecting Bayesianism in favor of Allais's theory is like going up to the top of a skyscraper in order to be nearer to the moon!

Chairman: I agree with the experimentalist. One must make a distinction between normative and descriptive theory. My opening question was meant descriptively. Decision theory and game theory have made tremendous progress on the clarification of the concept of ideal rationality. In this discussion we are not concerned with ideal rationality, but with actual human decision behavior.

However, we should not be too quick in the rejection of the optimization approach as nondescriptive. Evolutionary game theory—started by Maynard Smith and Price (1973)—is successful in biology and biology is thoroughly descriptive. The book by Maynard Smith (1982), *Evolution* and the Theory of Games, provides many examples. Game theory has been created as a theory of rational behavior, but it is now applied to animals and plants.

Among us is a biologist whom we call the "adaptationist" since he strongly beliefs in adaptation in the biological sense. The adaptationist thinks that the principle of fitness maximization is applicable to human behavior. I would like to ask him to explain his views.

Adaptationist: Let me explain what adaptation means in biology. Adaptation means fitness maximization. Fitness is reproductive success—roughly speaking the expected number of offspring in the next generation. Natural selection drives organisms toward fitness maximization. Fitness maximization is a powerful explanatory principle also for human behavior.

However, as far as human decisions are concerned—the same holds for animal decisions by the way—we must be aware of the fact that near to the optimum, selective pressures are weak. Therefore we often observe nearly optimal rules of thumb instead of truly optimal behavior.

The experimentalist has told us about deviations from rationality in experiments. I think that these deviations cannot be of great practical importance. Otherwise natural selection would have eliminated them long ago.

Chairman: Now several discussants want to say something. The next speaker is the Bayesian.

Bayesian: I find the remarks of the adaptationist very interesting. The principle of fitness maximization permits us to construct new kinds of theories in economics. Preferences can now be explained as a result of evolution. People like what is good for their fitness.

Let me make an additional remark on rules of thumb. What seems to be only nearly optimal may be truly optimal, if decision costs are taken into account. I think that with decision costs taken into account, many rules of thumb may turn out to be truly optimal upon closer inspection.

Chairman: I would like to ask the experimentalist for his opinion on rules of thumb and decision costs.

Experimentalist: As many people have pointed out, decision problems tend to become more difficult with decision costs taken into account. If one tries to save decision costs by taking them into account, one may easily end up with higher decision costs.

Adaptationist (interrupting): No! I must interrupt here. Excuse me for doing this. Remember that I think of rules of thumb as inherited. They are not made up on the spot. Evolution has already solved the optimization problem. The decision maker does not have to solve it any more.

Experimentalist: I concede this point. However, I wonder whether all rules of thumb which people use in everyday life are inherited. Sometimes rules of thumb have to be made up on the spot.

I would like to point out that nearly optimal rules of thumb can be far away from truly optimal policies. This may not matter for the decision maker very much, but it may be important for other people. I would like to tell you about an example in which the structure of experimentally observed behavior is dramatically different from that of the optimal policy.

Claus Berg (1973, 1974) published the results of an investment experiment. The subjects could invest in cash or a risky asset. In every period they had to divide total asset into cash and the risky asset. The end was decided by a stopping probability of 1%.

The risky asset yielded a positive interest, say 25%, or a negative interest, say -25%. The percentage was the same in both cases. It was fixed and known to the subject. There was a fixed probability p for the positive interest. This probability was not known to the subject. Profits and losses changed total assets from period to period.

Bayesian decision theory yields the following prediction about changes of the proportion of the risky asset in total assets:

After a positive interest the risky asset proportion is increased. After a negative interest the risky asset proportion is decreased.

This is due to an increase of the posterior probability for a positive interest after a positive interest and to a decrease of this posterior probability after a negative interest. However, in 75% of all cases the subjects changed the risky asset proportion in the direction opposite to the Bayesian prediction!

A closer look at the data revealed the reason for this phenomenon. The subjects tend to form an aspiration level for the total assets they want to obtain at the end of the period. Often this aspiration level is total assets at the beginning of the last period. This has the following consequences: After a loss as much is risked as necessary in order to recuperate the loss; this requires a raise of the risky asset proportion. After a gain not more is risked than has been won; this requires a decrease of the risky asset proportion.

Berg ran computer simulations with an idealized description of the aspiration guided behavior observed in the experiments. He compared the results with those for Bayesian optimal policies starting from beta-distributed priors with expected values near to the true probability. Strangely enough the results for the behavioral theory were often better than those for the optimal policies, particularly in the parameter range of the experiments. Maybe the priors were not appropriate. The success of a Bayesian policy may crucially depend on the prior. Only asymptotically the prior does not matter. But how should we choose the prior?

In any case Berg's theory performs very well. However, its structure of

aspiration guided behavior is very different from that of the optimal Bayesian policy. The Bayesian prediction of the change of the risky asset proportion goes in the wrong direction!

Chairman: I permit the Bayesian to make a short comment.

Bayesian: I only want to say that aspiration guided behavior may be truly optimal, if decision costs are taken into account.

Chairman: Well, this is a possibility, but I would say a remote one. I would like to come back to the statement of the adaptationist: Natural selection drives organisms toward fitness maximization. Among us is a population geneticist. I would like to ask him, what can be said about this statement from the point of view of population genetics?

Population geneticist: Fitness maximization is thought of as the result of a dynamic process of natural selection. We need a justification of this idea within explicit dynamic models. Such models are the subject matter of population genetics.

Ronald Fisher, the great population geneticist and econometrician proved a "fundamental theorem," which under certain conditions shows that natural selection increases fitness until a maximum is reached. Unfortunately Fisher's conditions are rarely satisfied in genetic systems.

In the following I shall rely heavily on a very illuminating paper by Ilan Eshel (1988) which has been made available as a preprint. Maynard Smith and Price (1973) have introduced the concept of an evolutionarily stable strategy as an attempt to give a static description of a dynamically stable result of natural selection in game situations. We may say that evolutionary stability is the generalization of fitness maximization to game situations. Eshel's paper is concerned with the dynamic foundations of evolutionary stability.

We must distinguish two mechanisms of natural selection:

- 1. adaptation of genotype frequencies without mutation,
- 2. gene substitution by mutation.

Mendelian inheritance and selective pressures combine to change the frequencies with which genotypes are represented in the population. This is the process called "adaptation of genotype frequencies without mutation."

Moran (1964) has shown that under realistic conditions adaptation of genotype frequencies without mutation does not necessarily maximize fitness, it may even decrease fitness, until a local minimum is reached. This result holds even without any game interaction. The explanation lies in the combined effects of Mendelian inheritance and linkage. Linkage is the phenomenon that genes near to each other on the same chromosome are likely to be inherited together.

After Moran's result the idea of fitness maximization fell into disrepute

among population geneticists. Only recently a new picture emerged, first in a paper by Eshel and Feldman (1984). They showed for two-locus systems that gene substitution by mutation works in the direction of evolutionary stability. This result has been generalized to an arbitrary number of gene loci by Liberman (1988).

This is very good news for all those who work on evolutionary game theory. Evolutionary game theory now has a more solid dynamic foundation.

However, successful mutants are very rare. Gene substitution by mutation is very slow. Therefore fitness maximization or evolutionary stability can only be expected as a long run equilibrium phenomenon. It is dubious whether any mutations have changed human economic behavior in the relatively short time since the beginning of the dispersion of agriculture about 10,000 years ago. This means that biologically man may still be a hunter and gatherer not very well adapted to the necessity of long run planning. This may be the reason why some Ph.D. dissertations take much longer than planned. In any case it would be silly to expect that man is genetically adapted to modern industrial society.

Chairman: The remarks of the population geneticist throw doubt on the near optimality of human economic behavior. At least we can say that natural selection did not necessarily produce this result. Among us is a naturalist, a man who knows animals and plants. He wants to make a comment.

Naturalist: We rely heavily on the principle of fitness maximization in the explanation of field phenomena. This principle has tremendous explanatory power. However, it must be used with care. Fitness maximization does not work absolutely but only under structural constraints. Let me give you an example in order to make it clear what I mean by structural constraints.

The example is the giraffe. As you all know the giraffe has a very long neck, but only 7 neck bones like every other mammal. This is very inconvenient for the giraffe. It has difficulties lying down to sleep and standing up quickly in danger. It actually sleeps very little. Why did evolution fail to increase the number of neck bones? The answer is simple: *Evolution cannot change many things at once*. A change of the number of neck bones alone would be disastrous. Muscles, nerves, and other things would have to be adjusted. This would require many simultaneous mutations. Therefore the number of neck bones acts as a *structural constraint* on the evolution of the giraffe.

A related problem is that of the correct strategy space in biological game models. Hammerstein and Riechert (1988) have modelled the fighting behavior of *agelonopsis aperta* (a spider). In this model the spiders ignore some useful information like the number of days passed in the

season. From the point of view of fitness maximization the spider's strategy should depend on this information. Hammerstein and Riechert do not permit this in their model. I have no objection! The data justify the restriction of the strategy space. Structural constraints on the spiders' behavior must be working here.

I am convinced of the principle of fitness maximization, but under structural constraints. Fitness maximization alone is not sufficient to explain natural phenomena. A thorough knowledge of nature cannot be replaced by abstract principles.

Chairman: We have heard three reasons against the biological deduction of human economic behavior:

1. the slowness of gene substitution by mutation,

2. the fact that genotype frequency adaptation without mutation does not necessarily optimize fitness,

3. structural constraints.

We have to gain empirical knowledge. We cannot derive human economic behavior from biological principles. I would like to ask the Bayesian what he thinks about this.

Bayesian: Well, maybe the discussion has overemphasized natural selection. Rationality needs training. Small children are not yet rational. Untrained grown-ups still make many mistakes. Maybe we all are not yet sufficiently trained. We have to change this. Bayesian methods should be taught to future executives much more than this is done now. In this way Bayesian methods will become more and more widespread in business and government. Haphazard natural decision behavior will be replaced by superior Bayesian methods. This process of cultural evolution will establish descriptive relevance of Bayesian decision theory, at least where it matters, in business and government.

Chairman: Thank you for mentioning cultural evolution. Up to now this topic has been neglected in our discussion. Two population geneticists, Cavalli-Sforza and Feldman, have created a fascinating mathematical theory of cultural evolution. In 1981 they published a book on the subject. Another useful systematic exposition can be found in the book of Boyd and Richerson (1985). The population geneticist is familiar with this literature. I would like to ask him to describe the basic ideas underlying the mathematical theory of cultural evolution.

Population geneticist: This is a difficult task. I shall give a highly simplified picture. The theory of cultural evolution focuses on cultural traits. A cultural trait is something like the use of a dialect or the adherence to a religion. We think of cultural traits as acquired in the formative years of childhood and adolescence and not changed later in life. Of course this is a simplification. One can model a cultural trait as absent or present, but it is often more adequate to think of a cultural trait as a quantifiable variable measured on a scale. Somebody may more or less adhere to a religion. He may go to church every Sunday or only occasionally. In the following I shall restrict my attention to the case of a quantifiable cultural trait.

Models of cultural evolution with quantifiable cultural traits are similar to models of quantitative inheritance. The theory of quantitative inheritance was initiated by Galton (1889) long before the beginnings of population genetics. Quantitative inheritance theory does not make use of Mendelian genetics. Nevertheless models of quantitative inheritance continue to be useful in animal and plant breeding.

Consider a trait like "height." Quantitative inheritance of height is as follows. Three components determine the height of an individual:

- 1. the parents' average height,
- 2. the population mean,
- 3. a random component.

First a weighted average with fixed coefficients is formed of the first two components and then the random component is added. In this way the height of an individual is determined.

Models of cultural evolution are similar. However, there may be many cultural parents. The biological parents may or may not be among them. The cultural parents exert their influence by teaching or setting an example. They transmit an average in which different cultural parents may have different weights reflecting differences of importance.

I shall now give a sketch of a possible model which is meant to be an illustration and should not be taken too seriously. We shall look at the cultural trait "conformance to work ethics." We have to make an assumption on who becomes a cultural parent of whom. We make the simplest possible assumption. The cultural parents of an individual are a random sample of fixed size taken from the previous generation. In the transmitted average the cultural parents have weights which increase with prestige. Economic success has a positive influence on prestige and conformance to work ethics has a positive influence on economic success. It can be seen how in a model along these lines a high level of conformance to work ethics can evolve in the population, even if in the beginning this level is very low.

I am now at the end of my short exposition. Of course much more could be said about the theory of cultural evolution.

Chairman: I can see that models of this type may be useful in the explanation of economic development. Cultural traits like values, ambitions, and left styles influence economic behavior and thereby economic conditions. Economic conditions exert selective pressure on the cultural

traits. In this way we obtain a feedback loop. Obviously the application of cultural evolution theory to economics offers some interesting possibilities.

However, we should not forget evolutionary theories which have their origin within economies. Already Schumpeter (1934) created an evolutionary theory of innovation and imitation. Nelson and Winter (1982) present formal models in their book on "an evolutionary theory of economic change." The economist is familiar with this work and I would like to ask him to explain the approach of Nelson and Winter.

Economist: Nelson and Winter focus on firms rather than individuals. In their models firms do not maximize profits. Instead of this they adapt to success and failure in a trial and error fashion. I think that the best way to explain the approach of Nelson and Winter is with a short sketch of a particular model described by Sidney Winter (1971).

I shall sketch a slightly simplified version of the model. A finite number of goods is produced by many firms. Production methods connect inputs and outputs by fixed proportions. Inputs and outputs differ from method to method. An output of one firm may be an input of another.

The market as a whole ends up with a surplus for some commodities and with a deficit for others. We may think of surplusses as sold to consumers and of deficits as imported. The surplusses and deficits determine the prices by a relationship technically named "inverse demand function."

I now come to the behavioral assumptions. After a profit, a firm expands by one unit. After a loss a firm contracts by one unit and in addition to this starts a search for a new production method which is found and adopted with a positive probability. A firm whose production method yields zero profits does not change anything. Essentially the same rules apply to potential firms which do not produce anything. These firms also have a production method. They enter in case of profitability. They search for a new production method if the present one is not profitable at current prices.

Sidney Winter has shown that under appropriate assumptions on the inverse demand function the stochastic process defined by the model converges to an absorbing state with the properties of a competitive equilibrium. In this way he provided a new foundation to the theory of competitive equilibrium without profit maximization assumptions.

I admire the work of Nelson and Winter and in particular the model which I just described. However, I cannot see a close analogy to biological and cultural evolution. Firms do not reproduce and do not die. There is no cultural transmission from firm to firm—at least not in this model. I have no objection! Nelson and Winter have created an evolutionary theory in its own right, much better adapted to economics than any literal translation from biology could be. *Chairman*: Maybe there is a closer analogy than one thinks at first glance. It is necessary to change the perspective. Not the firms but the production methods are the animals under selective pressure. The behavior of the firms is the environment of the production methods. Production methods are born and may die.

In equilibrium all active production methods have zero profitability. This is analogous to biological models with asexual reproduction where in equilibrium all surviving genotypes have fitness 1. Obviously profitability in one case has the same role as fitness in the other case. I would like to know whether a similar analogy can be established between cultural and biological evolution. Is it possible to define a cultural fitness? Maybe the population geneticist can answer this question.

Population geneticist: The idea of cultural fitness sometimes appears in the literature, for example, in the book of Cavalli-Sforza and Feldman (1981), but I cannot remember having seen a precise definition. Nevertheless it seems to be easy to give a meaning to the term. Cultural fitness could be defined as a measure of the expected influence exerted as a cultural parent on the next generation. However, it is unclear whether a cultural fitness concept could be useful.

In the models I described, cultural inheritance is similar to quantitative inheritance in biology and it is unclear whether the biological fitness concept is useful in models of quantitative inheritance, unless very special assumptions are made. Two parents of optimal height usually have children of nonoptimal height—due to the random component and maybe the influence of the population mean. In equilibrium not everybody will be of optimal height. The equilibrium height distribution has a positive variance. Generally there will not even be a monotonic relationship between fitness and representation in the equilibrium population. Only under very special assumptions on the way in which fitness depends on height do we obtain such a monotonic relationship. This indicates that a cultural fitness concept is not very useful in models of cultural evolution like those I have described. This may be different for other types of models.

Chairman: The adaptationist wants to make a remark.

Adaptationist: I want to say that cultural evolution tends to the maximization of biological fitness. We do not need a concept of cultural fitness. Mechanisms of cultural evolution are shaped by biological evolution. There must be fitness advantages for transmitters and receivers. Consider the following example:

Parents teach their children which fruits are edible and which are poisonous.

It is advantageous for the children to accept the transmission and it is advantageous for the parents to transmit—after all the children are their fitness. The shakers provide another example. The shakers have the cultural trait of not having children. Therefore they are dying out. In the long run a cultural trait which reduces biological fitness cannot persist.

Chairman: Is this really true? I would like to hear the opinion of the population geneticist.

Population Geneticist: My exposition of cultural evolution was highly simplified. I did not talk about the interaction of cultural and biological evolution. Actually this interaction is emphasized in the literature. Cavalli-Sforza and Feldman (1981) present a model in which family size is the cultural trait under consideration. In spite of the interaction with biological evolution, in this model cultural evolution stabilizes a small family size.

Adaptationist (interrupting): No! No! I must interrupt here. Excuse me. The assumptions of the model must be wrong. In the long run natural selection would favor a tendency not to accept the cultural transmission of a small family size. If a bigger family size offers a fitness advantage, a small family size cannot persist.

Population geneticist: Maybe in the very long run this is true, but not on the time scale of cultural evolution processes in recorded history. In the explanation of such processes we can safely ignore the interaction with biological evolution, at least if one avoids extreme examples like the shakers.

In my simplified exposition of cultural evolution I also omitted another kind of interaction—the interaction of cultural evolution with individual learning. It seems to me that the interaction with individual learning is much more important than the interaction with natural selection. Actually the interaction with individual learning is also emphasized in the literature. Boyd and Richerson (1985) describe the psychological learning model by Bush and Mosteller (1955) and they also explain Bayesian updating. However, they do not make any explicit use of these modelling possibilities. They only model the effect of learning. The precise mechanism is left open.

Boyd and Richerson assume that learning is guided by a criterion of success, for example, money income. Moreover they assume that at any point of time there is exactly one value of the trait under consideration, which is optimal with respect to the criterion of success. The value of an individual's trait is influenced by three components:

- 1. a culturally inherited value,
- 2. the optimal value,
- 3. a random error added to the optimal value.

The idea is that the optimal value is not correctly perceived. This is expressed by the random error. Learning shifts the culturally inherited value in the direction of the misperceived optimal value. The value of the trait is a weighted average with fixed coefficients for the culturally inherited value and the optimal value modified by the random error. The optimal value is endogenous rather than exogenous; it may depend on the distribution of the trait.

Chairman: This reminds me of Arrow's (1962) theory of learning by doing. In this theory, too, only the effect of learning is modelled. It is assumed that experience results in a downward shift of the cost curve, but how the firms learn to save costs is left open.

If we want to describe economic behavior, it is not enough to model the effects of learning. We must ask the question: What is the structure of learning? Maybe the experimentalist can help us to answer this question.

Experimentalist: It is necessary to distinguish at least three kinds of learning:

- 1. rote learning,
- 2. imitation,
- 3. belief learning.

In rote learning success and failure directly influence the choice probabilities. Rote learning does not require any insight into the situation. It is based on a general trust in the stability of the environment: What was good yesterday will be good today.

Imitation of successful others is similar to rote learning with the difference being that it is the success of others which directly influences choice probabilities.

Belief learning is very different. Here experiences strengthen or weaken beliefs. Belief learning has only an indirect influence on behavior.

We do not know very much about the structure of learning, but we know more about rote learning than belief learning. Up to now we do not yet have a sufficient understanding of belief learning. Therefore I shall restrict my comments to rote learning.

Psychological learning models like the one by Bush and Mosteller (1955) describe rote learning. These models have only two possible reward levels, success and failure. This is reasonable for some animal experiments: The rat either finds food in the maze or it does not find food. For the description of learning by an economic agent motivated by profits one needs a continuum of possible reward levels. In his book from 1983 John Cross has generalized the model of Bush and Mosteller to a continuum of reward levels. As far as I know the generalized model has not been confronted with data.

I would like to tell you about an experimental investigation by Malawski described in his unpublished Ph.D. dissertation (1989). He looked at game situations with minimal information. The subjects did not even know that they played games. They only knew that they had to choose one of several alternatives, say A, B, C. After each choice they were informed about their own payoff (sometimes they obtained additional information which will not be described here). They experienced the same decision situation about 70 times.

Malawski developed a theory of "learning through aspiration," which is in good agreement with his data.

His theory assumes that aspiration levels on payoffs are formed in the beginning of the session (the first 12 decisions). After this initial phase only two responses to the experience made last time are possible:

- (a) the same choice as last time or
- (b) a randomly picked different alternative.

Depending on the situation one of these responses is "normal" and the other one is "exceptional." The normal response is taken more often.

Under the condition that the aspiration level has been reached or surpassed last time, the normal response is the same choice as last time. Under the condition that the aspiration level has not been reached, the normal response is a randomly picked different alternative.

The normal response is taken with a probability $p > \frac{1}{2}$ and the exceptional response is taken with probability 1 - p. The probability p is a parameter which varies from subject to subject.

Once more the good fit of this theory shows the importance of aspiration levels in economic behavior. Of course Malawski's theory is not yet firmly established. More experiments have to be made, but his results look very promising.

Chairman: I would like to hear the opinion of the economist on the structure of learning.

Economist: I am surprised about the remark of the experimentalist, that we cannot say very much about belief learning. After all Bayesian updating is a plausible model of belief learning. Bayesian updating can be descriptively right, even if subjectively expected utility maximization is descriptively wrong. Somebody who does not maximize utility can still adjust his subjective probabilities by Bayes' rule.

We can replace subjectively expected utility maximization by alternative theories in the literature, for example, prospect theory by Kahneman and Tversky (1984) or regret theory by Loomes and Sugden (1982), Bell (1982), and Fishburn (1982). However, I cannot see any alternative to Bayesian updating. The experimentalist did not even sketch an alternative model of belief learning.

Chairman: The experimentalist should answer this remark. Is it really

true that we have no alternative to Bayesian updating as a model of belief learning?

Experimentalist: The trouble with Bayes' rule is that people do not obey it. It is, for example, well known that people overvalue the information content of small samples. In this connection Tversky and Kahneman speak of a "law of small numbers" (1982a).

I have been accused of not having sketched an alternative model of belief learning. Therefore I shall do this now. What I shall tell you is highly speculative, even if there is some experimental support. I refer to an oligopoly experiment described by Selten (1967).

If one wants to model belief learning, one first has to model a belief system. Bayes' rule operates on a belief system which is a probability distribution over all possible states of the world. However, belief systems of human decision makers may have a completely different structure. In my sketch of an alternative model of belief learning, beliefs are formed on causal links like

Advertising increases sales.

The belief structure has the form of a causal diagram composed of such causal links. The term causal diagram has been used by Selten (1967). Later very similar structures were called "cognitive maps" by Axelrod (1976). Axelrod has done interesting empirical research on cognitive maps of politicians expressed in speeches and writings.

A causal diagram shows chains which connect a decision variable like advertising with a goal variable like profits. Suppose that there are two causal chains in the diagram:

- 1. Advertising increases sales and sales increase profits.
- 2. Advertising increases cost and cost decreases profit.

Subjects reason qualitatively on the basis of such chains. The first chain is an argument for more advertising and the second chain is an argument for less advertising. Such conflicts are resolved by judgments on the relative importance of causal links. Suppose that the causal link from advertising to cost is judged to be the least important one. On the basis of this judgment the second chain is neglected and a decision to increase advertising is based on the first chain. Of course, this determines only the direction of change. The amount of change has to be determined in some other way.

Belief learning exerts its influence on importance judgments. Experience may show that the influence of advertising on cost is more important than the influence of advertising on sales. The mechanism of belief learning can be modelled in a fashion similar to that of the mechanism of rote learning. *Chairman*: These remarks show that belief learning is closely connected to boundedly rational reasoning. Not only belief systems must be modelled but also their use in reasoning. We need to know more about bounded rationality. In my view the development of a theory of bounded rationality is one of the most important tasks of economics in our time. Bayesian updating has been mentioned favorably and unfavorably. I think that the Bayesian should comment on Bayes' rule.

Bayesian: I would rather like to comment on bounded rationality. I may have no opportunity to do this later. Herbert Simon (1957) introduced the idea of satisficing. His view of bounded rationality did not exert a strong influence on economic theory. The work of Herbert Simon inspired the book by Cyert and March (1963) on the "behavioral theory of the firm." Many people were very impressed by this book, in particular by the surprising empirical success of the model which describes the behavior of a department store manager with high predictive accuracy. However, the book did not start a revolution of economic theory. What is the reason for this? Let me give my answer to this question.

A typical piece of research in the behavioral theory of the firm is a simulation study, based on a complex model, sometimes with hundreds of parameters and with many behavioral ad hoc assumptions. Generally no clear conclusion can be drawn from such simulation studies. What we see here is a *theory without theorems*. A theory without theorems cannot succeed.

In the behavioral theory of the firm and the evolutionary approach by Nelson and Winter, behavior is described by ad hoc assumptions, which vary from model to model. This is very unsatisfactory. Economic theory needs a description of economic behavior based on a few general principles which can be applied to every conceivable decision situation. Bayesian decision theory meets this requirement. Bayesian decision theory should not be thrown away in favor of ad hoc explanations of experimental phenomena.

Experiments are often too quickly interpreted as evidence against Bayesian decision theory. I would like to mention the example of the finitely repeated prisoner's dilemma. This game has only one equilibrium outcome, namely noncooperation in every period. Nevertheless one observes cooperation until shortly before the end. This seems to refute the rationality assumptions of game theory. However, a rational explanation has been given by Kreps *et al.* (1982). They introduced a small amount of incomplete information on payoffs of the other players. In the slightly modified game the usual pattern of behavior is an equilibrium outcome.

Even the limits of computational capabilities permit a Bayesian treatment. Repeated games with limited memory or limited complexity have been analyzed by Neyman (1985), Aumann and Sorin (1989), Kalai and Stanford (1987), and others. Another approach to problems connected to bounded rationality is the relaxation of common knowledge assumptions explored by Neyman. We see here the beginnings of a Bayesian theory of bounded rationality. It is not necessary to construct a theory of bounded rationality outside the Bayesian framework. It is much more fruitful to do this within the Bayesian framework.

Chairman: I would like to ask the experimentalist what he thinks about Bayesian bounded rationality.

Experimentalist: Let me first make a comment on the paper by Kreps *et al.* (1982). Game theoretically their work is very interesting, but behaviorally they miss the mark! Selten and Stoecker (1986) describe an experiment where each subject plays 25 supergames of 10 periods each against anonymous opponents changing from supergame to supergame. In these experiments the typical pattern of cooperation until shortly before the end did not emerge in the first supergame, but only after a considerable amount of learning. In the beginning behavior is chaotic. Only slowly cooperation is learned and after cooperation the end effect. According to Kreps *et al.*, the typical pattern is due to thinking rather than learning and therefore should emerge immediately. In their theory there is no room for chaotic behavior in the beginning and for slow learning afterward.

I now want to comment on infinite supergames with restricted memory. In such games only the operating memory is restricted, or in other words, the storage space available for the execution of a strategy. The computational capabilities for the analysis of the game remain unrestricted. In fact, the analysis of the game tends to become more difficult by constraints on the operating memory. I cannot see any contribution to a theory of bounded rationality in this kind of work.

Let me now say something about common knowledge or the lack of it. Consider a chain of the following kind:

I know, that he knows, that I know, that he knows,

Roughly speaking, common knowledge means, that such chains can be continued indefinitely. Does it really matter in practical decision situations whether I have common knowledge or whether I have to break off such chains after stage 4? I do not think so. As far as human decision behavior is concerned I dare say: A lack of common knowledge is not important; what often is important is a very common lack of knowledge.

I highly appreciate the behavioral theory of the firm and the evolutionary approach by Nelson and Winter. I do not accept the criticism against the use of ad hoc assumptions. Look at human anatomy and physiology: bones, muscles, nerves, and so on. Human anatomy and physiology cannot be derived from a few general principles.

Let me also say something else in defense of ad hoc assumptions.

Experiments show that human behavior is ad hoc. Different principles are applied to different decision tasks. Case distinctions determine which principles are used where. Successful explanations of experimental phenomena have been built up along these lines, for example, the theory of equal division payoff bounds for three-person games in characteristic function (Selten, 1987). Let me conclude my comments with a final remark: It is better to make many empirically supported ad hoc assumptions, than to rely on a few unrealistic principles of great generality and elegance.

Chairman: I would like to ask the economist whether he thinks that economic theory should abolish the optimization approach in favor of a more realistic description of economic behavior.

Economist: Many economic theorists are uneasy with the usual exaggerated rationality assumptions—but they continue to use them. They do not see a clear alternative. In recent years the interest in experimental economics has increased tremendously. This offers a hope for a new foundation of microeconomics. However, as long as this new foundation has not yet been established, we have to go on relying on exaggerated rationality assumptions. I do not think that present-day microeconomics will become completely obsolete. Market experiments by Smith, Plott, and others reviewed in the literature (Smith, 1980; Plott, 1982) confirm competitive equilibrium theory. What is now derived as a result of optimization may later be explained as a result of learning.

Bayesianism may be wrong descriptively, but this does not touch its great normative significance. Moreover as has been pointed out earlier in the discussion, teaching of Bayesian methods in universities will increase their use in business and government and thereby establish descriptive relevance for Bayesian decision theory.

Chairman: Yes, we did not yet sufficiently discuss the idea that in the future economic behavior will become more Bayesian than it is now due to the influence of teaching. I would like to ask the experimentalist what he thinks of the prospects of a cultural evolution toward a widespread use of Bayesian methods in business and government.

Experimentalist: The application of Bayesian methods makes sense in special contexts. For example, a life insurance company may adopt a utility function for its total assets; subjective probabilities may be based on actuarial tables. However, a general use of Bayesian methods meets serious difficulties. Subjective probabilities and utilities are needed as inputs. Usually these inputs are not readily available.

There is no probability and utility book in the brain which can be looked up like a telephone directory. Probability and preference judgments require information processing in the brain. They are outputs rather than inputs. There is no reason to suppose that information processing in the brain yields consistent probability and preference judgments. There is much experimental evidence to the contrary. Let me introduce an example concerning probability judgments, the conjunction effect, described by Tversky and Kahneman (1982b). They asked their subjects to rank a number of statements on some future events with respect to likelihood. One of the events was a tennis match involving Borg. Among the statements were the following two:

Statement A: Borg will lose the first set. Statement B: Borg will lose the first set, but he will win the match.

Subjects tend to judge statement B as more likely than statement A, in spite of the fact that statement B describes a subcase of statement A. This shows that probability judgments do not even have one of the most basic consistency properties, namely monotonicity of the probability measure with respect to set inclusion.

The phenomenon is a "representativeness effect." The statement "Borg will loose the first set" is not representative of the image of Borg in the mind of the subjects. Borg was a winner, not a loser. The additional detail "but he will win the match" is representative and therefore improves the impression of credibility.

Preference judgments are as inconsistent and unreliable as probability judgments. Martin Weber *et al.* (1988) have published an experimental investigation of "dimension splitting in multiattribute utility measurements." They show that the weight of an attribute is increased if it is split into several subattributes. This is important, since multiattribute utility measurement is recommended as an instrument of decision aid by Bayesian decision theorists interested in practical applications. We see here that the result of the method heavily depends on the way in which the problem is presented to those who have to make the preference judgments.

We can conclude: Normative Bayesianism is dubious in view of the unreliability of its inputs. If you ask people to be consistent, you ask for too much. Imagine a normative theory of sports which commands: Athlete, jump 100 m high!—It cannot be done.

I do not see an unavoidable cultural evolution toward a more and more widespread use of Bayesian methods in business and government.

Chairman: It seems to be necessary to make a distinction between a *practical normative theory* and an *ideal normative theory*. A practical normative theory can be used in order to help people to improve their decisions. Ideal normative theory has the purpose of clarifying the concept of rationality independent of the limitations of real persons. Even if Bayesianism may fail as a practical normative theory it still remains an ideal normative theory of great philosophical importance.

We now must come to the end of the discussion. I shall try to summarize the results as I see them. I opened the discussion with a question: What do we know about the structure of human economic behavior? I must admit that the answer is disappointing. We know very little.

We know that Bayesian decision theory is not a realistic description of human economic behavior. There is ample evidence for this, but we cannot be satisfied with negative knowledge—knowledge about what human behavior fails to be. We need more positive knowledge on the structure of human behavior. We need quantitative theories of bounded rationality, supported by experimental evidence, which can be used in economic modelling as an alternative to exaggerated rationality assumptions.

We have identified a hierarchy of dynamic processes which shape economic behavior. I name these processes in the order of increasing speed:

- 1. (the slowest process) gene substitution by mutation,
- 2. adaptation of genotype frequencies without mutation,
- 3. cultural transmission from generation to generation,
- 4. learning (including imitation).

The speed differences are so great that for many purposes an adiabatic approximation seems to be justified. Adiabatic approximation means that if we look at one of the four processes, results of slower processes can be taken as fixed and quicker processes can be assumed to reach equilibrium instantly.

Learning, the quickest process, is the most important one for economics. Day to day price movements on the stock exchange and other competitive processes involve learning and imitation. For slower dynamic phenomena like economic development, cultural transmission from generation to generation is also important.

The two processes of biological evolution have shaped the inherited components of economic behavior. Gene substitution by mutation maximizes fitness, but slowly and under structural constraints. Adaptation of genotype frequencies without mutation is nonoptimizing.

It is interesting to speculate on the evolution of behavioral tendencies. One may, for example, construct theories on the influence of prehistoric or even prehuman environmental factors on mechanisms of cultural evolution. However, such speculations, as interesting as they may be, are no substitute for empirical research. It makes no sense to speculate on the evolution of unicorns unless unicorns have been found in nature. Biological theory cannot be used as an instrument to discover facts by armchair reasoning.

We have to do empirical research if we want to gain knowledge on the structure of human economic behavior. In order to replace unrealistic rationality assumptions, we need theories of bounded rationality. As I have already said, we need quantitative theories which can replace the usual rationality assumptions in economic models.

In the near future theories of limited range which apply to restricted areas of experimental research have to be expanded. A number of such theories already can be found in the literature. Some of them have been mentioned in the discussion. It is hoped that eventually many theories of limited range will grow together and evolve into a comprehensive picture of the structure of human economic behavior. Only painstaking experimental research can bring us nearer to this goal.

I close the discussion now. Whoever wants to add something must do this in private conversations.

References

- ALLAIS, M. (1953). "Le Comportement de l'Homme Rationnel Devant le Risque Critique des Postulates et Axiomes de l'Ecole Americaine," *Econometrica* 21, 503-546.
- ARROW, K. J. (1962). "The Economic Implications of Learning by Doing," Rev. Econ. Stud. 29, 155-173.
- AUMANN, R. I., AND SORIN, S. (1989). "Cooperation and Bounded Recall," *Games Econ.* Behav. 1, 5–39.
- AXELROD, R. (1976). Structure of Decision: The Cognitive Maps of Political Elites. Princeton, NJ: Princeton Univ. Press.
- BELL, D. E. (1982). "Regret in Decision Making under Uncertainty," Oper. Res. 30, 961–981.
- BERG, C. C. (1973). "Individuelle Entscheidungsprozesse: Laborexperimente und Computersimulation," Wiesbaden.
- BERG, C. C. (1974). "Individual Decision Concerning the Allocation of Resources for Projects with Uncertain Consequences," Manag. Sci. 21, 98-105.
- BOYD, R., AND RICHERSON, P. I. (1985). Culture and the Evolutionary Process. Chicago/ London: Univ. of Chicago Press.
- BUSH, R. R., AND MOSTELLER, F. (1955). Stochastic Models of Learning. New York: Wiley.
- CAVALLI-SFORZA, L. L., AND FELDMAN, M. W. (1981). Cultural Transmission and Evolution: A Quantitative Approach. Princeton, NJ: Princeton Univ. Press.
- CROSS, J. G. (1983). A Theory of Adaptive Economic Behavior. Cambridge/London/New York: Cambridge Univ. Press.
- CYERT, R. M., AND MARCH, J. G. (1963). A Behavioral Theory of the Firm. Englewood Cliffs, NJ: Prentice-Hall.
- ESHEL, I. (1988). "Game Theory and Population Dynamics in Complex Genetical Systems: The Role of Sex in Short Term and in Long Term Evolution," Working Paper No. 24, Game Theory in the Behavioral Sciences, ZiF, Universitat Bielefeld.
- ESHEL, I., AND FELDMAN, W. M. (1984). "Initial Increase of New Mutants and Some Continuity Properties of ESS in Two Locus Systems," Amer. Naturalist **124**(5), 631-640.

- FISHBURN, P. C. (1982). "Nontransitive Measurable Utility," J. Math. Psychol. 26, 31-67.
- GALTON, F. (1989). Natural Inheritance. New York: Macmillan Co.
- HAMMERSTEIN, P., AND RIECHERT, S. E. (1988). "Payoffs and Strategies in Territorial Contests of Two Ecotypes of Spider, Agelonopsis Aperta," Evol. Ecol. 2, 115–138.
- KAHNEMAN, D., AND TVERSKY, A. (1984). "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, 263–291.
- KALAI, E., AND STANFORD, W. (1987). "Finite Rationality and Interpersonal Complexity in Finitely Repeated Games," *Econometrica* 56, 397–410.
- KREPS, D., MILGROM, P., ROBERTS, J., AND WILSON, R. (1982). "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," J. Econ. Theory 27, 245–252.
- LIVERMAN, V. (1988). "External Stability and ESS: Criteria for Initial Increase of New Mutant Allele," Math. Biol. 26, 477-485.
- LOOMES, G., AND SUGDEN, R. (1982). "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Econ. J.* **92**, 805–824.
- MALAWSKI, M. (1989). "Some Learning Processes in Population Games," Unpublished dissertation, University of Bonn.
- MAYNARD SMITH, J. (1982). Evolution and the Theory of Games. Cambridge: Cambridge Univ. Press.
- MAYNARD SMITH, J., AND PRICE, G. R. (1973). "The Logic of Animal Conflict," Nature 246, 15–18.
- MORAN, P. A. P. (1964). "On the Nonexistence of Adaptive Topographies," Amer. Human Genet. 27, 343-383.
- NELSON, R., AND WINTER, S. G. (1982). An Evolutionary Theory of Economic Change. Cambridge, MA/London: Belknap Press of Harvard Univ. Press.
- NEYMAN, A. (1985). "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma," Manuscript.
- PLOTT, C. R. (1982). "Industrial Organization Theory and Experimental Economics," J. Econ. Lit. 20, 1485–1527.
- SAVAGE, L. I. (1954). The Foundation of Statistics. New York: Wiley.
- SCHUMPETER, I. A. (1934). The Theory of Economic Development. Cambridge, MA: Harvard Univ. Press.
- SELTEN, R. (1967). "Invetitionsverhalten im Oligopolexperiment," in Beiträge zur experimentellen Wirtschaftsforschung (H. Sauermann, Ed.), Vol. I, pp. 60-102. Tübingen: Mohr (Paul Siebeck).
- SELTEN, R. (1987). "Equity and Coalition Bargaining," in Experimentation in Economics (A. Roth, Ed.). Cambridge/New york: Cambridge Univ. Press.
- SELTEN, R., AND STOECKER, R. (1986). "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames," J. Econ. Behav. Organ. 7, 47-70.
- SIMON, H. A. (1957). Models of Man. New York: Wiley.
- SMITH, V. L. (1980). "Microeconomic Systems as a Science," Amer. Econ. Rev. 5, 923-955.
- TVERSKY, A., AND KAHNEMAN, D. (1982a). "Belief in the Law of Small Numbers," in Judgment and Uncertainty, Heuristics and Biases (D. Kahneman, P. Slovic, and A. Tversky, Eds.), pp. 23-31. Cambridge/London/New York: Cambridge Univ. Press.
- TVERSKY, A., AND KAHNEMAN, D. (1982b). "Judgements of and by Representativeness," in Judgment and Uncertainty, Heuristics and Biases (D. Kahneman, P. Slovic, and A. Tversky, Eds.), pp. 84–98. Cambridge/London/New York: Cambridge Univ. Press.

- TYSZKA, T. (1983). "Contextual Multiattribute Decision Rules," in Human Decision Making (L. Sjöberg, T. Tyszka, and J. Wise, Eds.). Bodafors: Doxa.
- WEBER, M., EISENFÜHR, F., AND VON WINTERFELDT, D. (1988). "The Effects of Splitting Attributes on Weights in Multiattribute Utility Measurement," *Manag. Sci.* 34, 432-445.
- WINTER, S. G. (1971). "Satisficing, Selection and the Innovating Remnant," Quart. J. Econ. 85, 237-261.