Costly and Bounded Rationality in Individual and Team Decision-making

Roy Radner
(Henry Kaufman Management Center, Department of Information Systems, Leonard N. Stern School of Business, New York University, 44 West Fourth Street, Suite 9-170, New York, NY 10012-1126, USA.
Email: rradner@stern.nyu.edu)

The 'Savage paradigm' of rational decision-making under uncertainty has become the dominant model of human behavior in mainstream economics and game theory. However, under the rubric of 'bounded-rationality', this model has been criticized as inadequate from both normative and descriptive viewpoints. This paper sketches the historical roots and some current developments of this movement, distinguishing between attempts to extend the Savage paradigm ('costly rationality') and the need for more radical departures ('truly bounded rationality').

1. *Introduction*

Some kind of model of rational decision-making is at the base of most current economic analysis. This is certainly true of microeconomic theory and, in particular, recent theoretical research on the economics of organization. At the same time, there is a growing unease with the mainstream models of *Homo economicus*, and a small but growing body of economic theory in response to it. Such unease was already voiced at the beginning of the 20th century (if not before), but the work of Jacob Marschak and Herbert Simon provided the stimuli for a more intense level of activity. Interestingly enough, their work appeared almost at the same time that the publications of von Neumann and Morgenstern, L. J. Savage and John Nash established the foundations of the current rational-choice school of the theory of economic organization.

In this essay I shall sketch some of the roots of the current revolt against the mainstream, and also some of its branches. I shall confine myself to the roots and branches that relate most closely to what I call the 'Savage paradigm' of rational decision in the face of uncertainty. Thus I shall leave the

Industrial and Corporate Change Volume 9 Number 4 2000

game-theoretic aspects of organization theory aside.[1] I shall also confine myself to the economic theory literature, thus neglecting a large literature using a more sociological approach. Evidence of bounded rationality is not difficult to cite, at least the 'boundedness' part. Fortunately, I can refer the reader to the excellent review article by John Conlisk (1996).

My title, 'Costly and Bounded Rationality', reflects the two rather different approaches of Marschak and Simon. Marschak's work emphasized that the Savage paradigm could be extended to take account of costs and constraints associated with information acquisition and processing in organizations, without abandoning the notion of optimizing behavior. This approach found its most detailed expression in the theory of teams. Marschak was also interested in modeling the phenomenon of inconsistency in decision-makers' choices, and was one of the founders of the theory of 'stochastic choice'. However, even there he was interested in how, and to what degree, the real choices of decision-makers approximated, at least statistically, those of a fully rational person.

Simon was more concerned with behavior that could not so readily be interpreted, if at all, as 'optimizing'. As far as I know, he coined the term 'bounded rationality', and developed the concept in its relation to economics, psychology and artificial intelligence. However, in some of his publications he apparently considered bounded rationality to be a broader concept, subsuming costly rationality as a particular case. In order to avoid terminological disputes, I shall sometimes call the narrower concept 'truly bounded rationality'.

In this paper, I start with the theory of *unbounded* rationality of decision-making in the face of uncertainty, and with Savage's model as the archetype of such a theory. There is no doubt that Savage's purpose was *normative*. Near the beginning of *The Foundations of Statistics* he states:

> I am about to build up a highly idealized theory of the behavior of a 'rational' person with respect to decisions. In doing so I will, of course, have to ask you to agree with me that such and such maxims of behavior are 'rational'. In so far as 'rational' means logical, there is no live question; and if I ask your leave at all, it is only as a matter of form. But our person is going to have to make up his mind in situations in which criteria beyond the ordinary ones of logic will be necessary. So, when certain maxims are presented for your consideration, you must ask yourself whether you try to

[1] See, however, some inflammatory remarks on the latter topic in Radner (1997).

> behave in accordance with them, or, to put it differently, how would you react if you noticed yourself violating them. (Savage, 1954, p. 7)

This quote describes a philosophical exercise. If there is any empiricism here, it is armchair empiricism about what 'reasonable and thoughtful' persons might think.

My paper is written in the same spirit, as far as I can push it. But theorists of *costly and bounded* rationality face a dilemma. On the one hand, it is not rational to prescribe behavior that is hopelessly unrealistic. In particular, it is not realistic to expect even expert decision makers with computers to be fully logical (see Section 5.5). On the other hand, knowledge about cognitive costs and bounds is empirical. Nevertheless, my ultimate purpose here is normative. Given what we know about the cognitive costs and limits of decision-making in the face of uncertainty, what might it mean to make decisions rationally, if indeed there is any meaning at all?

One 'empirical' approach to rationality is to try to study how 'successful' decision-makers behave. I fear that this approach is fraught with difficulties, as Savage's own theory would explain. A discussion of this issue would fill a gap in the paper. For example, in an uncertain world, good decisions can often lead to bad consequences, and vice versa. Furthermore, in a world in which many uncertainties are never resolved, it is often impossible to tell *ex post* why the consequences of a decision were what they were. Recall that Savage's goal was to lay the foundations of inductive (i.e. statistical) inference, and in particular to rationalize many current statistical methods. At best, his model describes a set of 'rational' procedures for using information to make decisions. In this sense, his was a theory of *procedural rationality*.

I am, of course, aware that many economists use the Savage model, or something like it, as a *positive* theory to explain observed economic behavior and institutions. (This use of the model has even made inroads into other social sciences.) Mainstream positive economics holds on tenaciously to the Savage model, even though it is patently false if taken literally. However, this is not an issue that I address in this paper.

The term 'procedural rationality' brings with it connotations of dynamics and change. As explained in Section 3, in the Savage paradigm a decision-maker's preferences among actions and strategies are determined jointly by his or her 'tastes and beliefs'. The focus is on the dynamics of how a rational decision-maker revises his or her beliefs about the relative likelihoods of relevant events, in the light of new observations about the external world. However, *the decision-maker's underlying model of the world is fixed*, and the successive revisions of beliefs proceeds according to the calculus of conditional probability. (Technically, the underlying probability space remains unchanged.) In reality, decision-makers typically change their underlying models of the world from time to time, in response to stimuli other than observations consistent with the current underlying model. Whether there is a 'rational' way to do this in the spirit of the original Savage paradigm is an important and challenging question. I believe that there may be such a way, and in Section 5.4 I explore the idea in the context of a particular example.

The decision-maker's tastes are also subject to change. (In the jargon of economics, the decision-maker's tastes are represented by his or her *utility function*.) Some apparent changes in taste might be explained by the acquisition of new information about the consequences of action in a way that preserves the underlying model of the decision-maker, but there are often more deep-seated changes. For example, a recurrent theme in the management literature is the importance of the formation and maintenance of *values* in the firm (e.g. Mintzberg, 1989; Peters and Waterman, 1982). However, a serious examination of what has been written about the psychology and social psychology of human motivation is beyond the scope of this paper.

In Section 2, I sketch some of the roots of the ideas of costly and bounded rationality. In order to explain more clearly the branches that have subsequently developed, I found it desirable to sketch the elements of the Savage paradigm, which I do in Section 3. The topics of costly and truly bounded rationality follow in Sections 4 and 5 respectively. In particular, Section 5.4 sketches a proposed analysis of 'rational' model revision. Finally, in Section 6, I speculate on the senses (if at all) in which a 'boundedly rational' decision-maker can be said to be 'rational'. As I have already mentioned, my discussion is almost entirely confined to individual and team decision-making, as distinct from game-theoretic issues.

2. *Some Roots*

2.1 Uncertainty

Discussions of rational decision-making—unbounded and otherwise—have been closely tied to the consideration of uncertainty. The very beginnings of formal probability theory were in part stimulated by questions of how to act rationally in playing games of chance with cards and dice (Hald, 1990).

Kenneth Arrow began his 1951 article, 'Alternative Approaches to the Theory of Choice in Risk-taking Situations', with the paragraph:

> There is no need to enlarge upon the importance of a realistic theory explaining how individuals choose among alternative courses of action when the consequences of their actions are incompletely known to them. It is no exaggeration to say that every choice made by human beings would meet this description if attention were paid to the ultimate implications. Risk and the human reactions to it have been called upon to explain everything from the purchase of chances in a 'numbers' game to the capitalist structure of our economy; according to Professor Frank Knight, even human consciousness itself would disappear in absence of uncertainty. (Arrow, 1951, p. 204)

During the first half of this century a number of alternative views were developed concerning the nature of uncertainty, the possibly different kinds of uncertainty and whether, or in what circumstances, it could be measured. I cannot review this development in any detail, but by the time Arrow wrote his article, a number of competing views had been formulated (see Arrow, 1951; Savage, 1954, ch. 4).

The first question was: what is one uncertain about? As the above quote suggests, economists tended to focus on uncertainty about the consequences of actions. One could also be uncertain about the laws of nature, or about the truth or falsity of mathematical propositions. Of course, to the extent that knowing the consequences of actions depended on knowing the laws of nature or the implications of mathematical axioms, uncertainty about the latter would be important for economic actions as well.

A second question was: can uncertainty be meaningfully measured (quantified)? I believe that there was general (but not universal) agreement that, if uncertainty could be quantified, then that quantification should obey the mathematical laws of probability. However, the 'frequentist' school reserved the legitimacy of probabilistic quantification to 'experiments' (planned or naturally occurring) that were repeated indefinitely under identical conditions. Thus, according to them, it would not be meaningful on 1 July 1998, to assign a probability to the event, 'President Clinton will not serve out his full term of office', or to the event that 'an economical method of producing energy by controlled fusion will be developed by 1 January 2001'. This view was espoused by Frank Knight (1921), who made the distinction between 'risk' and 'uncertainty', the former being measurable in an actuarial sense, and the latter not. Knight went on to apply this distinction to 'organizational economics' by asserting that, since 'uncertainty' (in his sense) could not be insured against, it was the peculiar role of economic entrepreneurs to bear this uncertainty in the hope of being rewarded with economic profits (i.e. profits beyond the returns due to the standard factors of production). He went on to say:

> When uncertainty is present, and the task of deciding what do takes the ascendancy over that of execution, the internal organization of the productive groups is no longer a matter of indifference or a mechanical detail. (Knight, 1921, p. 268)

The 'personalist school,' which included a diverse set of methodologies, argued that the concept of probability was applicable to events outside of the frequentist realm. Some personalists went so far as to deny that the frequentist view could be applied meaningfully to any events at all, i.e. all probability judgements were 'personal'. (See Arrow's and Savage's accounts of the work of Ramsey, Keynes, Jeffries, Carnap and de Finetti. In some sense, the personalist view might also be ascribed to earlier authors, such as Bayes and Laplace.) The personalist view was given a solid foundation by Savage (1954), whose stated aim was to rationalize statistical practice. Ironically, his theory seems to have had only a modest effect on statistical practice. On the other hand, it was adopted by mainstream economists and game theorists as the dominant model of how decision-makers should (or even do) respond rationally to uncertainty and information. (See Section 3 for a summary of Savage's theory and its implications for economic decision-making.)

Central to the development of thinking about uncertainty was the simple idea that uncertainty about the consequences of an action could (or should) be traced to uncertainty about the 'state of the world' in which the action would be taken. An essential feature of the concept of the 'state of the world' is that it is beyond the control of the decision-maker in question. More precisely, its occurrence or non-occurrence is independent of the decision-maker's action. A further clarification was provided by the theory of games, put forward by von Neumann and Morgenstern (1947). In multiperson decision-making situations in which the participants have conflicting goals, the theory of games distinguishes between two aspects of the state of the world from the point of view of any single decision-maker, namely (i) the 'state of Nature', which is beyond the control of any of the persons involved, and (ii) the actions of the other persons. The latter has been called 'strategic uncertainty.' Economists (going back at least to Cournot, 1838), had struggled for a long time with how to incorporate strategic uncertainty into models of rational behavior in duopolistic and oligopolistic situations. Following the appearance of von Neumann and Morgenstern's book (1947), a number of different game-theoretic 'solution concepts' were put forward, and they generally predicted different outcomes, except in the special case of two-person zero-sum (or constant-sum) games. Nash's concept of 'non-cooperative equilibrium' (Nash, 1951), which was a generalization of Cournot's concept of oligopolistic market equilibrium, finally emerged as the 'solution' most popular with economists. However, in Nash's theory there was only strategic uncertainty, but no uncertainty about 'states of Nature'. Furthermore, his concept of non-cooperative equilibrium—or what is now usually called *Nash equilibrium*—did not entirely solve the problem of strategic uncertainty. The difficulty is that in a large number of game-theoretic models there are many Nash equilibria, and so the theory provides no good basis for predicting the actions of the other players.

Savage's book, *The Foundations of Statistics*, appeared in 1954, three years after Arrow's 1951 article. The synthesis of his approach with Nash's theory of games was initiated by Harsanyi (1967–68), and further developed by Selten (1975) and Kreps and Wilson (1982). The resulting theory of 'Nash–Harsanyi equilibrium' did not eliminate the difficulty of a multiplicity of equilibria. Nevertheless, I think that it is fair to say that this synthesis now provides the dominant formulation of economic rationality, especially in the theoretical development of the economics of organizations.

In the remainder of this paper, I shall concentrate on a critique of this theory of rational behavior *as it is applied to single-person decision-making, or to multiperson situations in which the persons do not have conflicting goals*. A critique of the theory of strategic rationality is beyond the scope of the present paper, but for some remarks on this see Radner (1997).

2.2 The Savage Paradigm

A fuller sketch of the *Savage paradigm* of individual rational decision-making under uncertainty will be given in Section 3, but a few elements are needed here in order to introduce the notions of costly and bounded rationality. The essential building blocks of the theory are (i) a set of alternative *states of the world*, or simply *states*, which are beyond the decision-maker's control; (ii) a set of alternative actions available to the decision-maker, or as Savage calls them, *acts*; and (iii) a set of alternative *consequences*. An act determines which consequence will be realized in each state (of the world). Hence a parsimonious way to think about an act is that it is a function from states to consequences.

The simplicity of this formulation hides a wealth of possible interpretations and potential complexity. I shall discuss these in more detail in subsequent sections, but here I shall allude only to a few of them. First, the decision-maker (DM) is assumed to have preferences among acts. These preferences reflect both the DM's *beliefs* about the relative likelihood of the different states, and the DM's *tastes* with regard to consequences. A few axioms about the independence of beliefs from consequences enable the DM to infer preferences among complicated acts from those among simpler ones. The DM is assumed (or advised) to choose an act that is most preferred among the available ones. I shall call this the assumption of *optimization*.

Second, if the decision problem has any dynamic aspects, then states can be quite complex. In fact, a full description of any particular state will typically require a full description of the *entire history* of those features of the DM's environment that are relevant to the decision problem at hand.

Third, the description of the set of available acts reveals—if only implicitly—the opportunities for the DM to acquire information about the state of the world and react to it. Together with the axioms about preferences among acts, this imposes a fairly stringent constraint on how the DM should learn from observation and experience. In fact, we shall see that this aspect of the model is what gives the Savage paradigm its real 'bite'. An act that describes how the DM acquires information and reacts to it dynamically is sometimes called a *strategy* (plan, policy). Savage described his point of view as follows:

> As has just been suggested, what in the ordinary way of thinking might be regarded as a chain of decisions, one leading to another in time, is in the formal description proposed here regarded as a single decision. To put it a little differently, it is proposed that the choice of a policy or plan be regarded as a single decision. This point of view, though not always in so explicit a form, has played a prominent role in the statistical advances of the present century. For example, the great majority of experimentalists, even today, suppose that the function of statistics and of statisticians is to decide what conclusions to draw from data gathered in an experiment or other observational program. But statisticians hold it to be lacking in foresight to gather data without a view to the method of analysis to be employed, that is, they hold that the design and analysis of an experiment should be decided upon as an articulated whole. (Savage, 1954, pp. 15, 16)

In a model of a sequential decision problem, the space of available strategies can, of course, be enormous and complex. This observation will be a dominant motif in what follows.

2.3 The Simon Critique

As Herbert Simon emphasized in his work, the cognitive activities required by the Savage paradigm (and its related precursors) are far beyond the capabilities of human decision-makers, or even modern human/computer systems, except in the simplest decision problems. This led him and his colleagues (especially at Carnegie–Mellon University) to investigate models of human decision-making that are more realistic from the point of view of cognitive demands, and yet do not entirely abandon the notion of rationality. As he put it,

> Theories that incorporate constraints on the information-processing capabilities of the actor may be called *theories of bounded rationality.* (Simon, 1972, p. 162)[2]

Simon was not the first to criticize the mainstream economic model of rational human behavior. The well-known economist J. M. Clark took time out during the First World War to study contemporary psychology, which effort resulted in 1918 in an article, 'Economics and Modern Psychology'. In a felicitous passage he wrote:

> If one wanted to be unfair to economists in general, he might select, for purposes of comparison with these psychological principles, a certain well known though fictitious character whose idiosyncrasies furnish alternate joy and irritation to modern readers of economics. He is a somewhat inhuman individual who, inconsistently enough, carries the critical weighing of hedonistic values to the point of mania. So completely is he absorbed in his irrationally rational passion for impassionate calculation that he often remains a day laborer at pitifully low wages from sheer devotion to the fine art of making the most out of his scanty income and getting the highest returns from his employers for such mediocre skill as he chooses to devote to their service. Yet he cannot fail to be aware that the actuarial talent he lavishes outside of working hours would suffice to earn him a relatively princely salary in the office of any insurance company. So intricate are

[2] Simon's interest in bounded rationality appears already in his early book, A*dministrative Behavior* (1947). See also Simon (1955), and March and Simon (1958, esp. ch. 6).

> the calculations he delights in that even trained economists occasionally blunder into errors in recording them. (Clark, 1918, p. 24)

Savage himself was aware of the problem of bounded rationality. Early in his book he wrote, regarding the choice of a strategy (plan, policy):

> The point of view under discussion may be symbolized by the proverb, 'Look before you leap', and the one to which it is opposed by the proverb, 'You can cross that bridge when you come to it.' When two proverbs conflict in this way, it is proverbially true that there is some truth in both of them. . . . One must indeed look before he leaps, in so far as the looking is not unreasonably time-consuming and otherwise expensive; but there are innumerable bridges one cannot afford to cross, unless he happens to come to them.
>
> Carried to its logical extreme, the 'Look before you leap' principle demands that one envisage every conceivable policy for the government of his whole life (at least from now on) in its most minute details, in the light of the vast number of unknown states of the world, and decide here and now on one policy. This is utterly ridiculous, not—as some might think—because there might later be cause for regret, if things did not turn out as had been anticipated, but because the task implied in making such a decision is not even remotely resembled by human possibility. It is even utterly beyond our power to plan a picnic or to play a game of chess in accordance with the principle, even when the world of states and the set of available acts to be envisaged are artificially reduced to the narrowest reasonable limits. (Savage, 1954, p. 16)

Nevertheless, Savage felt that his model was a useful one for thinking about rational decision-making.

> Though the 'Look before you leap' principle is preposterous if carried to extremes, I would argue that it is the proper subject of our further discussion, because to cross one's bridges when one comes to them means to attack relatively simple problems of decision by artificially confining attention to so small a world that the 'Look before you leap' principle can be applied there. I am unable to formulate criteria for selecting these small worlds and indeed believe that their selection may be a matter of judgement and experience about which it is impossible to enunciate complete and sharply defined general principles. . . . On the other hand, it is an operation in which we all necessarily have much experience, and one in

which there is in practice considerable agreement. (Savage, 1954, pp. 16, 17)

Savage's approach remains the dominant method of 'mainstream' economists' theorizing about economic organization. In this methodology, the theorist sets up a rather simple model of the situation, and then applies the Savage paradigm, or some variation thereof, in its full force. Of course, the model has to be simple enough for the analyst to able to be capable of deducing interesting implications from the assumptions of 'full rationality'.

A different kind of approach, sometimes called 'behavioral', abandons the assumption of optimization in favor of some other model of decision-making, supposedly empirically based. Examples include 'satisficing', 'rules of thumb', 'routines', etc.[3] The development of these approaches has also had some interesting and close connections with the field of artificial intelligence, and Simon himself was an early contributor to both fields (e.g. Simon, 1981; Newell and Simon, 1972). Unfortunately, a serious treatment of this subject is beyond the scope of this paper.

2.4 Costly and 'Truly Bounded' Rationality

Simon's description of theories of bounded rationality as 'incorporat[ing] constraints on the information-processing capabilities of the actor' is broad enough to cover a number of different approaches. One approach deals with *costly rationality* without giving up the assumption of optimization. In this extension of the mainstream economic methodology, one explicitly models the constraints and/or costs relating to the cognitive capabilities of the DM, and assumes that the DM optimizes fully, subject to those constraints and taking account of the relevant costs. I shall call this the *extended Savage paradigm* (see Section 4). Early work in this spirit included the statistical research of Abraham Wald (1947, 1950), and the 'theory of teams', introduced by Jacob Marschak (Marschak, 1955; Radner, 1962; Marschak and Radner, 1972).

More radical approaches, which give up the assumption of optimization, are required to deal with the problems caused by *vagueness*, *ambiguity*, and *'unawareness'* in the decision-making situation. Even more troublesome is the *failure of logical omniscience*, which refers to the fact that the DM cannot know all of the logical consequences of what he knows. To avoid confusion with Simon's original terminology, I shall lump these diverse phenomena together under the heading *truly bounded rationality*, and discuss them in Section 5.

[3] For early work in this spirit, see Simon (1947, 1955, 1957) and March and Simon (1958). For a further formal analysis, see Radner (1975a).

3. *A Trunk: The Savage Paradigm*

The *Savage paradigm* of individual decision-making in the presence of uncertainty has already been briefly introduced in Section 2. In this section I shall give a fuller sketch of the theory, but this sketch will be far from complete or systematic. My limited goal here is to provide a background for the subsequent discussion of bounded rationality.[4]

As noted in Section 2.2, there are three essential building blocks of the theory:

1. a set S of alternative *states of the world*, or simply *states*, which are beyond the decision-maker's control;
2. a set A of alternative actions available to the decision-maker, or as Savage calls them, *acts*; and
3. a set C of alternative *consequences*.

An act determines which consequence will be realized in each state (of the world). Hence a parsimonious way to think about an act is that it is a function from the set of states to the set of consequences.

Savage illustrates the model with the following example:

> Your wife has just broken five good eggs into a bowl when you come in and volunteer to finish making the omelet. A sixth egg, which for some reason must either be used for the omelet or wasted altogether, lies unbroken beside the bowl. You must decide what to do with this unbroken egg. Perhaps it is not too great an oversimplification to say that you must decide among three acts only, namely, to break it into the bowl containing the other five, to break it into a saucer for inspection, or to throw it away without inspection. Depending on the state of the egg, each of these three acts will have some consequence of concern to you, say that indicated by Table 1. (Savage, 1954, p. 14)

TABLE 1. An Example Illustrating Acts, States and Consequences

Act	State of the sixth egg	
	Good	Rotten
Break into bowl	six-egg omelet	no omelet, five good eggs destroyed
Break into saucer	six-egg omelet, saucer to wash	five-egg omelet, saucer to wash
Throw away	five-egg omelet, one good egg destroyed	five-egg omelet

[4] The best systematic treatment is still Savage's book (1954, 1972), which includes extensive historical and critical remarks. A briefer, and somewhat more accessible, account is given in Part I of Marschak and Radner (1972).

Note that:

1. The set *S* has two alternative states (of the 'world'), namely 'the sixth egg is good' and 'the sixth egg is rotten'.
2. The set *C* has six different consequences, described in the six entries in the body of the table.
3. (i)The set *A* has 3 alternative acts:
   (i) break the sixth egg into the bowl;
   (ii) break the sixth egg into the saucer, then add it to the bowl if it is good, and throw it away if it is bad;
   (iii) throw the sixth egg away without breaking it.

The example can be used to illustrate two points. First, the DM's uncertainty about the consequences of action has been modeled as uncertainty about the state of the sixth egg, which is beyond the DM's control.

Second, acts (i) and (iii) can be thought of as corresponding to 'simple actions', taken without further information about the state of the sixth egg. On the other hand, act (ii) corresponds to a 'strategy' (plan, policy), according to which some information is first obtained (by breaking the sixth egg into the saucer), and then some appropriate further action is taken, depending on the information obtained (about the sixth egg). Note that the information is 'costly', since the saucer must be washed.

Typically, the set of states of the world for a decision problem will have many states, not just two. It corresponds to our everyday language to call a set of states an *event*. Thus, in the omelet example, there might be three states of the sixth egg: very fresh, moderately fresh and rotten. The first two states together would then comprise the event, 'the sixth egg is usable'.

Returning to the general model, Savage makes a number of assumptions about the DM's behavior in decision problems of this type. In this paper I can only highlight a few of those assumptions, and sketch their most important implications. First, it is assumed that the DM can rank all of the acts in order of preference. Implicit in this assumption is the interpretation that the DM will choose an act that is ranked highest in this preference order. (Some further mathematical assumption is needed to ensure that there will be at least one best act, e.g. this would be the case if the set of acts were finite.) At first sight, one might think that this is much too demanding a requirement. After all, in a complicated decision problem the DM will want some help, i.e. some method for finding the best act. However, additional assumptions about the structure of the DM's preferences make it possible (in principle) for him to infer his preferences among complicated acts (e.g. strategies) from his preferences among relatively simple ones. Taken together, these assumptions express the idea that the DM's tastes concerning consequences are independent of the states in which they occur, and his beliefs about the relative likelihood of different states (or events) are independent of the accompanying consequences. I shall call this requirement *the independence of tastes and beliefs*.

To some extent the postulate of independence of tastes and beliefs is a convention, in that its validity will depend on a careful description of states, acts and consequences. Savage gives the following example to illustrate this point. A person is about to go on a picnic with friends, but does not yet know where the picnic will be held. Furthermore, he will have no influence on that decision. Provisionally, we may regard the place of the picnic as the relevant state of the world. Before going on the picnic, the person

> . . . decides to buy a bathing suit or a tennis racket, not having at the moment enough money for both. If we call possession of the tennis racket and possession of the bathing suit consequences, then we must say the consequences of his decision will be independent of where the picnic is actually held. If the person prefers the [possession of the] bathing suit, this decision would presumably be reversed if he learned that the picnic were not going to held near water. . . . But under the interpretation of 'act' and 'consequence' that I am trying to formulate, this is not the correct analysis of the situation. The possession of the tennis racket and the possession of the bathing suit are to be regarded as acts, not consequences. (It would be equivalent and more in accordance with ordinary discourse to say the coming into possession, or the buying, of them are acts.) The consequences relevant to the decision are such as these: a refreshing swim with friends, sitting on a shadeless beach twiddling a brand-new tennis racket while one's friends swim, etc. (Savage, 1954, p. 25)

Savage's assumptions about the structure of the DM's preferences among acts imply the following proposition, sometimes called the *expected utility hypothesis* (EUH). Implicit in the DM's preferences are two scales. The first scale assigns *probabilities to events*, and obeys the usual mathematical laws of the calculus of probabilities. The second scale assigns *utilities to consequences*. (It is uniquely determined except for the choice of the origin and the unit of utility.) Using these two scales, one can calculate an *expected utility* for each act, in the usual way, since an act associates a consequence with each state. Thus for each state one calculates the product of its probability times the utility of the associated consequence, and then adds all of the products to obtain the expected utility of the act. Finally, the DM's preferences among acts is *represented by expected utility* in the following sense: *one act is as good as or better than another (from the point of view of the DM) if and only if the expected utility of the first act is greater than or equal to the expected utility of the second act.*

In one sense, the EUH is quite weak. No restrictions are imposed on the probability measure and the utility function implicit in the DM's preferences. To paraphrase an old saying, 'neither beliefs nor tastes are subject to dispute'. Furthermore, it can be shown that the specific form of the utility function represents not only the DM's preference ranking of consequences, but also his attitude towards risk. Thus the Savage axioms put no restrictions on the DM's attitude towards risk—he can be averse to risk, a risk-lover or some combination of the two. Thus two DMs can 'rationally' differ in their tastes, their beliefs and their attitudes towards risk.

On the other hand, the *EUH completely determines how a DM must update his beliefs, and hence his expected utility of acts, in response to new information.* In fact, his updating of beliefs must be governed by the mathematical laws of conditional probability, which in this case takes the form of the celebrated 'Bayes's Rule'. Hence the updating of beliefs by a DM who conforms to the Savage axioms is sometimes called *Bayesian learning*.

A simple example will illustrate the strength of the EUH in this context. Suppose that a 'black box' produces a potentially infinite sequence of zeroes and ones; denote the sequence by $\{X(t)\}$. Two DMs disagree about the probability law that governs this stochastic process, $\{X(t)\}$, but they agree on the following: any two sequences of the same finite length that differ only by a permutation of the zeroes and ones in them have the same probability. Suppose that the black box produces an infinite sequence $\{X(t)\}$ such that the relative frequency of ones is some number, say p, and suppose further that the DMs observe successively longer finite initial segments of the sequence, say

$$[X(1), X(2), \ldots, X(T)]$$

Then as T increases without bound, the DMs will asymptotically approach agreement that the black box produces a sequence of independent and identically distributed random variables, taking the values zero or one, and that the probability of a one is equal to p.[5]

The above example illustrates the *merging of opinions* of DMs who start with different beliefs about the state of the world, but who receive larger and

[5] *Technical note*: The condition on which the DMs agree is called *symmetry* or *exchangeability.* In my loose statement of the proposition I have left out an important technical assumption, which is—roughly speaking—that the *a priori* beliefs of neither DM rule out as impossible any particular value of p (Savage, 1954, pp. 50–55).

larger 'chunks' of identical information. General conditions under which the merging of opinions occurs has been studied by a number of authors.[6] Space limitations do not allow me to review this material here, but I shall return briefly to the topic in Section 5.

4. *Branch I: Costly Rationality*

4.1 The Extended Savage Paradigm

As just sketched in the previous section, the Savage paradigm does not appear explicitly to take account of the costs of decision-making. However, nothing prevents the DM from incorporating into the description of the consequences of an act the costs—in terms of resources used—of implementing the corresponding actions. The costly activities involved in decision-making include:

1. Observation
2. Information processing, i.e. computation
3. Memory
4. Communication

The last category may be important when the decision-making process is assigned to a team of individuals.

If the resources used by these decision-making activities are limited, then those limits may impose binding constraints on the activities themselves, constraints that must be taken into account in the DM's optimization problem. If the constraints are on the *rate of resource use per unit time*, then more extensive decision-making activities may cause *delays* in the implementation of the eventual decisions. To the extent that a delay lowers the effectiveness of a decision (e.g. by making it more obsolete), we may think of delay as an 'indirect cost'.

We shall see that extending the Savage paradigm to incorporate the costs of decision-making may in some cases be natural, and in other cases problematic. The first class of cases I shall call *costly rationality*, which is the subject of the present section. My treatment will be more of an outline than a true discussion. Fortunately, there is a substantial literature on many of the subtopics, although the authors usually have not used the term 'costly rationality'.

[6] See Diaconis and Freedman (1986), for a review of the statistical literature up to that date; see also Sims (1971) and Jackson *et al*. (1999) for additional material.

4.2 Observation

The notion that observation is costly was implicit in the Neyman–Pearson theory of hypothesis testing, and was made explicit by Abraham Wald in his pioneering studies of sequential statistical procedures. (See Wald, 1950, for an influential codification of his general approach.) Examples of such sequential procedures include (i) the destructive testing of fuses and artillery shells, which falls under the more general rubric of *sequential analysis* (Wald, 1947); and (ii) clinical trials of new drugs, which is part of a more general topic with the curious name *bandit theory* (Basu *et al.*, 1990). The cost of observation also figures in more classical (non-sequential) statistical problems such as the design of sample surveys and agricultural experiment. Given some model of the costs of observation, the DM chooses the kind and amount of observation, optimally balancing the expected benefits of additional observations against their costs. Such decision problems fit naturally into the Savage paradigm, although taking account of these costs typically complicates the analysis. For example, in the case of clinical trials and other 'bandit' problems, the calculation of optimal policies quickly becomes computationally intractable for many problems of realistic size.

4.3 Information Processing and its Decentralization

Even after the information has been collected (e.g. by observation) it still must be further processed to produce the required decisions. This information-processing task may be quite demanding. Examples include (i) computing a weekly payroll, (ii) scheduling many jobs on many machines, (iii) managing multi-product inventories at many locations, and (iv) project selection and capital budgeting in a large firm. Such tasks are typically too complex to be handled by a single person, even with the aid of modern computers. In such circumstances the required processing of the information is *decentralized* among many persons in the organization.

The formal modeling of decentralized information processing is a relatively recent research topic. Computer science has provided a number of useful models of information processing by both computers and humans, and the decentralization of information processing in human organizations finds its counterpart in the theories of parallel and distributed processing in computer systems. Marschak and McGuire (1971) were probably the first to suggest the use of a particular model of a computer (the finite automaton) to represent the limited information processing capabilities of humans in economic organizations. Reiter and Mount were early contributors to this line of research, and went further in analyzing economic organizations as networks of computers (see Mount and Reiter, 1998, and references cited there).[7]

One conclusion from this literature is what I have called *the iron law of delay* for networks of processors of bounded individual capacity. This 'law' can be paraphrased in the following way: a*s the size of the information processing task increases, the minimum delay must also increase unboundedly, even for efficient networks and even if the number of available processors is unlimited.*[8]

4.4 Memory and Communication

Memory storage and communication among humans and computers are also resource-using activities, and cause further delays in decision-making. I shall not discuss these topics here, except to say that both the storage and transmission of information and the results of information processing seem to be relatively 'cheap' compared to observation and processing, at least if we consider computer-supported activities. The proliferation of large data banks, and the flood of junk mail, telephone calls and email, lend support to this impression. It appears that today it is much cheaper, in some sense, to send, receive and store memos and papers than it is to process them. Nonetheless, organization theorists should not, and have not, neglected the study of these activities. Game theorists have paid particular attention to models of players with limited memory (see Osborne and Rubinstein, 1994; Rubinstein, 1998, and references cited therein). For models of costly communication in organizations, and some implications for organizational structure, see, for example, Marschak and Reichelstein (1998) and Bolton and Dewatripont (1994).

4.5 Implications for Organizational Structure and Returns to Scale

The first implication of costly rationality for organizational structure is that in all but the very smallest economic organizations the process of decision making will be decentralized to some extent, i.e. distributed among several persons. Indeed, it is difficult to think of an economic organization with more than one member in which this is not the case. In itself, this says nothing about the 'decentralization of authority' in the every-day sense. In fact,

[7] For more recent developments, see Radner (1993), Radner and Van Zandt (1992) and Van Zandt (1998a–c, 1999). See also papers presented at the recent conference organized by the CEPR on 'Information Processing Organizations', Free University of Brussels, 25–26 June 1999.

[8] Of course, a precise statement would require a precise description of the models in which it has been verified. See Radner (1997) and references cited therein for a fuller discussion.

the concept of 'authority' has proved elusive in the economic theory of organization, and I shall have nothing to say about it in this paper. (For some remarks on this topic, see Radner, 1997, pp. 336–338.)

The fact that decision-making is costly implies that a DM or organization will economize on the costly activities that are involved. In particular, since the incremental value of additional information typically declines after some point, a DM will not want to use all possible information for any one decision or group of related decisions. This in turn implies that, in an efficient organization, *different decisions or groups of related decisions will be based on different sets of information.* Hence, in a large organization both the information-processing and the decision-making will be distributed among many persons. In other words, in a large organization, *the decentralization of both information-processing and decision-making is inevitable.*

The implications for the architecture of this decentralization are less clear. The theory of teams was an early attempt to grapple with this question (Marschak, 1955; Marschak and Radner, 1972). One handicap of this attempt was that at the time there was no readily available model of the 'technology' of information-processing. Thus team theory compared the *benefits* of different organizational structures, but not their *costs*. Perhaps the time is ripe to return to some of those issues with the help of more recent models of information-processing.

Another complexity in the analysis of organizational structure is that the decentralization of information leads to a 'decentralization of power' (Radner, 1975b). I shall not even try to give a precise definition of this concept, but more recent studies of principal-agent relationships and other models of distributed (or 'asymmetric') information give some content to the idea. In other words, decentralization of information, together with diversity of goals among DMs in the organization, generates *strategic uncertainty*, with all its game-theoretic ramifications. As I have said above, I have put such considerations outside of the purview of the present paper (see, however, Dutta and Radner, 1995; Radner, 1992, 1997). Recall that in the theory of teams it is postulated that all the team members have the same goal, although they may have different information and make different decisions.

The renewed interest in the theoretical study of organizational structure, using both the tools of game theory and models of information-processing technology, has made progress, especially in the study of 'hierarchical' organizations. However, in my view we are still far from understanding why and under what circumstances firms are—or are not—organized hierarchically, and why economic activities are sometimes organized in firms and sometimes in markets (Radner, 1992).

A perennial topic in economic theory concerns the sources of increasing and decreasing returns to scale in firms, and in particular what are the organizational sources of scale effects. Economists have generally concluded that decreasing returns to scale should arise only when some input is held fixed.

Otherwise, at least constant returns to scale should be achievable by simply replicating a firm of a given size, and calling the collection of replicated firms a single 'firm.' On the other hand, one might hypothesize that

> there are organizational limits on returns to scale due to the problems of coordinating large numbers of activities within a single firm and operating in diverse environments. Such limits are theoretically possible because replication cannot be used to extend the scale of a firm at constant or decreasing cost. Replicating the activities of several small firms means that the subunits cannot communicate, coordinate their activities, or allocate resources, except as independent firms would do, such as through markets. There can be no headquarters that controls the subunits because such control would incur managerial costs and delays that the independent firms avoid. Such informationally disintegrated units could hardly {be regarded as} a single firm. (Van Zandt and Radner, 2000, p. 8)

As early as 1934, Nicholas Kaldor argued that the supply of coordinating ability must be limited in any single firm.

> You cannot increase the supply of coordinating ability available to an enterprise alongside an increase in the supply of other factors, as it is the essence of coordination that every single decision should be make on a comparison with all the other decisions made or likely to be made; it must therefore pass through a single brain. (Kaldor, 1934, p. 68)

However, the value of information and information processing, and the associated technologies, do not always follow the laws of the consumption and production of material goods. They also interact in surprising ways. To begin with, depending on how it is measured, the marginal value of information can be increasing at low levels of information, and decreasing at higher levels. Indeed, it can be shown that, under very general conditions*, the marginal value of information is increasing at sufficiently low levels* (Radner and Stiglitz, 1983). On the other hand, the law of large numbers (and similar theorems) imply that in standard sampling situations the incremental value of increasing the sample size is eventually decreasing. Applying team theory, Radner (1961) studied how the value of specific organizational devices varied with scale, but without regard to their costs.

On the technology side, Kenneth Arrow pointed out some time ago that the production and dissemination of information would typically show *increasing* returns to scale (Arrow, 1974). This would be so because the production of information usually requires a fixed expenditure of resources independent of the scale of the use or dissemination of the information (as in the case of research and development). The situation is even more striking today in the case of electronic communication, where the marginal cost of dissemination of information (including software) is often almost negligible compared to the initial development cost. The 'iron law of delay' (see Section 4.3 above) further complicates matters because the loss due to delay depends crucially on the intertemporal statistical properties of the relevant environment.

Even without multiplying such examples, it should not be surprising that returns to scale in decision-making activities show complex patterns. Van Zandt and Radner identify three informational scale effects:

> diversification of heterogeneous risks (positive), sharing of information and of costs (positive), and crowding out of recent information due to information processing delay (negative). [On the other hand, because] decision rules are endogenous, delay does not inexorably lead to decreasing returns to scale. However, returns are more likely to be decreasing when computation constraints, rather than sampling costs, limit the information upon which decisions are conditioned. [These] results illustrate the fact that informational integration causes a breakdown of the replication arguments that are used to [predict] nondecreasing technological returns to scale. (Van Zandt and Radner, 2000)

Finally, distributed information together with conflicting goals may lead to 'loss of control' or 'agency costs', which have their own peculiar scale effects. However, this takes us into game-theoretic models, which (I repeat) are beyond the scope of this paper.

5. *Branch II: Truly Bounded Rationality*

5.1 Introduction

Recall that in the Savage paradigm the DM has a model of the decision problem that is comprised of four building blocks:

1. a set S of alternative *states of the world*, or simply *states*, which are beyond the DM's control; subsets of S are the *events* about which the DM may be uncertain;
2. a set C of alternative *consequences*;
3. a set A of alternative *acts* available to the decision-maker; each act is a function from the set of states to the set of consequences;
4. a complete preference ordering on the set of acts.

In addition, several axioms are postulated concerning the preference ordering on the set of acts. Taken together, these axioms define a notion of *consistency* that is required of the DM by the paradigm. A preference ordering that is consistent in this sense can be represented by two numerical scales, a *personal probability scale* (defined on events) and a *utility scale* (defined on consequences). Of any two acts, one is preferred to the other if and only if it has a higher *expected utility*.

Among the difficulties a DM typically faces in trying to apply the Savage paradigm to a real decision problem are:

1. Inconsistency
2. Ambiguity
3. Vagueness
4. Unawareness
5. Failure of logical omniscience

As will be seen, these difficulties are somewhat related and overlapping. I mention 'ambiguity' because it is a term one often finds in the literature on decision-making. On the other hand, it appears difficult to distinguish in practice between 'ambiguity' and 'vagueness.' Hence, following Savage, I shall stick with the latter term.

In what follows, I give only a sketch of these difficulties and how they might be met. Savage (1954) devotes Chapter 4 and Section 5.6 to a more detailed discussion of inconsistency and vagueness. An account of more recent attempts to deal with some of these problems is given by Machina (1987, 1889).

5.2 Inconsistency

As is so often the case in this paper, I introduce the present topic with a quote from Savage.

> According to the personalistic view, the role of the mathematical theory of probability is to enable the person using it to detect inconsistencies in his own real or envisaged behavior. It is also understood that, having detected an inconsistency, he will remove it. An inconsistency is typically removable in many different ways, and the theory gives no guidance for choosing. (Savage, 1954, p. 57)

Some 'inconsistencies' have been observed so frequently, and have been so 'appealing', that they have been used to criticize the Savage axioms, and to form a basis for a somewhat different set of axioms. In particular, I have in mind the so-called 'Allais Paradox' and 'Ellsberg Paradox,' which I shall not discuss here (see Machina, 1987).

In other cases, it has been argued that inconsistent preferences arise because the DM is forced to articulate preferences about which he is not 'sure'. (This explanation is related to 'vagueness'; see below.) To solve this problem, some authors have proposed the introduction of 'probabilities of a second order', according to which the DM thinks that some probability comparisons of events are more likely to be correct (for him) than others. Savage rejects this approach as 'leading to insurmountable difficulties . . . once second order probabilities are introduced, the introduction of an endless hierarchy seems inescapable. Such a hierarchy seems very difficult to interpret, and seems at best to make the theory less realistic, not more' (Savage, 1954, p. 58). The axioms of consistency can also be used to alleviate problems of 'unsureness' and the resulting inconsistencies. As de Finetti points out:

> The fact that a direct estimate of a probability is not always possible is just the reason that the logical rules of probability are useful. The practical object of these rules is simply to reduce an evaluation, scarcely accessible directly, to others by means of which the determination is rendered easier and more precise. (de Finetti, 1937, p. 60)

Another device to deal with 'unsureness' about preferences between acts is to relate this unsureness to uncertainty about some relevant state of the world. Going further, unsureness about preferences may also be due to uncertainty about what *are* the states of the world , i.e. about *what is the set of possible states*, a circumstance that can lead to a preference for 'flexibility' (see below).

It has been observed in experiments that inconsistencies in preferences are more frequent the closer the alternatives are in terms of preference. This observation led Marschak and others to the elaboration of models of 'stochastic choice'. Such models have become an important tool of the econometrics of demand (McFadden, 1999).

Finally, the DM's preferences can also be inconsistent (relative to the axioms) *because he makes mistakes in mathematics, or more generally, in logic*. This difficulty will be discussed below.

5.3 Vagueness and Unawareness

I have already alluded to the DM's possible vagueness about his preferences. However, he could also be vague about any aspect of his model of the decision problem, and is likely to be so if the problem is at all complex. Vagueness can be about the *interpretation* of a feature of the model, or about its *scope*, or both, as I shall explain in a moment.

A common source of vagueness is the need for the DM to limit artificially the scope of his model of the decision problem. One cannot expect the DM to plan his whole life as one grand decision problem. (See Savage's remarks about the 'Look before you leap' principle, quoted above in Section 2.3.) Instead, life is decomposed into a set of roughly independent or loosely linked decision problems, together with large parts that are more or less unplanned. As we shall see this creates problems of interpretation. Savage calls this the problem of 'small worlds', and provides some formal conditions under which a small-world decision problem can, without loss, be carved out of the grand decision problem (Savage, 1954, pp. 82–91). As one might expect, these conditions are quite restrictive, and will not provide much comfort in most realistic situations.

Typically, as the DM uses a particular model he will refine the interpretation and expand the its scope. He may also become aware of states, consequences and acts of which he was not previously aware. Furthermore, *although he knows that these revisions will take place in the future, he cannot predict what they will be.* This dynamic process of model revision is a key feature of most applications of the Savage paradigm to real problems, but the theory implicit in the paradigm does not deal with this feature, at least not directly. *I see the development of a theory of rational model revision as a major challenge posed by the problem of bounded rationality* (see Section 5.4).

Consider first the problems of vagueness with respect to consequences. Recall that, in the Savage paradigm, all uncertainty about the consequences of action is reduced to uncertainty about the 'state of the world' in which the action takes place. Thus, given the true state of the world, there is no further uncertainty about the consequence of an action. However, in applications of the theory to real decision problems, the DM will typically be forced to simplify the model to such a degree that consequences are in fact uncertain, even given the state of the world (in the model). For example, in a simple model of a decision whether to buy a lottery ticket, the 'consequence' of owning the winning number might be thought to be receiving a prize of, say, $1 million (less the cost of the ticket). To put it slightly more accurately, the 'consequence' of winning is an increase in the DM's wealth of $1 million less the cost of the ticket. However, in order to correctly assess the utility of this consequence, the DM would have to know many things, such as his future income and health, the rate of inflation, etc., which will in turn depend on future actions and the further evolution of the state of the world (a history). In other words, the 'winning the lottery' is not a 'pure consequence' in the strict sense required by the theory.

In principle, given enough time and patience, the DM might be able to describe many of the various 'pure consequences' of winning the lottery, which would involve describing in detail all the possible future histories of all of the aspects if his life that concern him. Another difficulty is that there will be possible futures that he is not aware of and cannot even imagine. Although I cannot prove it, I feel that in practice this is less of a problem with respect to consequences. For example, as far as I know, the development of the atomic bomb was not even imagined in the 15th century. However, the consequences of dropping the bomb on Japan in the Second World War, if cast in terms of persons killed, property destroyed, etc., was probably imaginable, even though people living in the 15th century could not imagine the particular actions and states of the world that produced those consequences in 1945. Other, more cheerful, examples could easily be drawn from the history of technology, medicine, etc.

I turn now to the problems of vagueness about states of the world. The first problem is analogous to the first one concerning consequences. It is the practical difficulty of describing in full detail all of the relevant aspects of the past and future history of the DM's environment that, together with his actions, determine the consequences that directly concern him. If a decision model is not soon discarded as useless, it is likely to be refined by expanding the scope of the set of states. This will involve both describing states in more detail, and adding states. Some of the added states will have, or could have, been envisaged originally, but were omitted in the process of model simplification. Others will not even have been envisaged or imagined. In other words, as time goes on the DM will become aware of states of which he was previously unaware. Nevertheless, the thoughtful DM will anticipate that this will happen, and will try to take account of this in his decision-making process.

An important case of this arises in the DM's construction of the model itself. The concept of 'state of the world' is quite flexible, and can encompass what in everyday language are called 'laws of nature' and 'model parameters'. For example, if the DM believes that a certain stochastic process is a Markov chain, but does not know its parameter values, then the set of states can be expanded to include a description of the parameter values as well as the actual realization of the process. (This will be discussed more fully in the next section.) As the DM revises and expands the model, new parameters will be envisaged, and the set of states must be correspondingly revised. In the next section I shall illustrate how this might be done 'rationally'.

Finally, consider the problems of vagueness and unawareness about acts. Any vagueness and/or unawareness about consequences and states will, of course entail corresponding problems for acts. In addition, the DM may in the course of time discover or invent new acts. The history of technical change provides a rich store of such examples.

If the DM anticipates that he will become aware of new acts in the future, then, other things equal, he may prefer present actions that allow for 'flexibility.' The concept of flexibility has had many interpretations in the economics literature, and the subject has a long history, in part associated with the theory of money.[9]

I have the impression that the revision of a decision model to incorporate new acts is most often associated with the incorporation of new states into the model, although a substantiation of this hypothesis would require further study. I have already alluded to a corresponding impression with regard to consequences. This motivates the analysis in the next section.

5.4 A Bayesian Analysis of Model Revision

In the previous section I suggested that some process of model revision was an important component of most real dynamic decision problems. This leads to the question of what it might mean to engage in this process 'rationally'. In what follows I propose a Bayesian approach to this question, in the context of a special case of statistical model revision (Radner, 2000). I should emphasize that a general approach to the question that goes beyond this example has not yet been worked out, so what I shall present is more of a research program than a full-blown theory.

Suppose that the DM must make successive predictions, one period ahead, of each of a sequence of random variables, $X(0)$, $X(1)$, . . ., *ad infinitum.* It will

[9] For recent contributions, and further references, see Kreps (1992) and Jones and Ostroy (1984).

be convenient to have a name for this sequence, so I shall call it the *X-process*. The prediction of each random variable $X(t)$ may be—but is not required to be—based on the past *history* of the X-process, defined by:

$$H(t-1) = \{X(0), \ldots, X(t-1)\}$$

Let us say, provisionally, that the state of the world, s, for this problem is a particular realization, x, of the X-process, i.e.

$$s = x = \{x(0), x(1), x(2), \ldots, \textit{ad infinitum}\}.$$

A *prediction strategy* for the DM is a complicated object. It determines

1. a prediction of the initial value, $X(0)$; and
2. for each subsequent period t and each history $H(t-1)$, a prediction of $X(t)$ as a function of $H(t-1)$.

Suppose that the DM is rewarded each period according to how accurate his prediction turns out to be. Then the relevant *consequence* is *the* (*infinite*) *sequence of eventual rewards.* Any particular policy determines a mapping from states (sequences of the random variables) to consequences (sequences of rewards); this mapping is an *act*.

In what follows it is not necessary to specify the DM's particular utility function on the set of consequences. Instead, we shall concentrate attention on DM's beliefs concerning the states, and how he updates his beliefs in the light of successively longer histories. To simplify the example further, suppose that the random variables $X(t)$ can take on only the values one or zero. According to the Savage paradigm, the DM's beliefs about the state can be scaled in terms of a probability law of the stochastic process of zeroes and ones.

For example, the DM might believe that the random variables $X(t)$ are independent and identically distributed (IID), with $\text{Prob}\{X(t) = 1\} = 1/6$, and $\text{Prob}\{X(t) = 0\} = 5/6$. [For example, as in tossing a 'fair die', with $X(t) = 1$ if the die comes up with one spot, and $X(t) = 0$ otherwise.] In this case, with a reasonable reward function, the DM will in each period ignore the previous history, and always predict a zero.

But what if the DM believes that the random variables are IID, but is uncertain about the probability that $X(t) = 1$? (For example, he thinks that the die might be biased.) In this case, the Bayesian approach would be to add a new state variable, call it p, and define the state as

$$s = (p,x)$$

[Recall that x denotes the infinite sequence, $\{x(t)\}$.] The DM's beliefs about the state would now have two components:

1. the marginal probability distribution of the 'parameter' p, called the *prior*; and
2. a family of conditional distributions of the X-process, given p, according to which, given p, the random variables $X(t)$ are IID, with Prob$\{X(t) = 1\} = p$.

Given any history, $H(t)$, the DM can calculate the conditional distribution of the parameter p given the history. This conditional distribution is called the *posterior distribution of p* (*given the history*). (The formula for this calculation is called 'Bayes's Rule.') The DM can then use this posterior distribution to calculate his prediction of the next random variable, $X(t + 1)$.

It is a striking theorem of Bayesian analysis that, if the DM's prior distribution of the parameter p is sufficiently 'open minded', then, if the true value of p is p^* (say), then the sequence of the DM's posterior distributions of p will become more and more concentrated in the neighborhood of p^*. In other words, the DM will asymptotically learn the true value of the parameter p. By 'open minded' I mean, roughly speaking, that the DM does not rule out as impossible any value of the parameter between zero and one.[10]

We are now in a position to consider the problem of model revision in the context of this example. Suppose that, in fact, the random variables in the X-process are not IID, but Markovian. This means that, in any period, the conditional distribution of the future random variables, given the history of the entire past, depends only on the value of the random variable in the preceding period. Will the DM ever discover this? To answer this question, I first observe that, under the hypothesis that the X-process is IID, in order to calculate the posterior distribution of the parameter after any history $H(t - 1)$, it is sufficient to know the number of ones in the history (in addition to the length, t, of the history). On the other hand, under the hypothesis that the X-process is Markovian, it is necessary (and sufficient) to know *three* numbers (in addition to t), namely, the initial value, $X(0)$, the number of ones that follow a zero, and the number of ones that follow a one. If the DM

[10] *Technical Note:* More precisely, 'open minded' means that the support of the prior is the entire unit interval. A more general statement of the theorem is that, if p^* is in the support of the prior, then conditional on $p = p^*$ the sequence of posterior probability distributions will converge weakly (i.e. in the weak* topology) to the point distribution with unit mass at p^*. A version of this theorem for more general sequences of IID random variables is due to Doob; see Diaconis and Freedman (1986).

believes that the X-process is IID, then he will not make these more detailed counts, except out of idle curiosity, and hence will not discover that the process is Markovian.

To see more clearly what is going on, recall that if the X-process is Markovian, then it can be parameterized by a starting probability and a transition matrix. To make things simpler, assume that the Markov process is *stationary and irreducible*, so we can dispense with the starting probability (it will be determined by the transition matrix). The transition matrix can be represented by a table:

		$X(t)$	
		1	0
$X(t-1)$	1	q	$1-q$
	0	r	$1-r$

The interpretation of the table is:

$$\text{Prob}\{X(t) = 1 \mid X(t-1) = 1\} = q$$

$$\text{Prob}\{X(t) = 0 \mid X(t-1) = 1\} = 1 - q$$

$$\text{Prob}\{X(t) = 1 \mid X(t-1) = 0\} = r$$

$$\text{Prob}\{X(t) = 0 \mid X(t-1) = 0\} = 1 - r$$

Here are some useful facts:

1. If $q = r$, then the Markov process is an IID process.
2. The Markov process is irreducible if q and r are strictly between 0 and 1.
3. If the Markov process is stationary, then for every period t:

$$\text{Prob}\{X(t) = 1\} = v = r/(1 - q + r)$$

It follows that *if the DM believes that the X-process is IID, but it is really Markovian with the above parameters* (*and with q distinct from r*), *then he will asymptotically come to believe the 'true parameter' of the IID process is $p = v$, and he will never learn the Markovian truth about the process.*

The only way out for this DM is for him not to be totally dogmatic about his belief that the X-process is IID, and experiment with the belief that it is Markovian. On the other hand, an extension of Doob's Theorem (see above) guarantees that *if the DM believes that the process is Markovian, and is open minded about the parameters, then if the process is in fact IID he will asymptotically learn that.*

We can, and should, consider further extensions of the process of model revision. It could be that the X-process is even outside the class of Markov processes, even though the DM does not envisage this at the start. For example, we might consider processes in which, for every period t, the conditional distribution of the future random variables, given the history of the entire past, depends only on the values of the random variable in the preceding k periods. This is sometimes called a *kth-order Markov process*. Clearly, for $h < k$, an hth-order Markov process is a special case of a kth-order Markov process. We can now imagine a never-ending process of model revision in which the DM experiments with kth-order Markov processes with larger and larger values of k.

We can now refine the question of how the DM might pursue such a process of model revision 'rationally'? I submit that ideally—if possible—*such a process should mimic the behavior of a Bayesian who is following the Savage paradigm from the beginning*, in other words whose prior beliefs do not rule out any of the models that the DM will eventually consider. This might seem to be an excessively demanding requirement, but the results that I have obtained so far suggest that it is attainable (Radner, 2000). Let $M(k)$ denote the model that the X-process is a kth-order Markov process, and recall that parameter space of the model $M(k)$ is essentially a subspace of that of the model $M(k + 1)$; in fact, I shall say that $M(k)$ *is* a subset of $M(k + 1)$. The basic idea is that every time the DM envisages a more general model, say from $M(k)$ to $M(k + 1)$, he extends his *prior distribution* to $M(k + 1)$ accordingly in a way that does not contradict the old prior within (i.e. conditional on) the subspace $M(k)$. He then continues to compute his posterior distributions as before, but with the new prior. However—and this is crucial for the interpretation of the process in terms of bounded rationality—he is not required to formulate a grand prior in advance, but only required to extend his prior incrementally, as he faces the prospect of using more and more general models.

Of course, in principle the DM can always eventually become aware of more general models than the one he is currently working with. For example, the class of all kth-order Markov processes does not exhaust the class of all stationary processes, and there are processes that are not even stationary. Here I must end the present discussion on a cautionary note. It can be shown that if the space of processes is 'too large' (i.e. if the parameter space is infinite-dimensional and 'very large' in some sense), it may not be possible to formulate a prior distribution on it that is 'open minded' in a useful way. Thus there will be limits to how well the process of model revision that I have roughly described above can guarantee that the DM will asymptotically learn the 'truth' about the process that he is observing.[11]

5.5 Failures of Logical Omniscience

Up to this point is has been assumed—if only implicitly—that the DM has no difficulty performing mathematical calculations or other logical operations. In particular, having formulated a decision model, he will be able to infer what it implies for his optimal strategy. As has already been pointed out, this assumption is absurd, even for small-world models, except for 'Mickey Mouse' problems that are constructed for textbooks and academic articles. The crux of the matter is that, in any even semi-realistic decision problem, *the DM does not know all of the relevant logical implications of what he knows.* This phenomenon is sometimes called *the failure of logical omniscience*.

In the Savage paradigm, DM's are uncertain about events such as 'It will rain tomorrow in New York City', or 'If most countries does not take action to reduce the emissions of carbon dioxide into the atmosphere, then the average atmospheric temperature of the Earth will increase by 5°F in the next fifty years.' Examples of the failure of logical omniscience are:

1. A DM who knows the axioms of arithmetic is uncertain about whether they imply that 'The 123rd digit in the decimal expansion of pi is 3', unless he has a long time to do the calculation and/or has a powerful computer with the appropriate software.
2. Twenty years ago, a DM who knew the axioms of arithmetic was still uncertain about whether they imply Fermat's Last Theorem.

The following examples are closer to practical life, and possibly more intimidating:

3. Given all that a DM know about the old and new drugs for treating a particular disease, what is the optimal policy for conducting clinical trials on the new ones?
4. Given all that is known, theoretically and empirically, about business organizations in general, and about telecommunications and AT&T

[11] *Technical Note:* Although the set of all models $M(k)$, with k finite, does not exhaust the class of all stationary processes, one can approximate any stationary X-process by a stationary $M(k)$ process, provided the random variables take values in some finite set. However, it is an open question whether it is possible to put a prior distribution on the class of all such stationary processes in such a way that a Bayesian will asymptotically learn what the true process is.

in particular, should AT&T reorganize itself internally, and if so, how? (These examples are taken from Radner, 1997, p. 133.)

In commenting on this problem in 1954, Savage wrote:

> The assumption that a person's behavior is logical is, of course, far from vacuous. In particular, such a person cannot be uncertain about decidable mathematical propositions. This suggests, at least to me, that the tempting program sketched by Polya of establishing a theory of the probability of mathematical conjectures cannot be fully successful in that it cannot lead to a truly formal theory . . . (Savage, 1954, fn, p. 7)[12]

In spite of some interesting efforts (Lipman, 1999; Dekel *et al.*, 1998; Modica and Rustichini, 2000), I am not aware of significant progress on what it means to be rational in the face of failures of logical omniscience.

6. *In What Sense(s) is a Boundedly Rational Decision-maker Rational?*

In the term 'bounded rationality', the first word reminds us that bounds on the cognitive abilities of human decision-makers, even when aided by modern computers, limit the extent to which their decision-making can conform to the Savage paradigm. On the other hand, the second word suggests that there is some sense in which such decision-making can be said to be rational. In this final section of the paper, I want to examine this presupposition.

As noted in Section 4, in some cases the Savage paradigm can be extended to incorporate the costs of, and constraints on, decision-making without abandoning the assumption of optimization. I called this 'costly rationality'. In fact, this is not really an extension of the paradigm, but merely a more sophisticated and detailed interpretation of the Savage model. In that sense, costly rationality retains the full meaning of 'rationality' in the Savage sense.

Nevertheless, this reinterpretation can be carried only so far in realistic situations. The DM must be content to work with highly simplified models, which in turn leads to a process of model revision—a potentially unending process. In Section 5.4 I sketched how, in some simple cases, the DM can devise a process of inference and model revision whose results eventually mimic those achieved by a Bayesian who is able to formulate a complete model from the beginning. In such cases, one may be justified in calling the

[12] For further discussion and references, see the 1972 edition of Savage's book.

process of decision-making 'rational'. I would like to think that, when applicable, this approach provides an instance of Simon's idea of 'rationality as process' (Simon, 1978).

Finally, I described the problems caused by 'failures of logical omniscience'. Here we come to the hard core of truly bounded rationality, which thus far presents the most difficult challenge to decision-theorists who strive to characterize the 'rationality' in 'bounded rationality'.

Acknowledgements

The present paper is a further development of part of Radner (1997), and as such it inevitably repeats many ideas in that paper. The reader will notice that I quote liberally from Savage's book, *The Foundations of Statistics*. It turns out that he anticipated most of the issues of costly and bounded rationality, and had something to say about them, even if he supplied few formal solutions. Thus in many instances I could not find a better way to express the relevant ideas than to use his own words. It would be too large a project to acknowledge my debts to all my friends and colleagues who have tried to help me think about this problem (some may judge unsuccessfully). At the very least, I must acknowledge the influences of Jacob Marschak, Herbert Simon and L. J. Savage, who introduced me to the topic. In addition, I thank Jim March and Sid Winter for stimulating discussions over the years, and specifically for comments on a previous draft of this paper.

References

The references cited in the body of the article have historical interest, provide an overview of a topic discussed and/or provide a key to other literature. The following references provide additional information on the application of notions of bounded rationality in economics and management: Arrow (1974), Majumdar (1998), McGuire and Radner (1986), Nelson and Winter (1982), Newell and Simon (1972), Radner (1992), Shapira (1997), Simon (1981), Williamson and Winter (1991) and Winter (1987).

Arrow, K. J. (1951), 'Alternative Approaches to the Theory of Choice in Risk-taking Situations,' *Econometrica*, 19, 404–437.

Arrow, K. J. (1974), *The Limits of Organization*. Norton: New York.

Basu, A., A. Bose and J. K. Gosh (1990), 'An Expository Review of Sequential Design and Allocation Rules,' Technical Report 90-08, Purdue University.

Bolton, P. and M. Dewatripont (1994), 'The Firm as a Communication Network,' *Quarterly Journal of Economics*, **109**, 809–839.

Clark, J. M. (1918), 'Economics and Modern Psychology, I,' *Journal of Political Economy,* **26**, 1–30.

Conlisk, J. (1996), 'Why Bounded Rationality,' *Journal of Economic Literature*, **34**, 669–700.

Cournot, A. (1838), *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Hachette: Paris; trans. N. T. Bacon. R. D. Irwing: Homewood, IL.

de Finetti, B. (1937), 'La Prévision: Ses Lois Logiques, ses Sources Subjectives,' *Annales de l'Institut Henri Poincaré*, **7**, 1–68.

Dekel, E., B. Lipman and A. Rustichini (1998), 'Recent Developments in Modeling Unforeseen Contingencies,' *European Economic Review*, **42**, 523–542.

Diaconis, P. and D. Freedman (1986), 'On the Consistency of Bayes Estimates,' *Annals of Statistics*, **14**, 1–26.

Dutta, P. K. and R. Radner (1995), 'Moral Hazard,' in R. Aumann and S. Hart (eds), *Handbook of Game Theory*, Vol. 2. North-Holland: Amsterdam, pp. 869–903.

Hald, A. (1990), *A History of Probability and Statistics and their Applications before 1750.* Wiley: New York.

Harsanyi, J. C. (1967–68), 'Games with Incomplete Information Played by Bayesian Players,' *Management Science*, **14**, 159–182, 320–334, 486–502.

Jackson, M. O., E. Kalai and R. Smorodinsky (1999), 'Bayesian Representation of Stochastic Processes under Learning: de Finetti Revisited,' *Econometrica*, **67**, 875–894.

Jones, R. A. and J. Ostroy (1984), 'Flexibility and Uncertainty,' *Review of Economic Studies*, **51**, 13–22.

Kaldor, N. (1934), 'The Equilibrium of the Firm,' *Economic Journal*, **44**, 70–71.

Knight, F. H. (1921), *Risk, Uncertainty, and Profits.* Houghton Mifflin: Cambridge, MA (reprinted 1957, Kelley and Millman: New York).

Kreps, D. (1979), 'A Representation Theory for "Preference for Flexibility",' *Econometrica*, **47**, 565–577.

Kreps, D. (1992), 'Static Choice in the Presence of Unforeseen Contingencies,' in P. Dasgupta *et al.* (eds), *Economic Analysis of Markets and Games*. MIT Press: Cambridge, MA, pp. 258–281.

Kreps, D. and R. B. Wilson (1982), 'Sequential Equilibria,' *Econometrica*, **50**, 863–894.

Lipman, B. L. (1999), 'Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality,' *Review of Economic Studies*, **66**, 339–361.

Machina, M. (1987), 'Choice Under Uncertainty: Problems Solved and Unsolved,' *Journal of Economic Perspectives*, **1**, 121–154.

Machina, M. (1989), 'Dynamic Consistency and Non-expected Utility Models of Choice Under Uncertainty,' *Journal of Economic Literature*, **27**, 1622–1668.

Majumdar, M. K. (ed.) (1998), *Organizations with Incomplete Information*. Cambridge University Press: Cambridge.

March, J. G. and H. A. Simon (1958), *Organizations*. Wiley: New York.

Marschak, J. (1955), 'Elements for a Theory of Teams,' *Management Science*, **1**, 127–137.

Marschak, J. and R. Radner (1972), *Economic Theory of Teams*. Cowles Foundation and Yale University Press: New Haven, CT.

Marschak, T. A. and C. B. McGuire (1971), 'Lecture Notes on Economic Models for Organization Design,' Unpublished manuscript, University of California: Berkeley, CA.

Marschak, T. A. and S. Reichelstein (1998), 'Network Mechanisms, Informational Efficiency, and Hierarchies,' *Journal of Economic Theory*, **79**, 106–141.

McFadden, Daniel (1999), 'Rationality for Economists?,' *Journal of Risk and Uncertainty*, **19**, 73–105.

McGuire, C. B. and R. Radner (eds) (1986), *Decision and Organization*, 2nd edn. University of Minnesota Press: Minneapolis, MN.

Mintzberg, H. (1989), *Mintzberg on Management*. Free Press: New York.

Modica, S. and A. Rustichini (1999), 'Unawareness: A Formal Theory of Unforeseen Contingencies,' *Games and Economic Behavior*, **29**, 15–37.

Mount, K. and S. Reiter (1998), 'A Modular Network Model of Bounded Rationality,' in M. K. Majumdar (ed.), *Organizations with Incomplete Information*. Cambridge University Press: Cambridge, pp. 306–340.

Nash, J. F. (1951), 'Non-cooperative Games,' *Annals of Mathematics*, **54**, 286–296.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton University Press: Princeton, NJ.

Nelson, R. R. and S. G. Winter (1982), *An Evolutionary Theory of Economic Change.* Harvard University Press: Cambridge, MA.

Newell, A. and H. A. Simon (1972), *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.

Osborne, M. and A. Rubinstein (1994), *A Course in Game Theory*. MIT Press: Cambridge, MA.

Peters, T. J. and R. H. Waterman (1982), *In Search of Excellence*. Harper & Row: New York.

Radner, R. (1961), 'The Evaluation of Information in Organizations,' in *Proceedings of the Fourth Berkeley Symposium on Probability and Statistics*. University of California Press: Berkeley, CA, Vol. 1, pp. 491–530.

Radner, R. (1962), 'Team Decision Problems,' *Annals of Mathematical Statistics*, **33**, 857–881.

Radner, R. (1975a), 'Satisficing,' *Journal of Mathematical Economics*, **2**, 253–262.

Radner, R. (1975b), 'Planning Under Uncertainty: Recent Theoretical Developments,' in M. Bornstein (ed.), *Economic Planning, East and West*. University of Michigan Press: Ann Arbor, MI.

Radner, R. (1992), 'Hierarchy: The Economics of Managing,' *Journal of Economic Literature*, **30**, 1382–1415.

Radner, R. (1997), 'Bounded Rationality, Indeterminacy, and the Managerial Theory of the Firm,' in Z. Shapira (ed.), *Organizational Decision Making*. Cambridge University Press: Cambridge, pp. 324–352.

Radner, R. (2000), 'A "Bayesian" Analysis of Model Revision for a Class of Stationary Processes,' Stern School, New York University (unpublished).

Radner, R. and J. E. Stiglitz (1983), 'A Nonconcavity in the Value of Information,' in M. Boyer and R. E. Kihlstrom (eds), *Bayesian Models in Economic Theory*. North-Holland: Amsterdam, pp. 33–52.

Radner, R. and T. Van Zandt (1992), 'Information Processing in Firms and Returns to Scale,' *Annales d'Economie et de Statistique*, no. 25/26, 265–298.

Rubinstein, A. (1998), *Modeling Bounded Rationality*. MIT Press: Cambridge, MA.

Savage, L. J. (1954), *The Foundations of Statistics*. Wiley: New York (2nd edn, 1972, Dover: New York).

Selten, R. (1975), 'Reexamination of the Perfectness Concept for Equlibrium Points in Extensive Games,' *International Journal of Game Theory*, **4**, 25–55.

Shapira, Z. (ed.) (1997), *Organizational Decision Making*, Cambridge University Press: Cambridge.

Simon, H. A. (1947), *Administrative Behavior*. Free Press: New York.

Simon, H. A. (1955), 'A Behavioral Model of Rational Choice,' *Quarterly Journal of Economics*, **69**, 99–118.

Simon, H. A. (1957), *Models of Man*. Wiley: New York.

Simon, H. A. (1972), 'Theories of Bounded Rationality,' in C. B. McGuire and R. Radner (eds), *Decision and Organization*, 2nd edn (1986). University of Minnesota Press: Minneapolis, MN, pp. 161–176.

Simon, H. A. (1978), 'Rationality as Process and Product of Thought,' *American Economic Review*, **68**, 1–16.

Simon, H. A. (1981), *The Sciences of the Artificial*, 2nd edn. MIT Press: Cambridge, MA.

Sims, C. (1971), 'Distributed Lag Estimation when the Parameter Space is Explicitly Infinite Dimensional,' *Annals of Mathematical Statistics*, **42**, 1622–1636.

Van Zandt, T. (1998), 'Decentralized Information Processing in the Theory of Organizations,' in M. Sertel (ed.), *Economic Design and Behavior*, *Proc. XIth World Congress of the International Economics Association*, Vol. IV. Macmillan: London.

Van Zandt, T. (1998), 'Organizations with an Endogenous Number of Information Processing Agents,' in M. K. Majumdar (ed.), *Organizations with Incomplete Information*. Cambridge University Press: Cambridge.

Van Zandt, T. (1998), 'Real-time Hierarchical Resource Allocation,' Unpublished manuscript Northwestern University.

Van Zandt, T. (1999), 'Real-time Decentralized Information Processing as a Model of Organizations with Boundedly Rational Agents,' *Review of Economic Studies*, **66**, 633–658.

Van Zandt, T. and R. Radner (2000), 'Real-time Decentralized Information Processing and Returns to Scale,' *Economic Theory* (forthcoming).

Wald, A. (1947), *Sequential Analysis*. Wiley: New York.

Wald, A. (1950), *Statistical Decision Functions*. Wiley: New York.

Williamson, O. E. and S. G. Winter (eds) (1991), *The Nature of the Firm*. Oxford University Press: New York.

Winter, S. G. (1987), 'Competition and Selection,' in J. Eatwell, J. Milgate and P. Newman (eds), *The New Palgrave*. Stockton Press: New York, pp. 545–548.