

How to Decide How to Decide How to...: Modeling Limited Rationality

Author(s): Barton L. Lipman

Reviewed work(s):

Source: *Econometrica*, Vol. 59, No. 4 (Jul., 1991), pp. 1105-1125

Published by: [The Econometric Society](#)

Stable URL: <http://www.jstor.org/stable/2938176>

Accessed: 12/06/2012 17:25

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The Econometric Society is collaborating with JSTOR to digitize, preserve and extend access to *Econometrica*.

HOW TO DECIDE HOW TO DECIDE HOW TO . . . : MODELING LIMITED RATIONALITY¹

BY BARTON L. LIPMAN

It seems inconsistent to model boundedly rational action choice by assuming that the agent chooses the optimal decision procedure. This criticism is not avoided by assuming that he chooses the optimal procedure to choose a procedure to . . . to choose an action. I show that, properly interpreted, this regress, continued transfinitely, generates a model representing the agent's perception of all his options including every way to refine his perceptions. In this model, the agent surely must choose the perceived best option. Hence it is not inconsistent to model limited rationality by assuming that the agent uses the "optimal" decision procedure.

KEYWORDS: Limited rationality, bounded rationality, infinite regress, optimal decision procedures.

"A person required to risk money on a remote digit of π would, in order to comply fully with the theory [of expected utility] have to compute that digit, though this would really be wasteful if the cost of computation were more than the prize involved. For the postulates of the theory imply that you should behave in accordance with the logical implications of all that you know. Is it possible to improve the theory in this respect, making allowance within it for the cost of thinking, or would that entail paradox?"

—Savage (1967).

1. INTRODUCTION

IT IS NO LONGER NOVEL to suggest that standard economic models make overly strong rationality assumptions.² Unfortunately, taking account of this criticism is no easy task. A recent approach in the literature is based on the idea that if finding the best action is costly, then the best way to decide on an action involves trading off these decision-making costs with the benefits to improving the choice of an action. Thus imperfect action choices are endogenously generated by optimal decision procedures.³ Unfortunately, this approach seems inconsistent. If the agent chooses actions imperfectly because it is costly to do better, why doesn't he choose decision procedures imperfectly as well? We might try to model imperfect choice of a decision procedure by analyzing the choice of decision procedures to pick decision procedures to pick actions. Again, though, it is not obvious why we would simply assume optimality at this

¹ This paper contains the material from Section II of my working paper "How to Decide How to Decide How to . . . : Limited Rationality in Decisions and Games." I thank Mark Bagnoli, Sugato Bhattacharyya, Larry Blume, Ed Green, Maria Herrero, Debra Holt, Steve Spear, Sanjay Srivastava, Spyros Vassilakis and two other referees, and Andreu Mas-Colell for helpful comments. Of course, I am responsible for any errors. I wish to acknowledge financial support from the National Science Foundation through Grant SES-8520296.

² See Simon (1955, 1976), Selten (1978), and Binmore (1987, 1988).

³ Papers in this spirit include Green (1982), Rubinstein (1986), Abreu and Rubinstein (1988), and Lipman and Srivastava (1990a, 1990b).

level either, driving us up yet another level. Clearly, we are headed for an infinite regress.⁴

Without a solution to this infinite regress problem, efforts to find a useful model of limited rationality are unlikely to succeed. Any model must either make direct assumptions on how agents respond to their environment or derive this behavior from optimality-based considerations.⁵ Direct assumptions on behavior are naturally subject to the criticism of ad hocery. Derivation of behavior from optimality either assumes away limited rationality or runs into this infinite regress problem.

This paper has two purposes. First, I present a model of limited rationality based on the approach described above. The main purpose of the paper is the second one: within the context of the model, I show that the infinite regress problem, properly interpreted, does not mean that modeling limited rationality by the choice of optimal decision procedures is doomed. That is, in a particular sense, we can “solve” the infinite regress problem.⁶

To explain the model and the solution to the infinite regress problem requires consideration of some difficult philosophical questions. In the remainder of this section, I address these issues, explain the idea of the model and the solution to the infinite regress problem, and briefly survey the relevant literature.

The first philosophical issue that must be addressed is the most basic: what do we mean when we say that rationality is “limited”? I believe most economists would reply that limited rationality means that the agent does not choose the optimal action because it is too difficult to compute that action. That is, in a complex environment, the agent may not know the optimal action. The obvious next question is: why is rationality limited? If the agent knows his preferences, knows the feasible set, and knows how the optimal action is defined given these objects, why does he not know the optimal action? The reason is that knowing a fact does not mean that one knows all the logical implications of that fact. For example, the reader is likely to know the axioms of set theory and the rules of logical inference, but is unlikely to know all theorems of set theory, though these are logically implied. Similarly, the agent can fail to recognize the appropriate action because this requires him to process his information. While his information mathematically defines the appropriate choice, it is often difficult to learn these implications.

The philosophy literature has long recognized the problem, dubbed *logical omniscience* by Hintikka (1975), that standard representations of knowledge imply that an agent knows all logical implications of his knowledge. The solution

⁴ This infinite regress problem has long been recognized. See, for example, Winter (1975) and the references therein.

⁵ To be sure, evolutionary models (see Blume and Easley (1989), for example), in a sense, derive behavior from considerations other than optimality. On the other hand, such models also must make assumptions on behavior to generate the evolutionary process.

⁶ After completing this paper, I discovered the related work of Mongin and Walliser (1988). Their approach to the infinite regress problem is quite different from mine as are their conclusions.

proposed by Hintikka, among others, is to allow *impossible possible worlds*.⁷ In other words, the agent believes possible something which is logically not possible. For example, the agent may know that p is true, know a rule of inference which yields q when applied to p , and yet believe that it is possible that q is false.

The approach I take is based on this idea. I assume that if the agent does not know what some fact implies, then he believes that the true implication depends on the state of the world. Among the states of the world given positive probability by the agent may be states which are logically inconsistent. Computation or thought in my model is the process by which the agent learns about the implications of his information. Essentially, he recognizes the inconsistency of some states and eliminates them as possible.⁸

In short, when we write down a model of an agent in some situation, we specify a feasible set and preferences for the agent, thus defining his optimal choice. If we believe that the real agent might not choose this optimum because he is only limitedly rational, we are arguing that the agent does not perceive the model in exactly the way we have written it down. We must be arguing either that the agent is less certain than the model assumed regarding his feasible set and preferences or else that he faces subjective uncertainty regarding what his information about the feasible set and his preferences tells him he should do. Either way, we have not correctly specified his information or, more loosely, his “perception” of the situation.

This immediately suggests another question: is that all that “limited rationality” means? In other words, if we correctly and completely specify an agent’s perception of his situation, can we assume that he is “completely rational” given this perception? I would argue that we have no other choice. We certainly cannot hope to predict behavior which is based purely on whim. Yet what else could explain the behavior of an agent who perceives some action to be the best one available for him and chooses something else anyway? Perceptions may be highly imperfect, even totally inaccurate, but given an agent’s perception of his world, we must assume that he chooses what he perceives to be best for him.⁹ Given this, the key to modeling limited rationality is modeling the agent’s perceptions, not postulating a “boundedly rational” choice procedure. This point is crucial to the motivation for everything which follows: I focus on these perceptions, taking as fundamental that they are what determines choice.

This leads to the final question: haven’t we just assumed away the infinite regress problem? If we only have to correctly specify an agent’s perception of

⁷ See also Fagin and Halpern (1985) and the references cited there.

⁸ Hacking’s (1967) reply to the Savage quote with which I began appears to be the first discussion of this approach.

⁹ The earliest discussion of this view known to me is Popper (1967). As he points out, we even try to “rationalize” the behavior of lunatics by supposing that their choices are what they perceive as best for them given their preferences and their sometimes rather bizarre perceptions of their environment.

his situation, where does infinite regress come in? Infinite regress arises when we try to construct the agent's perceptions. To see the point, suppose we are trying to model an agent who has a set of feasible actions $A = \{0, 1\}$. Suppose the agent gets \$100 if he chooses $a = 0$ and he gets $f(1)$ dollars if he picks $a = 1$, where $f(\cdot)$ is a very complex function that we write down for the agent.¹⁰ Clearly, if we assume complete rationality (and that the agent prefers more money to less), the prediction is trivial. The agent chooses $a = 0$ iff $100 \geq f(1)$. Since $f(\cdot)$ is written down in front of the agent, he must know the function.

Of course, if we believe that the agent's rationality is limited, we may believe that it is very difficult for him to compute the value of $f(1)$. What we are saying, then, is that he knows the function in the sense that he sees what it says, but that he does not know it in the sense of knowing all the ordered pairs the definition of the function implies. Hence to predict the agent's choice accurately, we must specify his perception of what $f(1)$ may be. That is, we must treat $f(\cdot)$ as random from the agent's point of view. Once we add in this uncertainty, though, we must recognize that there are more options available to the agent than simply choosing $a = 0$ or $a = 1$ —he may try to compute $f(1)$, for example. More generally, the agent has available ways to resolve his subjective uncertainty and this must be made part of the model if we are going to predict his behavior.

Bringing in these options involves constructing the set of decision procedures the agent has available and replacing the feasible set A with this new, larger feasible set. This new set, presumably, includes choices like computing $f(1)$ on a calculator or computer, trying to approximate the function $f(\cdot)$ by some more easily computed function, differentiating $f(\cdot)$ to see how the function behaves around 1, etc. One approach we could take would be to assume that the agent chooses the optimal decision procedure. However, we might suspect that he is quite uncertain about how hard it will be to perform some of these computations or that he doesn't know whether his calculator is capable of computing $f(1)$. In short, now that we've enlarged the set of options our model of the agent allows, we must recognize that the agent may be uncertain about these new options also. The important point to keep in mind is that it is precisely our intuitive sense that choosing the optimal decision procedure is "hard" that leads to the inclusion of this additional uncertainty. Naturally, once we include this uncertainty, we see that the agent may have options available to help him choose among decision procedures—for example, he might purchase a book which contains various algorithms, learn a new programming language, etc. This leads us to construct the set of decision procedures to pick decision procedures, enlarging the choice set yet again. If, again, we believe the agent to be uncertain about these new options, we must enlarge the uncertainty, and so on.

Thus the infinite regress problem, restated, is this: can we construct a model in which this regress stops? More precisely, can we construct a model in which

¹⁰ For example, we may tell him that $f(1) = 0$ if a given 500-digit number is prime and is 200 otherwise.

we fully represent the agent's uncertainty about each option available to him and, at the same time, make every way he could resolve his uncertainty an option? We can easily do this if we assume that the agent is perfectly rational, since there is no subjective uncertainty in this case. Similarly, we can easily do this if we assume that the agent has no options available to resolve his uncertainty. The real question is whether we can do this without making the problem trivial by either assuming away the subjective uncertainty generated by limited rationality or assuming away the agent's ability to compute to improve his choices.

To put it differently, we are looking for a fixed point.¹¹ Loosely, suppose we have an operator, say U , which gives us the uncertainty associated with a particular set of options and another operator D which gives us the options available to the agent to resolve some given uncertainty. The question then becomes whether there exists a set of options O such that $O = \mathcal{D}(O) \equiv D(U(O))$. Viewed this way, the "infinite regress" is simply the sequence $A, \mathcal{D}(A), \mathcal{D}(\mathcal{D}(A)), \dots$, where A is the set of actions available to the agent. The key question about this regress is whether it converges to a fixed point of $\mathcal{D}(\cdot)$. As this description suggests, the view taken here is *not* that the agent has to think through some infinite sequence to try to decide what action to pick. Instead, the consideration of the infinite sequence is our effort (as modellers) to find an appropriate model for predicting the agent's action choice. The "appropriate" model is one in which we have fully represented the agent's options and his perception of his options. Once we have done that, as argued above, we must close the model by assuming that the agent chooses the option which he perceives to be best.

The main result is that a fixed point does exist. One difficulty in demonstrating this fact is the following. Let $D_0 = A$, $D_1 = \mathcal{D}(A)$, $D_2 = \mathcal{D}(D_1)$, and so on for every finite number n . As I show in the next section, the countable sequence $\{D_n | n = 1, 2, \dots\}$ does *not* converge to a fixed point in general. One could impose a number of different assumptions which would guarantee that the countable sequence would converge, most of which are rather strong. However, the primary issue is the existence of a sequence converging to a fixed point, not how long the convergent sequence in question is. Hence the important point is that, even without these additional assumptions, if we consider "long enough" (transfinite) sequences, we always get convergence to a fixed point.¹² This is shown in Section 3.

Related Literature. There is a rapidly growing literature on bounded rationality. Much of this work focuses on the computational complexity of strategies or finding equilibria. See, e.g., Abreu and Rubinstein (1988), Kalai and Stanford

¹¹ I am grateful to a referee, Spyros Vassilakis, for recommending this explanation.

¹² This does not mean that there is no reason at all to be interested in the length of the convergent sequence. For example, if one is willing to impose a certain structure (such as a specific topology) on each of the D 's in the sequence, one may wish to know whether the fixed point also has this structure. For many kinds of structure, the answer to this question will hinge on the length of the convergent sequence. However, these questions are not the ones on which I focus here.

(1988), Ben-Porath (1989), Canning (1988), Gilboa and Zemel (1989), and Spear (1989).¹³ Other approaches include evolutionary or dynamic learning models (such as Fudenberg and Kreps (1988), Blume and Easley (1989), and Canning (1989)) and models with limited reasoning (Aumann (1988), Geanakoplos (1989) and Brandenburger, Dekel, and Geanakoplos (1989)). Also, several authors have provided “resolutions” of apparent infinite regress problems. The primary use of this technique in economics and game theory is based on Mertens and Zamir’s (1985) pioneering work on beliefs about the beliefs of others. For the reader familiar with Mertens and Zamir, I give a comparison of my construction with theirs in the following section. Rationalizability (Bernheim (1984), Pearce (1984)) is also based explicitly on analyzing an infinite regress.¹⁴ Such constructions, including transfinite sequences, have been used frequently in the philosophy literature. For example, Kripke (1975) used this approach to deal with an infinite regress problem which arises in defining truth. Finally, Vassilakis (1989) uses category theory to give a general theory of fixed points based on recursive constructions, such as the construction here or that of Mertens and Zamir (1985).

2. THE MODEL: COUNTABLE SEQUENCES

A precise construction of the sequences involved requires attention to some uninteresting technical issues. Rather than explain these points in detail, I give a more heuristic construction in the text and relegate the formal details to the Appendix.

The agent has a set of feasible actions, which I denote A or, interchangeably, D_0 . The agent is unsure about which action is best for him, but has some perception of how “good” each action is. If he had to choose an action without computation, he would choose the action which is best according to his initial perception. It is not obvious what kind of structure the agent’s perceptions should be assumed to have since this is the agent’s primitive, gut-level view of his situation. For expositional purposes, I impose a great deal of structure on the agent’s perceptions. However, as we will see, essentially none of this structure is needed for the main result. I represent the agent’s perception with a set S_0 of possible states and a utility function $u_0: D_0 \times S_0 \rightarrow \mathbf{R}$. With this formalization, the uncertainty of the agent regarding the best choice in D_0 is represented by his uncertainty about the true state in S_0 and how this affects u_0 . Through the remainder of the paper, I avoid additional assumptions on the agent’s perceptions. As we will see, this is partly responsible for the fact that countable sequences are not enough to converge to a fixed point. Fortunately, though, longer sequences do converge without strong assumptions on perceptions. I let $P_0 = (D_0, u_0, S_0)$ denote the initial problem—i.e., our initial attempt at a model describing the agent’s problem.

¹³ I do not discuss computational issues per se, but there is clearly room to consider such issues in this framework. The complexity literature characterizes the costs of various computations, while I focus on characterizing choice of computations as a function, in part, of their costs.

¹⁴ Lipman (1990a) shows that with infinite strategy sets and discontinuous payoff functions, rationalizability also requires a transfinite construction.

To model the computation the agent may carry out to refine his perceptions, I introduce a set L , which will be called the *language*. This set is the language in which thought or computation is carried out. For example, if we think of computation as the use of a computer, we might think of L as the set of strings of 0's and 1's. If we view computation as deriving implications in a formal logical system, we would assume that L is the set of well-formed formulas of the logical system. For convenience, an element of L will be referred to as a *word*. To avoid trivialities, I assume that L has at least two elements. A computation is a function c which produces a word as a function of the state.¹⁵ To explain the intuition, let us return to the example in the introduction. If the agent is uncertain about the value of $f(1)$, he may wish to try computing this value. More formally, the agent perceives the value of $f(1)$ as being a function, say $g(s)$, of the state of the world s , where S_0 is the domain of g . If he computes the value of $f(1)$ exactly, this corresponds to learning exactly what the true s is. That is, it is as if he observes the value of $c(s)$ where c is invertible.

For any set of states, say S , let the set of possible computations using S be $C(S)$. (An exact definition is given in a moment.) Using this, we can construct the set of decision procedures, which I write as $D(S)$. The details of the construction are given in the Appendix. Intuitively, a decision procedure begins with an initial choice of either an action $a \in A$ or a computation. If the procedure chooses an action, it is complete and the procedure is interpreted as choosing that action without performing computations.¹⁶ Thus this "procedure" is also an element of D_0 . If, instead, the procedure's initial choice is a computation, the outcome of that computation is observed. This generates a set of possible outcomes, each of which is a possible "history" of the computation procedure. I refer to these as histories of length one. After each such history, the procedure must again specify either an action or another computation. If actions are chosen for each possible history, again, the procedure is completely specified. Otherwise, we continue as above.¹⁷ The set of procedures, $D(S)$, is the set of functions that can be constructed this way.

By construction, every history generated by a decision procedure has finite length. However, I do not assume that there is a finite bound on the length of the histories generated by a given procedure. Hence I allow procedures that compute without stopping.¹⁸ This assumption simplifies but is not necessary for the analysis. It is not hard to show that all the results below hold if we only allow procedures for which the number of steps of computation is bounded,

¹⁵ An equivalent formulation would be to let computations take a word as input and give some word as output as a function of the state. My approach simply treats computing the same function on different inputs as different computations.

¹⁶ Throughout, I treat choice and evaluation as separate. An agent can always choose any action without computation. However, making a *good* choice may require evaluating several options—computing—before choosing. Since computation is costly, making good choices is costly.

¹⁷ It is important that a decision procedure not include a specification of what it does on histories it cannot generate. Without this, Theorem 2 does not hold.

¹⁸ One argument in favor of this approach is that in standard computability theory, there is no algorithm for determining whether a given procedure halts. While I do not focus on computability considerations, this certainly suggests that one should not assume that the agent realizes which procedures halt and which do not.

even if we require a uniform bound across procedures. As discussed below, there are some differences in the interpretation of the results with such bounds.

Finally, then, the set of decision procedures for improving the choice of an action is $D_1 = D(S_0)$, where $D_0 \subset D_1$. (Throughout, I use \subset to denote strict inclusion.) Since the agent can choose any decision procedure in D_1 , we must extend “perceptions” to this set. In doing so, it is important to keep in mind that S_0 , by assumption, reflected the uncertainty over the appropriate action to pick and over the outcomes of computations; there is no reason why it must reflect all uncertainty regarding the best choice of a decision procedure. For example, the payoff to a decision procedure should include whatever costs are associated with the computations the procedure uses and the agent may be uncertain about these costs. It is precisely this “extra” subjective uncertainty which leaves us uncomfortable with the assumption that the agent chooses the optimal decision procedure. Thus the agent’s perception of his options in D_1 are given by a utility function $u_1: D_1 \times S_0 \times S_1 \rightarrow \mathbf{R}$ where S_1 is a state set reflecting the additional subjective uncertainty that arises when we enlarge the set of options to D_1 . I require that the agent’s perceptions be well defined in the sense that the enlargement of the set of options does not affect the agent’s perceptions of the original options. That is, u_1 is an extension of u_0 in the sense that for any $d \in D_0$ and $s = (s_0, s_1) \in S_0 \times S_1$, $u_1(d, s) = u_0(d, s_0)$. Recall that this construction is being carried out by the modeller, not the agent, and so the process of construction should not affect the agent’s perceptions. I let $P_1 = (D_1, u_1, S^1)$, where $S^1 = S_0 \times S_1$.

Once we add in S_1 , this enlarges the set of computations to $C(S^1)$ and correspondingly enlarges the set of decision procedures to $D_2 = D(S^1)$. Again, this necessitates introducing an additional state set S_2 and extending preferences to u_3 . I will not explicitly construct an operator $U(D)$ giving the agent’s perceptions of each D . Instead, I take a sequence of state sets as given and extend the u function as necessary. As we will see, this approach is sufficient for ensuring that a fixed point exists.

I have delayed giving a precise definition of $C(S)$ to be able to better motivate the definition. A natural candidate for the definition is to let

$$C(S) = \{c | c: S \rightarrow L\}.$$

However, it is easy to show that no fixed point exists if we define $C(S)$ this way. To see this, suppose that using this definition, there is a fixed point, D , where u and S summarize all the uncertainty the agent has about D . With $C(S)$ defined this way, for any $\hat{S} \subseteq S$ and any $w^* \in L$, there is a computation c with $c(s) = w^*$ iff $s \in \hat{S}$. Hence the set of computations is at least as large as the power set of S . For each possible computation, though, we can construct a decision procedure which uses only that computation. Hence D is strictly larger than the power set of S . But then how can S capture all the uncertainty the agent might have about D ? I restrict the set of feasible computations by requiring that no computation generate “too much” information. This can be viewed as a restriction on the complexity of the calculations the agent can carry out.

More precisely, suppose $S = \prod_{\beta \in \alpha} S_\beta$ where α is some set of indices. For each $\gamma \in \alpha$, let $\rho_\gamma(s)$ give the projection of s onto S_γ . For any $c: S \rightarrow L$, I will say that c uses information in S_β if there exists $s, s' \in S$ with $\rho_\gamma(s) = \rho_\gamma(s')$ for all $\gamma \neq \beta$ but $c(s) \neq c(s')$. That is, c uses information in S_β if the outcome of the computation may depend on the β th component of s . For any set B , let $\#B$ denote the cardinality of B . I let

$$C(S) = \{c | c: S \rightarrow L \text{ and } \#\{\beta | c \text{ uses information in } S_\beta\} \leq \xi\}$$

where ξ is some cardinal number. Intuitively, ξ is a bound on the number of aspects of his uncertainty the agent can learn about using one computation. Since ξ can be a very large infinite cardinal, this assumption does not seem to be too restrictive. In Remark 5 below, I briefly explain how one can substantially relax this assumption.

Clearly, these definitions generate a countable sequence of models, P_n , $n = 1, 2, \dots$. We begin with $A = D_0$. This generates u_0 and S_0 , in turn giving us $C(S_0)$ and $D_1 = D(S_0)$. We then enlarge S_0 to capture the additional uncertainty, giving us $S^1 = S_0 \times S_1$ and u_1 . This in turn gives us $D_2 = D(S^1)$, etc.

REMARK 1: The reader familiar with Mertens and Zamir's (1985) work on beliefs about beliefs may find a comparison useful at this point. For ease of exposition, I only describe a simple two-player version of their construction. Mertens and Zamir begin with a set of possible states, say Θ , analogous to my A . They then consider the set of probability distributions over Θ , say $\Delta(\Theta)$. This is the set of possible beliefs the players may have regarding the true state, called first-level beliefs, and is analogous to my D_1 . Next, Mertens and Zamir consider the set of second-level beliefs, or beliefs about the other player's beliefs, which is just $\Delta(\Delta(\Theta))$ or, in the obvious notation, $\Delta^2(\Theta)$, analogous to my D_2 . More generally, their $\Delta^n(\Theta)$ is the analogue of my D_n . My construction is made slightly more complex by the fact that one has to go through two steps to go from D_n to D_{n+1} —construct the extension of S^{n-1} and then extend the set of decision procedures. Notice that one brings in a new object at this step, namely S_n . As we will see shortly, in part because of this, my construction also differs from theirs in the continuity of the operators involved.

As discussed in the introduction, we would like this sequence to converge to a fixed point, giving us a model in which each way the agent could resolve his subjective uncertainty is an option and the uncertainty about each option is represented. Unfortunately, though, this sequence does not converge to a fixed point in general. To see this intuitively, notice that at each level, the agent is only able to ask questions about the levels below him. Nowhere in the structure can the agent think about the entire sequence. That is, the agent can learn about S^n for any finite n , but he cannot learn about the limiting state set. On the other hand, this is an option at the limit.

To see this formally, let ω denote the first infinite ordinal—that is,

$$\omega = \{1, 2, \dots\}.$$

Let D_ω be the limit of D_n as $n \rightarrow \infty$, i.e.,

$$D_\omega = \bigcup_{n < \omega} D_n.$$

Similarly, let

$$S^\omega = \prod_{n < \omega} S_n$$

and let u_ω be the limit of $\{u_n | n = 1, 2, \dots\}$. It is easy to see that u_ω and S^ω do capture all the uncertainty about D_ω in a very natural sense. More precisely, consider any $d \in D_\omega$. Clearly, d is in D_n for some n , so the agent’s perception of d is defined by u_n and S^n . Since u_ω and S^ω simply extend u_n and S^n , they contain this information as well, so that uncertainty about every $d \in D_\omega$ is captured by these objects. However, D_ω does not contain all the agent’s options for learning about his uncertainty. That is, in general,

$$(1) \quad D_\omega \subset D(S^\omega).$$

(Recall that \subset denotes strict inclusion.) Hence D_ω is not the fixed point we sought to construct. Put loosely, the problem is that

$$(2) \quad \lim_{n \rightarrow \infty} D(S^n) \neq D\left(\lim_{n \rightarrow \infty} S^n\right).$$

The reason why (1) holds is quite simple. Consider the decision procedure, $d^* \in D(S^\omega)$, defined as follows. The procedure begins by performing some computation in $C(S_0)$. Then it performs some computation in $C(S^1) \setminus C(S_0)$, then some computation in $C(S^2) \setminus C(S^1)$, etc. If

$$(3) \quad C(S^n) \subset C(S^{n+1})$$

for all $n < \omega$,¹⁹ d^* is well-defined. Clearly, though, $d^* \notin D_n$ for any finite n and so $d^* \notin D_\omega$, even though $d^* \in D(S^\omega)$. It is not hard to show that (3) holds for all n whenever S_n has at least two elements for all n . Clearly, with a slightly more involved construction, essentially the same argument works under the weaker hypothesis that infinitely many of the S_n sets have at least two elements.

There are several approaches one might take to avoid this problem. An obvious possibility is to assume that all but finitely many of the S_n ’s are singletons. Recall, though, that the S_n ’s represent the agent’s subjective uncertainty about decision procedures or, more loosely, the limits of his rationality. Given this, such an assumption seems far too strong. For this reason, I assume that S_{n+1} is a singleton iff $D_{n+1} = D_n$. In other words, if $D_{n+1} = D_n$, there are no new decision procedures introduced at level $n + 1$, so there is no need to

¹⁹ To be more precise, $C(S^n)$ and $C(S^{n+1})$ are sets of functions on different domains and so one cannot be a subset of the other. However, it is natural to view $c: S^n \rightarrow L$ as the same as its trivial extension to $\hat{c}: S^{n+1} \rightarrow L$ (i.e., where \hat{c} does not use information in S_{n+1}). I adopt this convention throughout.

represent any new uncertainty. If, however, the set of options expands, the agent's uncertainty must increase to reflect his uncertainty about these options.

A second alternative would be to bound the number of steps of computation a procedure may perform. This assumption is clearly appropriate for some contexts. For example, if the agent must make a decision within a finite amount of time and the length of time it takes to perform a computation is bounded away from zero, then such a bound must exist. On the other hand, for many problems, there is no obvious maximum on the amount of time the agent can take. In such a case, it seems rather artificial to impose a bound—typically, there should be time for one more computation.²⁰ Aside from this objection, though, this approach to overturning (1) may not work. To see this, suppose L is infinite. This would be the case if, for example, L is the set of all finite length strings generated by some finite set of characters or the set of possible sentences in the English language. Suppose also that at least one S_n is infinite. Then, even if $\xi = 1$, there is a computation whose range is countable. Consider the decision procedure d^{**} which begins by performing any such computation. Clearly, d^{**} produces countably many histories of length 1. So if (3) holds for infinitely many n , we can have d^{**} perform any computation in $C(S^1) \setminus C(S_0)$ on the first of these histories, any computation in $C(S^2) \setminus C(S^1)$ on the second, etc. This procedure does not even need to perform more than two steps of computation!

Even aside from this, if $\xi \geq \omega$, we will not have $D_\omega = D(S^\omega)$. If ξ is not finite, then $C(S^\omega)$ contains computations using information in *every* S_n . Clearly, such a computation is not in any $C(S^n)$. Hence any decision procedure $d \in D(S^\omega)$ which uses such a computation is not in D_ω .

One alternative which does work is to restrict the computation sets to contain only computations whose ranges are finite. This can be viewed as an additional restriction on the complexity of the computations the agent can perform and so seems less objectionable than requiring L or the S_n 's to be finite. If we make this assumption, assume a finite bound on the number of steps of computation, and assume that ξ is finite, then $D_\omega = D(S^\omega)$. (This claim is a simple corollary of Theorems 1 and 2 in the next section.) On the other hand, as we will see in the next section, these assumptions are unnecessary if we are willing to use longer sequences to find our fixed point.

REMARK 2: A reader familiar with Mertens and Zamir may be quite surprised by the need to go beyond countable sequences. Why does $\{\Delta^n(\theta) | n < \omega\}$ converge to a fixed point when $\{D_n | n < \omega\}$ does not? As (2) should make clear, the answer is simply that $\Delta(\cdot)$ is continuous at $\prod_{n < \omega} \Delta^n(\theta)$ while $D(\cdot)$ is not continuous at S^ω . The reason for this, intuitively, is that the process of constructing decision procedures over S^ω is quite different from the process of constructing probability distributions over $\prod_{n < \omega} \Delta^n(\theta)$. Every probability distri-

²⁰ On the other hand, this criticism only applies to a uniform bound across procedures. There seems to be no contradiction between the view that one more computation is always possible and the view that each procedure must eventually stop computing. The comments below apply to uniform or nonuniform bounds, however.

bution over the infinite space implies a sequence of probability distributions over the smaller spaces. In this sense, no new probability distributions come in at the limit. On the other hand, a decision procedure in $D(S^\omega)$ is not equivalent to a sequence of decision procedures in $D(S^n)$, $n = 1, 2, \dots$.

To see the point more concretely, fix any measurable set, $B \subset \prod_{n < \omega} \Delta^n(\Theta)$. Suppose that for every n , the projection of B onto $\Delta^n(\Theta)$ is a proper subset of $\Delta^n(\Theta)$. There is no sense in which the probability of B is defined at some finite level N . However, its probability is determined by the infinite sequence; the probability of B is the limit of the probabilities of “truncations” of B to N dimensions.

Clearly, the procedure d^* constructed above is the analogue of B . Just as B could be defined as a limit of truncations, we can define d^* as a limit. To see this, fix some action a^* . For each n , let d_n^* be the procedure which follows d^* for n steps, but then picks action a^* at step $n + 1$. By the construction of d^* and d_n^* , we have $d_n^* \in D_n$ for all n . Furthermore, there is an obvious sense in which $\lim_{n \rightarrow \infty} d_n^* = d^*$. However, here is where the analogy to Mertens and Zamir breaks down: while probabilities must be continuous in this kind of limit, there is no obvious reason why the agent’s perceptions must be. In other words, one can show that if ξ is finite, then, with this alternative definition of the limit, formalized appropriately, $\lim_{n \rightarrow \infty} D_n = D(S^\omega)$. However, unless we assume that the agent’s perceptions of decision procedures are continuous in this topology, then the agent’s perceptions of the procedures in D_ω need not be completely specified by S^ω and u_ω . Since it is not obvious why this kind of continuity assumption would be a sensible restriction, it is fortunate that it is not needed if we consider longer sequences.

3. TRANSFINITE SEQUENCES

Summarizing, then, we can find a fixed point using a countable sequence, but only if we make stringent assumptions on the amount of subjective uncertainty facing the agent or on the ways he can resolve his uncertainty. If we do not wish to make these assumptions, we can still find a fixed point, but must work a bit harder. As noted in the introduction, it is the existence of the fixed point that I am concerned with here; the length of the sequence that finds a fixed point for us is of secondary importance. Analyzing the model the fixed point gives us is not made more complex by the fact that we used a transfinite argument to establish the existence of the fixed point.

To define longer sequences, we must continue the construction used above transfinitely. Thus I define P_n for $n \in \omega$ as above. P_ω is defined to be the limit of P_n —that is, it is the collection $(D_\omega, u^\omega, S^\omega)$. As noted above, $D_\omega \neq D(S^\omega)$, so let $D_{\omega+1} = D(S^\omega)$. Extending the set of options in this way again requires us to expand the state set and extend the utility function, giving us $u_{\omega+1}$ and $S^{\omega+1}$, where $S^{\omega+1} = S^\omega \times S_{\omega+1}$. Let $P_{\omega+1} = (D_{\omega+1}, u_{\omega+1}, S^{\omega+1})$. Clearly, we can continue this process to generate $P_{\omega+2}, P_{\omega+3}, \dots$, or $P_{\omega+n}$ for every $n \in \omega$. Again, we can take limits in the same way to define $P_{\omega+\omega}$, and then continue on to

$P_{\omega+\omega+1}$, etc. This generates a problem P_α for every ordinal number α . I denote this transfinite sequence $\{P_\alpha | \alpha \in ON\}$, where ON is the ordinal numbers, and will occasionally refer to it as the *hierarchy*.

Does this transfinite recursion find a fixed point for us? More precisely, if we consider some longer infinite sequence, does this sequence converge to a fixed point? To make the question precise, we require some definitions. A *limit ordinal* is an ordinal α for which there is no β such that $\alpha = \beta + 1$. In other words, a limit ordinal is an ordinal we must take limits to “get to,” such as ω or $\omega + \omega$. For a limit ordinal α , an α -*sequence* is a function with domain α . This is just the generalization of the usual notion of a sequence, which in this terminology is an ω -sequence. In short, “longer sequences” are just α -sequences where $\alpha > \omega$. The question, then, is whether there exists a limit ordinal α such that the α -sequence given by $\{P_\beta | \beta < \alpha\}$ converges to a fixed point.

Just as when $\alpha = \omega$, the key is the convergence of D_β . To see the point, suppose α is a limit ordinal. Then, just as when $\alpha = \omega$, there is a very natural sense in which S^α and u_α summarize the agent’s uncertainty about D_α . Since any $d \in D_\alpha$ is in D_β for some $\beta < \alpha$, the agent’s perception of d is summarized by S^β and u_β . Since S^α and u_α simply extend S^β and u_β , they contain this information as well. Hence the key question is whether $D_\alpha = D(S^\alpha)$.

For a decision procedure d , let H_d denote the set of histories generated by d . The reason that we cannot have $D_\omega = D(S^\omega)$ without the additional assumptions discussed in Section 2 is that otherwise, it is easy to find a decision procedure d for which H_d is at least countably infinite. Such a decision procedure can use computations which are in countably many different levels of the hierarchy and so the procedure is not in D_n for any finite n . Hence to find a fixed point, we need to be sure that for every d , H_d is not too large. “Too large” can be made precise using the notion of cofinality. Intuitively, the cofinality of a limit ordinal α is the length of the shortest sequence of ordinals converging to α from below. Formally, if α is a limit ordinal and $\{\beta_\nu | \nu < \theta\}$ is an increasing θ -sequence of ordinals with $\beta_\nu < \alpha$ for every ν , we say that the sequence is *cofinal* in α if its limit is α —that is, if

$$\bigcup_{\nu < \theta} \beta_\nu = \alpha.$$

The *cofinality* of α , denoted $\text{cf}(\alpha)$, is the smallest limit ordinal θ such that there is an increasing θ -sequence cofinal in α . For example, it is easy to see that $\text{cf}(\omega) = \omega$. This is true because the collection of ordinals smaller than ω is certainly an ω -sequence converging to ω and there is no smaller limit ordinal. More generally, $\omega \leq \text{cf}(\alpha) \leq \alpha$ for any limit ordinal α .²¹

To see why this concept is useful for characterizing the notion of “too large” a set of histories, suppose we can find a decision procedure in $D(S^\alpha)$ which generates more histories than $\text{cf}(\alpha)$. Just as in the countable case, this procedure can use computations at enough different levels of the hierarchy to

²¹ Devlin (1979, Chapter III) is a good source of information on cofinality.

guarantee that the procedure is not contained at any level below α . If so, $D_\alpha \neq D(S^\alpha)$ and thus the α -sequence $\{P_\beta | \beta < \alpha\}$ will not converge to a fixed point. Similarly, if $\xi > \text{cf}(\alpha)$, then there are computations in $C(S^\alpha)$ which are not in $C(S^\beta)$ for any $\beta < \alpha$. Hence a decision procedure $d \in D(S^\alpha)$ using such a computation is not in D_α . This is the reasoning behind the following theorem, proven in the Appendix.

THEOREM 1: *If α is a limit ordinal, then $D_\alpha = D(S^\alpha)$ iff for all $d \in D(S^\alpha)$, $\text{cf}(\alpha) > \max\{\xi, \#H_d\}$.*

As noted above, $\text{cf}(\omega) = \omega$. Hence Theorem 1 indicates why the countable sequence $\{P_n | n < \omega\}$ does not converge to a fixed point in general: except under very stringent assumptions, there are always decision procedures with $\#H_d \geq \omega$.

While there are generally decision procedures with $\#H_d \geq \omega$, one can construct an upper bound on $\#H_d$. Clearly, if a procedure begins by performing a computation, this computation cannot have more possible outcomes than the number of feasible outputs, $\#L$. Hence the set of histories of length one cannot be any larger than $\#L$. If a computation is performed after each of these histories, again, each cannot generate more than $\#L$ different outcomes, so the cardinality of the set of histories of length two cannot be larger than $(\#L)^2$. Continuing with this reasoning and recalling that every history of a decision procedure has finite length, we obtain a bound on the total number of histories of all lengths a procedure can generate.

THEOREM 2: *For any $\alpha \in ON$, for every $d \in D_\alpha$, $\#H_d \leq \sum_{n < \omega} (\#L)^n$.*²²

The upper bound of this theorem demonstrates one point made in the previous section: if we assume that L is finite and bound the number of steps of computation (uniformly or not), then no decision procedure can produce more than a finite number of histories. Hence, by Theorem 1, these restrictions imply that the countable sequence $\{P_n | n < \omega\}$ converges to a fixed point if ξ is finite. On the other hand, even without these assumptions, Theorems 1 and 2 together imply the existence of a fixed point. It is well known that for any β , there is a limit ordinal with cofinality larger than β .²³ Hence there is a limit ordinal α

²² The summation and exponentiation here both refer to cardinal arithmetic. The cardinal sum of two or more sets is simply the cardinality of the union of the sets (where we treat the sets as if they were disjoint by "labelling" each element according to the set from which it is drawn). The cardinal product of a family of sets is the cardinality of the product of the sets. Finally, cardinal exponentiation is defined as the cardinal product of a set with itself the requisite number of times. For finite numbers, these operations are the same as standard arithmetic.

²³ For example, let β' be any cardinal larger than β . It is easy to show that $\text{cf}(2^{\beta'}) > \beta' \geq \beta$ (see Devlin (1979, page 116)), so that $2^{\beta'}$ is such a limit ordinal.

such that

$$\text{cf}(\alpha) > \max \left\{ \xi, \sum_{n < \omega} (\#L)^n \right\}.$$

Thus Theorems 1 and 2 imply the following corollary.

COROLLARY: *There is a limit ordinal α such that $D_\alpha = D(S^\alpha)$.*

Thus if we consider long enough sequences, we can always find a fixed point. As discussed above, S^α necessarily describes all the agent's uncertainty regarding D_α . Hence when $D_\alpha = D(S^\alpha)$, we have a fixed point: D_α gives every option the agent might consider for resolving his uncertainty S^α and S^α describes all the agent's uncertainty about D_α . Hence $D_{\alpha+1} = D_\alpha$, so that $S_{\alpha+1}$ is a singleton, making $D_{\alpha+2} = D_\alpha$, etc. In short, for all $\gamma > \alpha$, $D_\gamma = D_\alpha$ and S_γ is a singleton. As argued in the introduction, given that this model completely characterizes the agent's perception of his problem, including his perception of how he can improve his perceptions, we must complete the model by assuming that the agent chooses the option he perceives to be the best one. While the agent's perception of what is best may be seriously flawed, why would he not choose what he thinks is best? In this sense, the assumption that the agent chooses the "optimal" decision procedure in D_α is without loss of generality. As noted above, the α itself is irrelevant—it is only part of the process the modeller, not the agent, goes through to construct the appropriate representation of the agent's perception. The main point is that since a fixed point exists, we know that there is a set of options D , a state set S , and a "perception" or utility function u such that the agent chooses what he perceives to be the best $d \in D$.

Of course, the Corollary does not imply that we can assume this for *any* D , S , and u . Instead, it says that such objects always exist for which this assumption is without loss of generality. It is also important to emphasize that this result is nonconstructive in the sense that it does not tell us how to construct the appropriate objects. Put differently, this model does not provide a mapping giving the agent's subjective perception of the world as a function of the objective data of the problem. Instead, it takes the agent's subjective perceptions as given and focuses on how we can work with them.

Still, the result is quite useful. First, it implies that the infinite regress problem which has seemed to prevent the development of a useful model of limited rationality is not necessarily a problem at all. Furthermore, it provides an approach to modeling limited rationality which may generate useful results. Of course, one must add further assumptions on the D , S , and u which are the fixed point in order to obtain such results. In Lipman (1989), I describe one way to extend the approach to games and give some results characterizing equilibria with limitedly rational players under alternative assumptions on the u 's. I also use this approach in Lipman (1990b) to explain the use of incomplete contracts.

REMARK 3: It is important to note that very few restrictions are needed to obtain the fixed point. Virtually no structure is needed on u or S . I have only required the agent's perceptions to be consistent in the sense that adding options does not change the way the agent perceives the existing options. Since the adding of options is done by the modeller, not the agent, this does not seem unreasonable. The importance of this point stems, in part, from the fact that u is supposed to represent the agent's view of his options prior to computation. This is quite different from the kind of preference we normally work with. Many common assumptions on preferences seem quite sensible when we have in mind preferences based on careful consideration. However, when we model the process behind such consideration, we ask for more primitive, "gut-level" preferences. It is not obvious what kind of structure these preferences should be assumed to have. The fact that these results use such weak requirements is thus very important.

REMARK 4: As emphasized above, the model generated by the fixed point has the property that it contains all the uncertainty the agent has about each of his options and contains all options for resolving the agent's uncertainty. One point of interest is that no procedure can resolve the uncertainty about itself. In this sense, the agent is never able to resolve all the uncertainty he faces.

REMARK 5: The requirement that no computation uses information in more than ξ of the state sets can be substantially relaxed. Suppose we simply assume that a computation c is feasible only if the collection of β 's such that c uses information in S_β is a set. Under this assumption, in general, there will not be any limit ordinal α such that $D_\alpha = D(S^\alpha)$. However, it will always be true that

$$\bigcup_{\alpha \in ON} D_\alpha = D\left(\prod_{\alpha \in ON} S_\alpha\right).$$

This is essentially the result of Theorem 2 of an earlier version of this paper, Lipman (1989). As discussed there, this version of the fixed point is more delicate since the objects involved may be proper classes. However, much the same interpretation can be given.

5. CONCLUSION

While many have argued the need to study bounded rationality, finding the appropriate model is quite difficult. Perhaps the most obvious objection to the recent approach of modeling limitations on rationality in the form of constraints or costs has been to the way these models have been closed: by the assumption that agents deal optimally with their limitations. I have shown here that this assumption is, in a particular sense, without loss of generality. Given that we have described the agent's perception of each of his options where his options include all feasible ways to refine his perceptions, the only appropriate assumption about choice is that the agent chooses what he perceives to be best.

While the agent's perception of what is best may be far from correct, it is surely what determines his choice. Hence in analyzing the fixed point, there is no inconsistency between supposing the agent is limitedly rational and assuming that he chooses optimally given his limitations.

There are many interesting potential applications of the model in game theory and economics. For example, it has long been suggested that the fact that real contracts are simpler and less complete than seems optimal might be due to bounded rationality.²⁴ In Lipman (1990b), incomplete contracts emerge endogenously in a model using this approach to limited rationality. As another example, the use of a model of limited rationality to (rigorously) explain deviations of experimental/econometric evidence from theories based on perfect rationality seems quite exciting.²⁵

An important and difficult problem is the formation of priors. As noted above, a complete model should go from the objective data of the problem to the individual's subjective perception of it and then to the individual's choice of action. This model only addresses the second step. The modeling involved in the first step raises some difficult problems. For example, the fact that the agent cannot be assumed to know all logical implications of his knowledge means that the form in which he knows the problem is crucial. This means much more than simply that games in normal form, extensive form, reduced form, etc., are not equivalent. It means that even permuting rows in a normal form game may lead to changes in behavior.²⁶ More generally, all sorts of framing effects can be expected to affect behavior—precisely as occurs in experimental settings.²⁷

*Graduate School of Industrial Administration, Carnegie Mellon University,
Schenley Park, Pittsburgh, PA 15213, U.S.A.*

Manuscript received February, 1989; final revision received June, 1990.

APPENDIX

FORMAL CONSTRUCTION OF THE TRANSFINITE HIERARCHY

I define the set of decision procedures by identifying procedures with certain trees. The idea is to let paths in the tree represent histories of computation and let the nodes of the tree represent the agent's choice of either an action or a computation given the history represented by the path to this node. A node at which the agent chooses an action will be a terminal node. If the agent chooses a computation at some node, then there is one arc from that node for each possible outcome of the computation.

²⁴ See Hart and Moore (1988).

²⁵ For recent work motivated by similar considerations, see Holt (1990).

²⁶ For an intuitively plausible example, consider an agent with 45,000 strategies. If the first strategy is dominated by the second one, this fact is likely to be noticed, so the first strategy is very unlikely to be used. On the other hand, if we move the second strategy to the 20,000th row and the first one is dominated by no other strategy, it seems more likely that the first strategy will be used.

²⁷ See, e.g., Tversky and Kahneman (1987). Interestingly, they note that the standard rationality axioms "are obeyed when their application is transparent and often violated in other situations," just as one would expect if making good choices is costly to the agent.

More formally, we have a set of states S and let $C = C(S)$. A *labelled tree in S* consists of the following collection. First, we have a set of nodes, N , which is a subset of $C \cup A$. Second, we have a partial order $<$ on N interpreted as precedence, where $<$ totally orders the class $\{n \in N \mid n < n'\}$ for all n' . If $n < n'$ and there is no n'' such that $n < n'' < n'$ then n' is a *successor of n* . Third, we have a function, γ , which labels ordered pairs (n, n') such that n' is a successor of n with elements of L . Finally, we have an initial node $n_0 \in N$ satisfying $n_0 < n$ for all $n \in N, n \neq n_0$. For any node n' , I refer to the set $\{n \in N \mid n < n'\} \cup \{n'\}$ as the *path to n'* .

I will say that a labelled tree $\mathcal{T} = (N, <, \gamma, n_0)$ is *acceptable* iff:

- (A1) If $n \in A$, then n has no successors.
- (A2) If $n \in C$, say $n = c$, then for every w' such that $c(s) = w'$ for some $s \in S$, there exists exactly one successor of n , say n' , such that $\gamma(n, n') = w'$. The node n has no other successors.
- (A3) For every $n \in N$, the path to n is finite.

Every acceptable labelled tree, or a.l.t., corresponds to a decision procedure. The initial node gives the first choice made by the agent. If an action is chosen, the procedure stops there. If a computation is chosen, there is a successor of the initial node for each possible output. Each of these possibilities corresponds to a possible history of length one. At each such history, the decision procedure makes another choice and so on. The path to each node must be finite because every history of computation has finite length. The fact that there is no upper bound on the length of a path means that there is no requirement that the procedure terminates.

I will say that a label w follows the path to n_k if there is a successor of n_k , say n_{k+1} , such that $\gamma(n_k, n_{k+1}) = w$. A *history* is an ordered pair consisting of a path to some node and a label which follows that path. (Hence if the path leads to a terminal node, it is not a history.) That is, the history gives the sequence of computations chosen and, through the labeling, the outcome of each. The *length* of a history is the number of nodes in the path, or the number of computations performed so far. It is convenient to also define the *ex ante history*, which I denote e , as the history preceding the choices we are modeling. This is the unique history of length zero. For a given a.l.t., \mathcal{T} , let $H(\mathcal{T})$ denote the set of histories in \mathcal{T} . The decision procedure corresponding to \mathcal{T} , or $d(\mathcal{T})$, is the function $d: H(\mathcal{T}) \rightarrow N(\mathcal{T})$ giving the node following the history. That is, for any path, d gives the next choice made. Thus d gives the same information as \mathcal{T} , but in a different form. For convenience, I write the set of histories $H(\mathcal{T})$ as H_d where $d = d(\mathcal{T})$. Finally, the set of decision procedures $D(S)$ is the set of $d(\mathcal{T})$ such that \mathcal{T} is an a.l.t. in S .

The construction of the hierarchy is by transfinite recursion. To provide the basis for the recursion, I first define P_0 . Let D_0 be the set of maps from $\{e\}$ to A . We also have a probability space $(S_0, \mathcal{F}_0, p_0)$, where S_0 is a compact metric space and \mathcal{F}_0 is the σ -algebra generated by the open sets. Finally, we have a utility function $u_0: D_0 \times S_0 \rightarrow \mathbf{R}$. P_0 is defined to be the collection $(D_0, u_0, (S_0, \mathcal{F}_0, p_0))$.

For any ordinal α , let

$$P_{\alpha+1} = (D_{\alpha+1}, u_{\alpha+1}, (S^{\alpha+1}, \mathcal{F}^{\alpha+1}, p_{\alpha+1})),$$

where $D_{\alpha+1} = D(S^\alpha)$. The utility function $u_{\alpha+1}: D_{\alpha+1} \times S^{\alpha+1} \rightarrow \mathbf{R}$ extends u_α in the sense that for any $d \in D_\alpha$ and $s \in S^{\alpha+1}$, $u_{\alpha+1}(d, s) = u_\alpha(d, s')$ where s' is the projection of s onto S^α . $S^{\alpha+1} = S^\alpha \times S_{\alpha+1}$ where $S^0 = S_0$ and $S_{\alpha+1}$ is a compact metric space. The σ -algebra, $\mathcal{F}^{\alpha+1}$, is the direct product of \mathcal{F}^α and $\mathcal{F}_{\alpha+1}$ where the latter is the σ -algebra of $S_{\alpha+1}$ generated by the open sets. The measure $p_{\alpha+1}$ is consistent with p_α in the sense that for any $B \in \mathcal{F}^\alpha$, $p_{\alpha+1}(B \times S_{\alpha+1}) = p_\alpha(B)$. (These assumptions are used only to guarantee that the limits are well-defined; they are not used in the proofs of Theorems 1 and 2.)

For any limit ordinal α , we define

$$P_\alpha = (D_\alpha, u_\alpha, (S^\alpha, \mathcal{F}^\alpha, p_\alpha)),$$

where

$$D_\alpha = \bigcup_{\beta < \alpha} D_\beta,$$

$$S^\alpha = \prod_{\beta < \alpha} S_\beta,$$

and \mathcal{F}^α is the limit of the direct products \mathcal{F}^β , $\beta < \alpha$. The utility function u_α is the uniquely defined limit of $\{u_\beta | \beta < \alpha\}$. (It is easy to show that the limit exists and is unique.) Finally, Kolmogorov's existence theorem guarantees the existence of a unique p_α consistent with p_β for all $\beta < \alpha$. (See Shirayev (1984, Theorem II.3.4 and the remarks following the proof.) Note that the consistency conditions required are implied by the consistency of $p_{\beta+1}$ with p_β . Hence P_α is well-defined for every ordinal α . I assume that $S_{\alpha+1}$ is a singleton iff $D_{\alpha+1} = D_\alpha$.

Some additional notation is needed for the proof of Theorem 1. The set of length k histories in H_d is denoted $H_d(k)$. Also, the rank of a computation c is the least ordinal α such that $c \in C(S^\alpha)$. Similarly, the rank of a decision procedure d is the least ordinal α such that $d \in D_\alpha$. Finally, for any d , $\mathcal{C}(d)$ denotes the set of computations used by d —that is, the set of c such that $d(h) = c$ for some $h \in H_d$.

PROOF OF THEOREM 1: (If.) For any decision procedure $d \in D(S^\alpha)$, let $R(\mathcal{C}(d))$ denote the set of ranks of the computations in $\mathcal{C}(d)$. It is easy to see that the rank of d is simply the largest ordinal in $R(\mathcal{C}(d))$. Also, $\#\mathcal{C}(d) \leq \#H_d$. Hence if for every $d \in D(S^\alpha)$, $\text{cf}(\alpha) > \#H_d$, then $\text{cf}(\alpha) > \#\mathcal{C}(d)$ for all $d \in D(S^\alpha)$. But then any $\#\mathcal{C}(d)$ -sequence of ordinals each strictly less than α has a limit strictly less than α . Hence if $\alpha \notin R(\mathcal{C}(d))$, then the $\#\mathcal{C}(d)$ -sequence given by the sequence of ordinals in $R(\mathcal{C}(d))$ has a limit strictly less than α . Thus if there is no computation in $C(S^\alpha)$ of rank α , then the rank of each $d \in D(S^\alpha)$ is strictly smaller than α , implying

$$D(S^\alpha) = \bigcup_{\beta < \alpha} D_\beta = D_\alpha.$$

I now show that $\text{cf}(\alpha) > \xi$ implies that there is no computation of rank α , completing this part of the proof. So suppose that $\text{cf}(\alpha) > \xi$ and consider any $c \in C(S^\alpha)$. Consider the sequence of γ 's such that c uses information in S_γ . By definition, the length of this sequence is less than ξ . Since $\xi < \text{cf}(\alpha)$, the limit of this sequence, say β , is necessarily strictly less than α . Hence $c \in C(S^\beta)$ for $\beta < \alpha$ and so $\text{rank}(c) < \alpha$.

(Only if.) First, suppose $\text{cf}(\alpha) \leq \xi$. Then consider any increasing ξ -sequence $\{\beta_\gamma | \gamma < \xi\}$ converging to α and consider a computation $c \in C(S^\alpha)$ which uses information in S_{β_γ} for all $\gamma < \xi$. Clearly, $\text{rank}(c) = \alpha$. But let d^* be any decision procedure in $D(S^\alpha)$ using such a computation. Clearly, the rank of d^* is α so that $d^* \notin D_\beta$ for any $\beta < \alpha$, implying $D_\alpha \subset D(S^\alpha)$.

So suppose $\text{cf}(\alpha) > \xi$ but there exists $d \in D(S^\alpha)$ such that $\#H_d \geq \text{cf}(\alpha)$. I will show that we again have $D_\alpha \subset D(S^\alpha)$. The following lemma is useful for this purpose.

LEMMA: For every limit ordinal α , $\text{cf}(\text{cf}(\alpha)) = \text{cf}(\alpha)$.

PROOF: Let $\theta = \text{cf}(\alpha)$ and suppose $\text{cf}(\theta) \neq \theta$. Recall that the cofinality of an ordinal is always weakly less than the ordinal itself, so we must have $\text{cf}(\theta) < \theta$. Let $\lambda = \text{cf}(\theta)$, let $\{\beta_\nu | \nu < \theta\}$ be an increasing θ -sequence converging to α , and let $\{\gamma_\kappa | \kappa < \lambda\}$ be an increasing λ -sequence converging to θ . Consider the sequence $\{\beta_{\gamma_\kappa} | \kappa < \lambda\}$. It is easy to see that this is an increasing λ -sequence which converges to α . But since $\lambda < \theta = \text{cf}(\alpha)$, this is impossible. Q.E.D.

To complete the proof, first, suppose $\text{cf}(\alpha) = \omega$. Then there is an increasing sequence of ordinals $\{\gamma_n | n \in \omega\}$ with $\gamma_n < \alpha$ for all n such that the limit of the sequence is α . If γ_n is not a limit ordinal, then there must be a computation of rank γ_n . To see this, recall that none of the state sets are singletons by assumption, so that there must be a computation which only uses information at S_{γ_n} . If γ_n is a limit ordinal, then there may be no computation with rank γ_n . However, if γ_n is a limit ordinal, there must be an ordinal β which is not a limit ordinal with $\gamma_{n-1} < \beta < \gamma_n$. In this case, we can certainly find a computation of rank β . Hence we can construct a procedure which chooses a computation of rank between γ_{n-1} and γ_n on each history of length n . It is easy to see that the rank of this procedure is also the limit of the sequence $(\gamma_n | n < \omega) = \alpha$. Hence $D_\alpha \subset D(S^\alpha)$. Therefore, we may as well assume $\text{cf}(\alpha) > \omega$.

Without loss of generality, we can assume that the d maximizing $\#H_d$ computes forever in every state. That is, for every $h \in H_d$, $d(h) \notin A$. This is without loss of generality because if d does not compute forever in every state, we can construct an alternative procedure which follows d until d terminates and then continues to compute. Obviously, this will (weakly) increase the cardinality of the set of histories generated. Since d does not terminate in any state, $\#H_d(n+1) \geq \#H_d(n)$. Let $N \subseteq \omega$ denote the set of n such that $\#H_d(n+1) > \#H_d(n)$. I now show that there must be a finite k such that $\#H_d(k) \geq \text{cf}(\alpha)$.

Case 1: Suppose N is finite. Then there exists a k such that $\#H_d(n+1) = \#H_d(n)$ for all $n \geq k$. The cardinality of H_d is the cardinal sum of the cardinalities of the $H_d(n)$ sets. Hence if $\#H_d(k)$ is finite, then $\#H_d = \omega$. But this implies $\text{cf}(\alpha) \leq \omega$, a contradiction. Hence $\#H_d(k) \geq \omega$. Since the sum of an infinite cardinal and any other cardinal is just the larger of the two, we see that $\#H_d = \#H_d(k)$. This implies $\#H_d(k) \geq \text{cf}(\alpha)$.

Case 2: Suppose N is countable. Then the sequence $\{\#H_d(n) \mid n \in N\}$ is an increasing ω -sequence converging to $\#H_d$, implying that $\#H_d$ is a limit ordinal with cofinality ω . Suppose that $\#H_d = \text{cf}(\alpha)$. Then $\text{cf}(\#H_d) = \text{cf}(\text{cf}(\alpha))$. By the Lemma, the right-hand side is $\text{cf}(\alpha)$, implying $\text{cf}(\alpha) = \omega$, a contradiction. Hence $\#H_d > \text{cf}(\alpha)$. But then we must be able to find an $k \in \omega$ such that $\#H_d(k) \geq \text{cf}(\alpha)$.

Since there exists a finite k such that $\#H_d(k) \geq \text{cf}(\alpha)$, we can construct a decision procedure, d' , which is identical to d on histories of length less than k . On each different k length history, d' chooses a computation of a different rank, where the sequence of the ranks has limit α . (This sequence can be constructed just as in the case where $\text{cf}(\alpha) = \omega$.) Since $\#H_d(k) \geq \text{cf}(\alpha)$, this is possible. Again, though, this implies that there is no $\beta < \alpha$ such that $d' \in D_\beta$, so $D_\alpha \subset D(S^\alpha)$. Q.E.D.

REFERENCES

- ABREU, D., AND A. RUBINSTEIN (1988): "The Structure of Nash Equilibrium in Repeated Games with Finite Automata," *Econometrica*, 56, 1259–1282.
- AUMANN, R. (1988): "Preliminary Notes on Irrationality in Game Theory," Working Paper.
- BEN-PORATH, E. (1989): "The Complexity of Computing a Best Response Automaton in Repeated Games with Mixed Strategies," Stanford University Working Paper.
- BERNHEIM, B. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028.
- BINMORE, K. (1987 and 1988): "Modeling Rational Players: Parts I and II," *Economics and Philosophy*, 3 and 4.
- BLUME, L., AND D. EASLEY (1989): "Wealth Dynamics and the Market Selection Hypothesis," Cornell University Working Paper.
- BRANDENBURGER, A., E. DEKEL, AND J. GEANAKOPOLOS (1989): "Correlated Equilibrium with Generalized Information Structures," Working Paper.
- CANNING, D. (1988): "Rationality and Game Theory when Players are Turing Machines," London School of Economics Working Paper.
- (1989): "Convergence to Equilibrium in a Sequence of Games with Learning," London School of Economics Working Paper.
- DEVLIN, K. (1979): *Fundamentals of Contemporary Set Theory*. New York: Springer-Verlag.
- FAGIN, R., AND J. HALPERN (1985): "Belief, Awareness, and Limited Reasoning: Preliminary Report," IBM Research Report 4657.
- FUDENBERG, D., AND D. KREPS (1988): "A Theory of Learning, Experimentation, and Equilibrium in Games," Working Paper.
- GEANAKOPOLOS, J. (1989): "Game Theory without Partitions, and Applications to Speculation and Consensus," Yale University Working Paper.
- GILBOA, I., AND E. ZEMEL (1989): "Nash and Correlated Equilibria: Some Complexity Considerations," *Games and Economic Behavior*, 1, 80–93.
- GREEN, E. (1982): "Internal Costs and Equilibrium: The Case of Repeated Prisoners' Dilemmas," University of Pittsburgh Working Paper.
- HACKING, I. (1967): "Slightly More Realistic Personal Probabilities," *Philosophy of Science*, 34, 311–325.
- HART, O., AND J. MOORE (1988): "Incomplete Contracts and Renegotiation," *Econometrica*, 56, 755–786.
- HINTIKKA, J. (1975): "Impossible Possible Worlds Vindicated," *Journal of Philosophical Logic*, 4, 475–484.
- HOLT, D. (1990): "Understanding Strategic Choice: The Statistical Analysis of Experimental Games," Carnegie Mellon University Working Paper.
- KALAI, E., AND W. STANFORD (1988): "Finite Rationality and Interpersonal Complexity in Repeated Games," *Econometrica*, 56, 397–410.
- KRIPKE, S. (1975): "Outline of a Theory of Truth," *Journal of Philosophy*, 72, 690–716.
- LIPMAN, B. (1989): "How to Decide How to Decide How to...: Limited Rationality in Decisions and Games," Carnegie Mellon University Working Paper.

- (1990a): "A Note on the Meaning of Common Knowledge of Rationality," Carnegie Mellon University Working Paper.
- (1990b): "Limited Rationality and Endogenously Incomplete Contracts," Carnegie Mellon University Working Paper.
- LIPMAN, B., AND S. SRIVASTAVA (1990a): "Computation as a Correlation Device," *Games and Economic Behavior*, 2, 154–172.
- (1990b): "Informational Requirements and Strategic Complexity in Repeated Games," *Games and Economic Behavior*, 2, 273–290.
- MERTENS, J.-F., AND S. ZAMIR (1985): "Formalization of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29.
- MONGIN, P., AND B. WALLISER (1988): "Infinite Regressions in the Optimizing Theory of Decision," in *Risk, Decision and Rationality*, ed. by B. Munier. Boston: D. Reidel Publishing, pp. 435–457.
- PEARCE, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.
- POPPER, K. (1985): "The Rationality Principle," written in 1967, reprinted as Chapter 29 of *Popper Selections*, ed. by D. Miller. Princeton: Princeton University Press, pp. 357–365.
- RUBINSTEIN, A. (1986): "Finite Automata Play the Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 39, 83–96.
- SAVAGE, L. (1967): "Difficulties in the Theory of Personal Probability," *Philosophy of Science*, 34, 305–310.
- SELTEN, R. (1978): "The Chain-Store Paradox," *Theory and Decision*, 9, 127–159.
- SHIRYAYEV, A. (1984): *Probability*. New York: Springer-Verlag.
- SIMON, H. (1955): "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69, 99–118.
- (1976): "From Substantive to Procedural Rationality," in *Method and Appraisal in Economics*, ed. by S. J. Latsis. Cambridge: Cambridge University Press, pp. 129–148.
- SPEAR, S. (1989): "Learning Rational Expectations Under Computability Constraints," *Econometrica*, 57, 889–910.
- TVERSKY, A., AND D. KAHNEMAN (1987): "Rational Choice and the Framing of Decisions," in *Rational Choice: The Contrast between Economics and Psychology*, ed. by R. Hogarth and M. Reder. Chicago: University of Chicago Press, 1987.
- VASSILAKIS, S. (1989): "Economic Data Types," University of Pittsburgh Working Paper.
- WINTER, S. (1975): "Optimization and Evolution in the Theory of the Firm," in *Adaptive Economic Models*, ed. by R. H. Day and T. Groves. New York: Academic Press, pp. 73–118.