Consequentialism, Structural Rationality, and Game Theory

Peter J. Hammond

Permanent office address (to be printed with paper):

Department of Economics, Stanford University, CA 94305-6072, U.S.A.

Temporary office address, October 1993–July 1994 (for correspondence and proofs):

Institute for World Economics, University of Kiel, Dürsternbrooker Weg 120,

Postfach 4309, D-24105 Kiel 1, GERMANY.

November 1993 revision of the paper presented to the International Economic Association conference on "Rationality in Economics" at the International Centre for Economics Research, Villa Gualino, Turin from October 16th to 18th.

Abstract: Previous work on consequentialism (especially in *Theory and Decision*, 1988, pp. 25–78) has provided some justification for regarding an agent's behaviour as "structurally rational" if and only if there are subjective probabilities, and expected utility is maximized. The key axiom is that rational behaviour should be explicable as the choice of good consequences. This and other axioms will be re-assessed critically, together with their logical implications. Their applicability to behaviour in n-person games will also be discussed. The paper concludes with some discussion of modelling bounded rationality.

I. THREE CONSEQUENTIALIST AXIOMS

In the little space and brief time allowed to me, I shall try to impart some of the key ideas of the "consequentialist" approach to rational behaviour. At the same time, I shall try to assess its significance and to explain its limitations.

Consequentialism relies on the presumption that behaviour is rational if and only if it is explicable by its consequences. More specifically, the set of consequences which can result from behaviour should depend only on the set of feasible consequences. And, as the key assumption, this should be true for an (almost) unrestricted domain of finite decision trees whose terminal nodes have specified consequences.

More formally, let Y denote a fixed domain of possible consequences. Let $\mathcal{T}(Y)$ denote the domain of finite decision trees with consequences in Y.[1] Each member T of $\mathcal{T}(Y)$ takes the form of a list:

$$\langle N, N^*, X, n_0, N_{+1}(\cdot), \gamma(\cdot)\rangle$$

whose six components are:

(i) the finite set of nodes N;

(ii) the subset $N^* \subset N$ of decision nodes;

(iii) the complementary set $X = N \setminus N^*$ of terminal nodes;

(iv) the initial node $n_0 \in N$;

(v) the correspondence $N_{+1}(\cdot) : N \twoheadrightarrow N$ determining what set $N_{+1}(n)$ of nodes immediately succeeds each node $n \in N$, which satisfies obvious properties ensuring that N has a tree structure with X as the set of terminal nodes because $N_{+1}(x)$ is empty for all $x \in X$;

(vi) the mapping $\gamma : X \to Y$ from terminal nodes to associated consequences.

At any decision node $n \in N^*$, each immediate successor $n' \in N_{+1}(n)$ of node n corresponds to a particular move from n to n' which the agent can make. Thus behaviour at $n \in N^*$ can be described by a non-empty "chosen" subset $\beta(T, n)$ of the set $N_{+1}(n)$ of all immediate successors of node n.[2] The first assumption is:

Axiom 1 (unrestricted domain). *There is a* **behaviour correspondence** β *whose values satisfy* $\emptyset \neq \beta(T, n) \subset N_{+1}(n)$ *at every decision node* $n \in N^*$ *of every decision tree* $T \in \mathcal{T}(Y)$.

[1] Of course, it is restrictive in general to consider only finite decision trees. But not when discussing the implications of consequentialism, as is done here. The issue with infinite decision trees is whether behaviour can be well defined in a way that naturally extends consequentialist behaviour in finite trees. Obviously, this will require a technical analysis of compactness and continuity conditions.

[2] This formulation does exclude stochastic behaviour, according to which $\beta(T, n)$ is a probability distribution over $N_{+1}(n)$. Under the consequentialist axioms set out below, it turns out that a very strong transitivity axiom — or what Luce (1958, 1959) describes as a "choice axiom" — must be satisfied. In fact, there must exist a preference ordering over consequences which is maximized by all the consequences that can occur with positive probability given the agent's stochastic behaviour. Thus randomization occurs only over consequences in the highest indifference class of feasible consequences available to the agent. Furthermore, each consequence must have a positive real number attached indicating the relative likelihood of that consequence occurring, in case it belongs to the highest indifference class that the agent can reach.

It should be noted that Axiom 1 is not entirely innocuous. For example, suppose that the consequences were extended to include a list of what would result from the different actions available to the agent at each moment of time. Then important facts about the structure of the decision tree could be inferred from these extended consequences. So the domain of decision trees on which behaviour is defined would be limited accordingly. Thus, Axiom 1 makes sense only when the consequences themselves are do not depend on the structure of the tree.[3] Of course, such independence has been the standard assumption in classical decision theory.

An important and natural property of the behaviour correspondence concerns *subtrees*, which take the form

$$T(n) = \langle N(n), N^*(n), X(n), n, N_{+1}(\cdot), \gamma(\cdot) \rangle$$

for some initial node n which is any node of T. Here $N(n)$ consists of all nodes in N which succeed n, including n itself. Of course n becomes the initial node of $T(n)$. Also $N^*(n) = N^* \cap N(n), X(n) = X \cap N(n)$, while the mappings $N_{+1}(\cdot)$ and $\gamma(\cdot)$ apply to the restricted domains $N(n)$ and $X(n)$ respectively. The relevant property I shall assume is:

Axiom 2 (dynamic consistency). *In every subtree $T(n)$ of each decision tree $T \in \mathcal{T}(Y)$, and at every decision node n^* of $T(n)$, one has $\beta(T(n), n^*) = \beta(T, n^*)$.*

The justification for this second assumption is that both $\beta(T(n), n^*)$ and $\beta(T, n^*)$ describe behaviour at n^*; whether n^* is regarded as a decision node of the full tree T or of the subtree $T(n)$ should be irrelevant. It turns out that, for an agent whose tastes are changing endogenously, even naïve behaviour in decision trees is dynamically consistent in this sense; for such an agent the inconsistency will be between plans and actual behaviour.[4]

The most important axiom of consequentialism is the third, which will be stated next. It involves considering, for any tree $T \in \mathcal{T}(Y)$, the *feasible set* $F(T)$ of all possible consequences

[3] Actually, Prasanta Pattanaik and Robert Sugden tried to convince me of essentially this point many years ago. Jean-Michel Grandmont and Bertrand Munier were finally more successful during the conference. I am grateful to all four, and would like to apologize to the first two.

[4] See Hammond (1976) for more discussion of naïve behaviour in decision trees, especially the "potential addict" example. Also, Amos Tversky has suggested that Axiom 2 be called "subtree consistency." I am sympathetic, but have not followed this suggestion for two reasons: (i) earlier papers used the term "dynamic consistency," and I would like to be consistent myself; (ii) subtrees do represent dynamic choice possibilities when the tree models decisions which will be made in real time.

which can result from the agent's decisions in T. This set, and the feasible sets $F(T(n))$ for each subtree $T(n)$ in turn, can be constructed by backward recursion. To do so, let $F(T(x))$ be $\{\gamma(x)\}$ when $x \in X$ is any terminal node, and then define

$$F(T(n)) = \bigcup_{n' \in N_{+1}(n)} F(T(n'))$$

for all decision trees $T(n)$ starting at the other nodes of T. Of course, $F(T) = F(T(n_0))$.

The following statement of the third axiom also involves the set $\Phi_\beta(T)$ of all possible consequences which can result from the agent's decisions when they lie in the behaviour set $\beta(T, n)$ at each decision node $n \in N$. This set can also be constructed by backward recursion through the subtrees of T. The construction starts with $\Phi_\beta(T(x)) = \{\gamma(x)\}$ for the trivial subtree $T(x)$ starting at any terminal node $x \in X$, and then proceeds by defining

$$\Phi_\beta(T(n)) = \bigcup_{n' \in \beta(T,n)} \Phi_\beta(T(n'))$$

for all other nodes $n \in N \setminus X$, until it arrives at $\Phi_\beta(T) = \Phi_\beta(T(n_0))$.

After these preliminary constructions, I can state:

AXIOM 3 (CONSEQUENTIALIST BEHAVIOUR). *On the domain of all non-empty finite subsets of Y, there exists a* **revealed consequence choice function** C_β *which is non-empty valued and satisfies the property that* $\Phi_\beta(T) = C_\beta(F(T))$ *for all* $T \in \mathcal{T}(Y)$.

This means that the set of possible consequences of behaviour should depend only the feasible set of consequences, so that behaviour can be interpreted as the pursuit of chosen consequences. The assumption that $C_\beta(F)$ is non-empty for every F in its domain loses no generality. For, given any non-empty finite set of consequences $F \subset Y$, one can construct a tree whose only decision node is the initial node n_0, while each consequence $y \in F$ has a corresponding terminal node x_y for which $\gamma(x_y) = y$. Then, of course, T is a finite decision tree in $\mathcal{T}(Y)$ for which $F(T) = F$ and $\Phi_\beta(T) = \{\, y \in F \mid x_y \in \beta(T, n_0) \,\} \neq \emptyset$.

It is the three "consequentialist" axioms presented above which have such strong implications,[5] and so play such a crucial role in all that follows. I want to emphasize very strongly that none of the other usual rationality hypotheses are being invoked. There is no

[5] For obvious reasons, one does not speak of (logical) "consequences."

presumption that a preference ordering exists. When we come to risk and uncertainty, there will be no assumption of independence or of Savage's sure thing principle. These standard rationality hypotheses will turn out to be implications of the three axioms set out above, however. That is what is so striking about them.

II. A REVEALED PREFERENCE ORDERING

In simple decision trees without risk or uncertainty, behaviour satisfying Axioms 1–3 must reveal a preference ordering[6] over the consequence domain. Specifically, given the revealed consequence choice function C_β defined on the domain consisting of all non-empty finite subsets of Y, there is a unique corresponding weak preference relation R_β defined by

$$y \; R_\beta \; y' \iff y \in C_\beta(\{ y, y' \})$$

This is the preference revealed by behaviour in any decision tree T whose feasible set $F(T)$ consists of a pair of consequences $\{ y, y' \}$ — or of a single consequence in case $y = y'$. The relation R_β is complete because $C_\beta(F)$ is non-empty whenever $F = \{ y, y' \}$, so either $y \; R_\beta \; y'$, or $y' \; R_\beta \; y$, or both.

The above definition of R_β already ensures that

$$C_\beta(F) = \{ y \in F \mid y' \in F \Longrightarrow y \; R_\beta \; y' \} \tag{1}$$

whenever F is a singleton or a pair set. Then $C_\beta(F)$ consists of precisely those members of F which maximize the preference relation R_β. In fact it is not too difficult to prove that (1) also holds for every finite set $F \subset Y$, no matter how large, and also that R_β is a preference ordering. Thus Axioms 1–3 imply:

Ordinal choice. *The revealed choice function C_β defined on the domain of all finite non-empty subsets of Y corresponds to a preference ordering R_β defined on the whole of Y.*

To save space, however, the proof is not given here — see Hammond (1988, Section 5). Instead, I just state:

[6] Following the terminology of social choice theory, a "preference ordering" means a complete and transitive binary weak preference relation.

THEOREM 1. *Suppose that Y is any consequence domain. Then:*

(A) Any behaviour correspondence β satisfying Axioms 1–3 reveals an ordinal choice function C_β.

(B) Conversely, given any preference ordering R on Y, there exists an associated behaviour correspondence β satisfying Axioms 1–3 whose revealed preference ordering R_β is equal to R.

Thus, one contentious axiom of structural rationality has become an implication of other axioms that may be harder to question. Above all, those who argue that a preference ordering is not necessary for structural rationality owe us an explanation of how they expect agents to depart from the consequentialist axioms, and how they should behave in decision trees.[7] Perhaps the structure of the tree, as well as the set of feasible consequences, should be allowed to affect the set of chosen consequences. But would this really be rational?

Note that Theorem 1 includes the converse of its first part. This is important because it confirms that ordinality is a *complete* characterization of behaviour satisfying the three consequentialist axioms. The converse is easily proved by ordinal dynamic programming arguments, as in Section 8 of Hammond (1988).

III. UNCERTAINTY AND THE SURE THING PRINCIPLE

Formally, let Y denote the consequence domain, and E a fixed non-empty finite set of possible states of the world. A *decision tree with uncertainty* is then defined as a list:

$$T = \langle N, N^*, N^1, X, n_0, N_{+1}(\cdot), \gamma(\cdot), S(\cdot)\rangle$$

Compared to the definition in Section I, the new features are: (i) N^1, the set of *natural nodes* at which nature's move reveals information; and (ii) the *event correspondence* $S : N \twoheadrightarrow E$ specifying what non-empty set $S(n) \subset E$ of states of the world is possible after reaching node $n \in N$. Note that N is now partitioned into the three disjoint sets N^*, N^1 and X. At any natural node $n \in N^1$, nature's move partitions the set $S(n)$ into the collection pairwise

[7] Some critics of consequentialism, notably Machina (1989) and McClennen (1990), have offered such explanations, especially for behaviour that maximizes a preference ordering and yet violates the independence condition which, as will be discussed in Section V, is one other implication of consequentialism.

disjoint subsets $S(n')$ ($n' \in N_{+1}(n)$). At any decision node $n \in N^*$, however, the agent's move cannot refine what information there is already about the state of the world, and so it is required that $S(n') = S(n)$ whenever $n' \in N_{+1}(n)$. Finally, the consequence mapping γ takes its value $\gamma(x)$ in the relevant Cartesian product set $Y^{S(x)} := \prod_{s \in S(x)} Y_s$ at each terminal node $x \in X$, where each Y_s ($s \in E$) is just a copy of the consequence domain Y.[8] Thus terminal nodes are associated with profiles $\gamma(x) = \langle y_s \rangle_{s \in S(x)} \in Y^{S(x)}$ of uncertain contingent consequences rather than with certain consequences $y \in Y$.

An obvious modification of the unrestricted domain Axiom 1 is to require that $\beta(T, n)$ be defined throughout all finite decision trees with uncertainty. Axioms 2–3 (dynamic consistency and consequentialist behaviour) can be applied virtually without change. The revealed consequence choice function, however, should really be replaced by a collection C^S_β of such functions, one for each non-empty set $S \subset E$. The reason is that feasible sets and revealed consequence choice sets will be subsets of different Cartesian product sets Y^S ($\emptyset \neq S \subset E$) rather than just of Y.

The backward recursion construction of Section I must also be adapted to treat natural nodes $n \in N^1$. At such nodes, because choices at different later decision nodes can be made independently, it is natural to construct the following Cartesian products:

$$F(T(n)) = \prod_{n' \in N_{+1}(n)} F(T(n')) \subset \prod_{n' \in N_{+1}(n)} Y^{S(n')} = Y^{S(n)};$$
$$\Phi_\beta(T(n)) = \prod_{n' \in N_{+1}(n)} \Phi_\beta(T(n')) \subset \prod_{n' \in N_{+1}(n)} F(T(n')) = F(T(n)).$$

Given any non-empty $S \subset E$, consider the restricted domain $\mathcal{T}^S(Y)$ of finite decision trees with no natural nodes, and with $S(x) = S$ at every terminal node. This restricted domain is effectively equivalent to $\mathcal{T}(Y^S)$, the domain of finite decision trees with consequences in Y^S. Applying Theorem 1 to this domain establishes the existence of a *conditional preference ordering* R^S_β on Y^S that is revealed by consequentialist behaviour in such trees.

Consequentialism also establishes an important relationship between the different revealed preference orderings R^S_β ($\emptyset \neq S \subset E$). Indeed, let S_1 and S_2 be two non-empty disjoint subsets of E, and let $S = S_1 \cup S_2$. Let a^{S_1}, b^{S_1} be two contingent consequences

[8] An important generalization allows the consequence domain Y_s to depend on the state s. As one might expect, it then becomes harder to ensure that behaviour reveals subjective probabilities. However, both the sure thing principle and independence do still follow from the consequentialist axioms.

in Y^{S_1}, and c^{S_2} a contingent consequence in Y^{S_2}. Consider the decision tree illustrated in Figure 1, with an initial natural node n_0 whose two successors are: (i) decision node n_1, at which $S(n_1) = S_1$; and (ii) terminal node x_c, at which $S(x_c) = S_2$ and $\gamma(x_c) = c^{S_2}$. Suppose that n_1 offers the choice of going to either of the terminal nodes x_a and x_b, at which $S(x_a) = S(x_b) = S_1$ while $\gamma(x_a) = a^{S_1}$ and $\gamma(x_b) = b^{S_1}$.

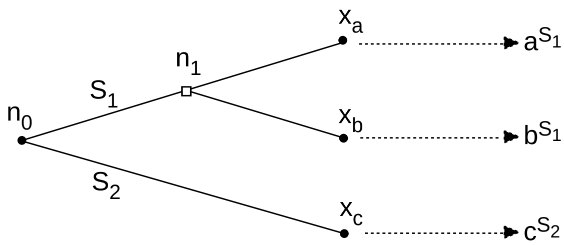

Figure 1

In this tree the non-trivial feasible sets are $F(T(n_1)) = \{ a^{S_1}, b^{S_1} \}$ and

$$F(T) = F(T(n_0)) = F(T(n_1)) \times F(T(x_c)) = \{ a^{S_1}, b^{S_1} \} \times \{c^{S_2}\} = \{ (a^{S_1}, c^{S_2}), (b^{S_1}, c^{S_2}) \}.$$

Using the definition of the revealed preference relations R^S_β on Y^S and $R^{S_1}_\beta$ on Y^{S_1}, as well as the dynamic consistency assumption that $\beta(T, n_1) = \beta(T(n_1), n_1)$, leads to the following chain of logical equivalences:

$$\begin{aligned} a^{S_1} \ R^{S_1}_\beta \ b^{S_1} &\iff a^{S_1} \in C^{S_1}_\beta(\{ a^{S_1}, b^{S_1} \}) = C^{S_1}_\beta(F(T(n_1))) \iff a^{S_1} \in \Phi^{S_1}_\beta(T(n_1)) \\ &\iff x_a \in \beta(T(n_1), n_1) \iff x_a \in \beta(T, n_1) \iff (a^{S_1}, c^{S_2}) \in \Phi^S_\beta(T) \\ &\iff (a^{S_1}, c^{S_2}) \in C^S_\beta(\{ (a^{S_1}, c^{S_2}), (b^{S_1}, c^{S_2}) \}) \iff (a^{S_1}, c^{S_2}) \ R^S_\beta \ (b^{S_1}, c^{S_2}). \end{aligned}$$

Thus, to summarize, consequentialism implies:

Savage's sure thing principle. *Whenever $S \subset E$ is partitioned into non-empty sets $S_1 \cup S_2$, and $a^{S_1}, b^{S_1} \in Y^{S_1}$ while $c^{S_2} \in Y^{S_2}$, then*

$$a^{S_1} \ R^{S_1}_\beta \ b^{S_1} \iff (a^{S_1}, c^{S_2}) \ R^S_\beta \ (b^{S_1}, c^{S_2}).$$

The main result of this section is:

THEOREM 2. *Suppose that Y is any consequence domain and that E is any finite non-empty set of uncertain states of the world. Then:*

(A) For any behaviour correspondence β satisfying Axioms 1–3, the hypothesized revealed choice functions C_β^S, defined for all non-empty $S \subset E$ on the domain of finite non-empty subsets of the appropriate Cartesian product set Y^S, must correspond to preference orderings R_β^S which together satisfy the sure-thing principle.

(B) Conversely, given any family of orderings R^S (for all non-empty $S \subset E$) defined on each product set Y^S that together satisfy the sure thing principle, there is an associated behaviour correspondence β satisfying Axioms 1–3 whose revealed preference ordering R_β is equal to R.

Part (A) has already been explained, though not proved formally. Like part (B) of Theorem 1, the converse can be proved by ordinal dynamic programming arguments. See Hammond (1988) for complete proofs.

IV. UNORDERED EVENTS

One of the principal axioms used by Savage (1954) actually goes back to earlier work by Keynes, Ramsey, de Finetti and others. This is the *ordering of events*. It requires the existence of a weak ordering $\geq$ on subsets of E, with the idea that the associated strong ordering $S_1 > S_2$ should mean that S_1 is more likely, or more probable, than S_2. More specifically, suppose that $\bar{y}$ and $\underline{y}$ are respectively "good" and "bad" consequences in Y, and that an agent is given the choice between the following two alternative contingent consequences:

(i) $(\bar{y}\,1^{S_1}, \underline{y}\,1^{S_2})$, representing $\bar{y}$ for sure if S_1 occurs, and $\underline{y}$ for sure if S_2 occurs;

(ii) $(\underline{y}\,1^{S_1}, \bar{y}\,1^{S_2})$, representing $\underline{y}$ for sure if S_1 occurs, and $\bar{y}$ for sure if S_2 occurs.

Then $S_1 > S_2$ should be equivalent to $(\bar{y}\,1^{S_1}, \underline{y}\,1^{S_2})$ being preferred to $(\underline{y}\,1^{S_1}, \bar{y}\,1^{S_2})$ for all pairs $\bar{y}$, $\underline{y}$ with $\bar{y}$ preferred to $\underline{y}$, because it is always preferable to "win" a good consequence contingent upon a more likely event.

An important corollary of Theorem 2 is that this ordering of events property is not an implication of the consequentialist axioms, because those axioms imply only the existence of conditional preference orderings satisfying the sure-thing principle, and no more. To establish this, it is enough to exhibit a family of contingent preference orderings satisfying the sure-thing principle but not the ordering of events.

Indeed, suppose that $Y = \{ y_1, y_2, y_3 \}$. Also let $E = \{ s_1, s_2 \}$. Then let $v : Y \to \Re$ be the real valued utility function defined by $v(y_i) := i - 2$ $(i = 1, 2, 3)$ whose values will be assumed to represent the preference ordering R^s_β on Y for each single state $s \in E$. Suppose too that R^E on Y^E is represented by the utility function $\phi_{s_1}(v(y_{s_1})) + \phi_{s_2}(v(y_{s_2}))$, where

$$\phi_{s_1}(v) = \begin{cases} v & \text{if } v \geq 0 \\ 2\,v & \text{if } v < 0 \end{cases} \quad \text{and} \quad \phi_{s_2}(v) = \begin{cases} 2\,v & \text{if } v \geq 0 \\ v & \text{if } v < 0 \end{cases}$$

Because this utility function is additive, it is obvious that the sure thing principle is satisfied.

In this case the total utility of the pair (y_1, y_2) (meaning y_1 if state s_1, and y_2 if state s_2) is $\phi_{s_1}(-1) + \phi_{s_2}(0) = -2$, while that of the pair (y_2, y_1) is $\phi_{s_1}(0) + \phi_{s_2}(-1) = -1$. Hence (y_2, y_1) is preferred to (y_1, y_2), thus suggesting that s_1 is more likely than s_2. On the other hand, the total utility of (y_2, y_3) is $\phi_{s_1}(0) + \phi_{s_2}(1) = 2$, while that of the pair (y_3, y_2) is $\phi_{s_1}(1) + \phi_{s_2}(0) = 1$. Hence (y_2, y_3) is preferred to (y_3, y_2), thus suggesting that s_2 is more likely than s_1. There is no well-defined ordering of the two events $\{s_1\}$ and $\{s_2\}$.

V. RISK AND UNCERTAINTY COMBINED

Without objective probabilities, consequentialism in decision trees with uncertainty does not imply the ordering of events, and so *a fortiori* does not imply that subjective probabilities are revealed by the agent's behaviour. With objective probabilities, however, and given some additional structural assumptions, it will be true that consequentialist behaviour must maximize expected utility, where expectations are represented by appropriate combinations of objective and subjective probabilities.

The first step is to extend once again the definition of a decision tree to accommodate a set N^0 of *chance nodes* at each of which there is a collection $\pi(n'|n)$ $(n' \in N_{+1}(n))$ of *transition probabilities*. Here $\pi(n'|n)$ is the probability of reaching n' conditional on having reached n already. Of course $\pi(n'|n) \geq 0$ for all $n' \in N_{+1}(n)$ and $\sum_{n' \in N_{+1}(n)} \pi(n'|n) = 1$ for all $n \in N^0$. For reasons that will emerge in due course, in fact I shall consider only probabilities satisfying $\pi(n'|n) > 0$ everywhere. One could argue that parts of the decision tree that are reached with only zero probability should be pruned off anyway, though in game theory this leads to severe difficulties in subgames.

So a finite decision tree becomes a list

$$T = \langle N, N^*, N^0, N^1, X, n_0, N_{+1}(\cdot), \pi(\cdot|\cdot), S(\cdot), \gamma(\cdot)\rangle$$

Now N is partitioned into the four sets N^*, N^0, N^1 and X (though any of the first three sets could be empty). It is also appropriate to re-define the range of the consequence mapping γ on X to consist of (simple) probability distributions on the appropriate set $Y^{S(x)}$, rather than single members of $Y^{S(x)}$. Thus probabilities $\gamma(y^{S(x)}|x)$ must be defined for each $x \in X$ and $y^{S(x)} \in Y^{S(x)}$. Moreover, it must be true that $\gamma(y^{S(x)}|x) > 0$ only for those $y^{S(x)} \in Y^{S(x)}$ in the finite *support* of $\gamma(\cdot|x)$.

It is also necessary to modify once again the recursive construction of the two sets $F(T(n))$ and $\Phi_\beta(T(n))$ for the subtree $T(n)$ starting at each node $n \in N$ so as to allow for chance nodes and probabilistic contingent consequences. In fact, it is natural to take

$$F(T(n)) = \sum\nolimits_{n' \in \beta(T,n)} \pi(n'|n)\, F(T(n')); \quad \Phi_\beta(T(n)) = \sum\nolimits_{n' \in \beta(T,n)} \pi(n'|n)\, \Phi_\beta(T(n'))$$

at any chance node $n \in N^0$ where the transition probabilities are $\pi(n'|n)$ $(n' \in N_{+1}(n))$. Thus each set consists of all possible appropriately probability weighted sums, as n' ranges over $N_{+1}(n)$, of the probability distributions belonging to the succeeding sets $F(T(n'))$ or $\Phi_\beta(T(n'))$. Finally, at natural nodes the appropriate product sets are no longer Cartesian products of sets, but rather sets consisting of all the independent joint probability distributions that can be created by multiplying members of the appropriate sets of probabilities in every way possible.

The consequentialist axioms 1–3 of Section I, as modified in Section III, can now be applied to this new extended domain of decision trees. They imply that there exists a conditional revealed preference ordering R^S_β for every non-empty $S \subset E$. The domain of R^S_β is no longer Y^S, however, but the set $\Delta(Y^S)$ of all "simple" probability distributions over Y^S — i.e., all distributions having finite support. Moreover, the *sure thing principle* should now be re-stated so that, whenever S_1, S_2 are non-empty disjoint subsets of E, and $\lambda, \mu \in \Delta(Y^{S_1})$ while $\nu \in \Delta(Y^{S_2})$, then

$$(\lambda \times \nu)\ R^S_\beta\ (\mu \times \nu) \iff \lambda\ R^{S_1}_\beta\ \mu$$

where $S = S_1 \cup S_2$ and $\times$ denotes the usual product of the probability distributions.

This does not exhaust the implications of consequentialism, however. In this new framework with some objective probabilities, it is also possible to adapt the argument used in Section III to derive the sure thing principle and establish the following:

Independence condition. *Suppose that* $\lambda, \mu, \nu \in Y^S$ *for some non-empty* $S \subset E$*, and that* $0 < \alpha \leq 1$*. Then* $\alpha\,\lambda + (1-\alpha)\,\nu \; R^S_\beta \; \alpha\,\mu + (1-\alpha)\,\nu \iff \lambda \; R^S_\beta \; \mu$*.*

In fact, were zero probabilities allowed in decision trees, one could prove the same result even when $\alpha = 0$. The implication would be that $\nu \; R^S_\beta \; \nu \iff \lambda \; R^S_\beta \; \mu$ for all $\lambda, \mu, \nu \in Y^S$. Since the left hand side of this equivalence is always true, so is the right hand side, and so there must be universal indifference over Y^S, for every non-empty $S \subset E$! To avoid this absurdity with ordinary probabilities requires either restricting the domain to exclude zero probability moves at chance nodes, as has been done in this paper, or else not imposing the dynamic consistency condition $\beta(T, n') = \beta(T(n), n')$ in decision subtrees $T(n)$ which can only be reached with zero probability. Neither escape is really appropriate in game theory, however, which has led me to consider decision trees embodying non-Archimedean probabilities in Hammond (1993a, 1994). Here, though, the exclusion of zero probabilities can perhaps be forgiven as a simplifying assumption.

By now the following result, whose proof is once again to be found in Hammond (1988), should come as no surprise:

Theorem 3. *Suppose that* Y *is a given consequence domain, that* E *is a given non-empty finite set of possible states of the world. Then:*

(A) If behaviour satisfies axioms 1–3 for the domain of finite decision trees with chance and/or natural nodes in which there are no zero probability moves at any chance node, then there must exist a family of revealed conditional preference orderings R^S_β *on* $\Delta(Y^S)$*, one for each non-empty* $S \subset E$*, satisfying both independence and the sure thing principle.*

(B) Conversely, given any family R^S *of orderings on the sets* $\Delta(Y^S)$ *(* $\emptyset \neq S \subset E$*) that satisfies independence and the sure thing principle, there is an associated behaviour correspondence* β *satisfying axioms 1–3 in all decision trees without zero probabilities at any chance node, whose revealed conditional preference ordering* R^S_β *is equal to* R^S *for each non-empty* $S \subset E$*.*

VI. SUBJECTIVE PROBABILITY

We remain some way short of Anscombe and Aumann's (1963) formulation of subjective probabilities. For one thing, the ratios of those probabilities amount to marginal rates of substitution between expected von Neumann–Morgenstern utilities conditional on different events; so far, we have not imposed any continuity on behaviour in a way that even ensures the existence of a utility function. Another problem is that Anscombe and Aumann assumed that two probability distributions $\lambda, \mu \in \Delta(Y^S)$ would be equivalent, and so indifferent, whenever the marginal distributions $\lambda_s, \mu_s \in \Delta(Y_s)$ were equal for all $s \in S$. As pointed out in Section 12 of Hammond (1988), this assumption is crucial in ruling out the kind of preference pattern observed in Ellsberg's (1961) "paradox." Indeed, such patterns cannot be excluded by the consequentialist axioms on their own without some help from additional plausible assumptions.

The continuity issue is easily treated. Consider a family of finite decision trees T^π which are all identical except for the collection $\pi = \langle \pi(\cdot|n) \rangle_{n \in N^0} \in \prod_{n \in N^0} \Delta^0(N_{+1}(n))$ of strictly positive probability distributions at each chance node. For each common decision node $n^* \in N^*$, there is an induced correspondence $\pi \mapsto\!\!\!\rightarrow \beta(T^\pi, n^*)$ from the domain $\prod_{n \in N^0} \Delta^0(N_{+1}(n))$ of allowable transition probability distributions at different chance nodes to the range of non-empty subsets of $N_{+1}(n^*)$. Now, continuity of behaviour generally requires such a correspondence to be upper hemi-continuous. Where the domain and the range are both compact sets, as they are here, upper hemi-continuity is equivalent to the following closed graph property:

Axiom 4 (continuous behaviour as probabilities vary). *For each decision node $n^* \in N^*$ the graph $\{ (\pi, n') \in \prod_{n \in N^0} \Delta^0(N_{+1}(n)) \times N_{+1}(n^*) \mid n' \in \beta(T^\pi, n^*) \}$ of the correspondence $\pi \mapsto\!\!\!\rightarrow \beta(T^\pi, n^*)$ is a relatively closed set.*

As is fairly easy to show, it then follows that for each non-empty $S \subset E$, the revealed conditional preference relation R^S_β on $\Delta(Y^S)$ must have the property that the two sets

$$\{ \alpha \in [0,1] \mid \alpha\, \lambda + (1-\alpha)\, \mu \; R^S_\beta \; \nu \} \quad \text{and} \quad \{ \alpha \in [0,1] \mid \nu \; R^S_\beta \; \alpha\, \lambda + (1-\alpha)\, \mu \}$$

are closed, for each triple $\lambda, \mu, \nu \in \Delta(Y^S)$. This continuity property is one of Herstein and Milnor's (1953) three axioms, applied to the binary relation R^S_β. The other two are that

R^S_β is an ordering and that the independence condition is satisfied. The implication of their main theorem is the existence, for each non-empty $S \subset E$, of a unique cardinal equivalence class of *conditional von Neumann–Morgenstern utility functions* (NMUFs) v^S such that R^S_β is represented on $\Delta(Y^S)$ by the expected value $\mathbb{E} v^S := \sum_{y^S \in Y^S} p^S(y^S)\, v^S(y^S)$ of v^S, where $p^S(y^S)$ denotes the probability of y^S.

The second issue, regarding the sufficiency of considering only marginal probability distributions, is much less straightforward. In fact, two extra assumptions are generally needed to ensure the existence of subjective probabilities. Of these, the first is:

Axiom 5 (certainty equivalence). *Suppose that T and T' are two decision trees without any natural nodes in which the only differences are in the event sets $S(n)$, $S'(n)$ at each node n of the common set of nodes N, and in the associated consequences $\gamma(x)$, $\gamma'(x)$ which occur at each terminal node x of the common set of terminal nodes $X = X'$. Specifically, suppose that in tree T there exists a single state $e \in E$ such that $S(n) = \{e\}$ for all $n \in N$, while $\gamma(x)$ is a riskless consequence in Y for all $x \in X$. On the other hand, suppose that in tree T' one has $S'(n) = S$ for all $n \in N$, while $\gamma'(x) = \gamma(x) 1^S$ for all $x \in X$, where $\gamma(x) 1^S \in Y^S$ denotes the particular constant contingent consequence function whose value is $\gamma(x)$ in each state of the world $s \in S$. Then the behaviour sets $\beta(T, n^*)$ and $\beta(T', n^*)$ are equal at each common decision node $n^* \in N^*$ of the two trees T and T'.*

Thus, the decision tree T' in which there is no uncertainty about the consequence at each terminal node, even though there may be uncertainty about the state of the world, is regarded as equivalent to the tree T in which there is not even any uncertainty about the state of the world. In fact, given the special property of tree T', one can regard T as an alternative "certainty equivalent" decision tree. Because of Theorem 3 it should be no surprise that, together with Axioms 1–3, this new assumption implies the existence of a state-independent revealed preference ordering R^*_β on $\Delta(Y)$ which is equal to $R^{\{s\}}_\beta$ on $\Delta(Y_s)$ for every $s \in E$, and also equal to the restriction of R^S_β to $\Delta(Y\, 1^S)$ for every non-empty $S \subset E$ (where $Y\, 1^S$ denotes the set of constant contingent consequences of the form $y\, 1^S$, for some $y \in Y$). In particular, Axiom 5 rules out awkward examples such as that presented in Section IV.

The second additional assumption is:

Axiom 6 (three consequences). *There exist at least three consequences $y_1, y_2, y_3 \in Y$ such that $y_3 \; P_\beta \; y_2 \; P_\beta \; y_1$, where P_β denotes the strict preference relation corresponding to the revealed preference ordering R_β.*

Axioms 5 and 6 together rule out rather strange conditional NMUFs such as $v^S(y^S) = \prod_{s \in S} v_s(y_s)$ or $-\prod_{s \in S} [-v_s(y_s)]$, and instead imply that $v^S(y^S) = \sum_{s \in S} v_s(y_s)$ for a unique co-cardinal equivalence class of state contingent NMUFs $v_s : Y_s \to \Re$ $(s \in E)$. Moreover, Axiom 5 in particular, Axiom 5 implies that there exists a unique cardinal equivalence class of state-independent NMUFs $v^* : Y \to \Re$ with the property that each v_s is cardinally equivalent to v^*. Hence there must exist additive constants α_s and positive multiplicative constants ρ_s (all $s \in E$) such that $v_s(y) \equiv \alpha_s + \rho_s \, v^*(y)$. In fact, the family of constants ρ_s $(s \in E)$ is unique up to a common multiplicative constant, and so there are well defined *revealed conditional probabilities* given by $P(s|S) := \rho_s / \sum_{e \in S} \rho_e$ whenever $s \in S \subset E$. Finally, as in Weller (1978), whenever $s \in S \subset S' \subset E$ these revealed conditional probabilities must satisfy

$$P(s|S') = \frac{\rho_s}{\sum_{e \in S'} \rho_e} \times \frac{\sum_{e \in S} \rho_e}{\sum_{e \in S'} \rho_e} = P(s|S) \, P(S|S').$$

The important implication is:

Bayes' rule. *Whenever $s \in S \subset S' \subset E$, it must be true that $P(s|S') = P(s|S)P(S|S')$.*

These arguments help to justify the following main theorem, also proved in Hammond (1988):

Theorem 4 (subjective expected utility maximization). *Suppose that Y is a given consequence domain, that E is a given non-empty finite set of possible states of the world. Then:*

(A) Suppose that behaviour satisfies Axioms 1–6 for the domain of all finite decision trees with chance and/or natural nodes in which there are no zero probability moves at any chance node. Then there must exist a family of revealed conditional preference orderings R^S_β on $\Delta(Y^S)$, one for each non-empty $S \subset E$. Moreover, these orderings must in turn reveal a unique cardinal equivalence class of state independent NMUFs $v^* : Y \to \Re$ and unique positive conditional probabilities $P(s|S)$ $(s \in S \subset E)$ satisfying

Bayes' rule, such that R^S_β is represented by the subjective expected utility expression

$$\sum\nolimits_{y^S \in Y^S} p^S(y^S) \sum\nolimits_{s \in S} P(s|S)\, v^*(y^s) \equiv \sum\nolimits_{s \in S} P(s|S)\, p_s(y_s)\, v^*(y^s)$$

where p_s denotes the marginal probability distribution on Y_s generated by p^S on Y^S.

(B) *Conversely, given any state independent NMUF $v^* : Y \to \Re$ and family of positive conditional probabilities $P(s|S)$ ($s \in S \subset E$) satisfying Bayes' rule, the associated behaviour correspondence β that maximizes subjectively expected utility must satisfy Axioms 1–5 in all decision trees without zero probabilities at any chance node.*[9]

VII. GAME THEORY AND RATIONALIZABILITY

Of late the most widespread use in economics of the expected utility model of decision-making under uncertainty has been in non-cooperative game theory. Following prominent works such as Aumann (1987) and Tan and Werlang (1988), most game theorists now take the view that players in a game should have beliefs about other players' strategies described by subjective probabilities, and that they should then choose their strategies to maximize their respective expected utilities. This seems at first to be an entirely natural use of orthodox decision theory. Yet there is an important difference in extensive form games between, on the one hand, natural nodes at which nature moves exogenously, and on the other hand, players' information sets at which moves are determined endogenously by maximizing the relevant player's expected utility. In particular, I have often seen it claimed that, in this game theoretic context, the existence of subjective probabilities and the maximization of subjectively expected utility are justified by Savage's axioms. Apart from betraying a fondness for Savage's particular set of axioms which may be hard to justify, this overlooks the fundamental issue of whether it makes any sense at all to apply Savage's axioms, the Anscombe–Aumann axioms, or some similar collection such as Axioms 1–6 above, to strategic behaviour in non-cooperative games.[10]

In fact there is one very clear difference between classical decision theory and orthodox game theory. In decision theory, all subjective probability distributions over unknown states

[9] Evidently Axiom 6 cannot be an implication of expected utility maximization if there are only one or two consequences.

[10] A similar point is the main concern of the paper by Mariotti presented to the same conference.

of the world are regarded as equally valid and equally rational. Game theory, by contrast, started out by attempting to determine players' beliefs endogenously. Indeed, it appears that the need Morgenstern (1928, 1935) had perceived to close certain economic models by determining expectations was what aroused his interest in von Neumann's (1928) early work on the theory of "party games." Much later, Johansen (1982) still felt able to argue that, for games with a unique Nash equilibrium, that equilibrium would entirely determine rational behaviour and rational beliefs within the game.

Since the work of Bernheim (1984, 1986) and Pearce (1984) on rationalizability, and of Aumann (1987) on correlated equilibria, game theorists have begun to pay more careful attention to the question of what rational beliefs players should hold about each other. Aumann follows the equilibrium tradition of Nash in looking for a set of common expectations over everybody's strategies (a "common prior") that attaches probability one to all players choosing optimal strategies given those common expectations. Bernheim and Pearce relax this condition and look for "rationalizable" expectations which can differ between players, but must attach probability one to all players choosing strategies that are rationalizable — i.e., optimal given their own rationalizable expectations. Thus, in the work on rationalizability, what becomes endogenous is each player's set of rationalizable strategies, rather than expectations about those strategies.

Rationalizable expectations are especially interesting because of the way in which they, and the associated rationalizable strategies, can be constructed recursively. First order rationalizable expectations are arbitrary; first order rationalizable strategies are optimal given first order rationalizable expectations; second order rationalizable expectations attach probability one to players choosing first order rationalizable strategies; second order rationalizable strategies are optimal given second order rationalizable expectations; third order rationalizable expectations attach probability one to players choosing second order rationalizable strategies; and so on. The result of this recursive construction is a diminishing sequence of sets of n-th order rationalizable strategies and of associated sets of rationalizable beliefs or expectations. In the end, rationalizable strategies are those which are n-th order rationalizable, for all natural numbers n. What is happening here is that our concept of rationality (or rationalizability) is becoming more refined as n increases and we progress further up the hierarchy.

Though the usual consequentialist axioms cannot be applied to multi-person games, it seems that a variant of them can be. This involves the idea of *conditionally rational behaviour* for each player, based on *hypothetical probabilities* attached to the other players' strategies. Specifically, it is assumed that each player's behaviour satisfies the consequentialist axioms in all decision trees that result from extensive games by attaching specific hypothetical probabilities to all other players' moves at each of their information sets. This implies that each player maximizes expected utility, given these hypothetical probabilities. It then remains to determine, as far as possible, what hypothetical probabilities represent beliefs that are rational, or at least rationalizable. This, of course, is more or less what non-cooperative game theory has been trying to accomplish since its inception.

At the moment, then, it seems that game theory says only that it is irrational to attach positive probabilities to rational players choosing strategies that are not rationalizable. Further restrictions on rational beliefs may emerge from subgame perfection and "forward induction" arguments.[11] But other restrictions beyond these are hard to motivate, at least in the current state of game theory. Above all, only if rationalizable strategies are unique should we expect rational beliefs to be determined uniquely.

VIII. STRUCTURAL RATIONALITY VERSUS BOUNDED RATIONALITY

Section VI concluded by showing that Axioms 1–5 and subjective expected utility maximization are equivalent for a consequence domain satisfying Axiom 6. The behaviour so described satisfies some well-known "structural" conditions which have often been regarded as required for logical consistency or coherence. Hence I call such behaviour *structurally rational*. It is important to understand that structural rationality is neither sufficient nor necessary for full rationality.

Insufficiency of structural rationality is fairly evident. Maximizing the expected value of *any* utility function with respect to *any* subjective probabilities is structurally rational. So is *minimizing* expected utility! More seriously, a fuller concept of rationality clearly requires the pursuit of appropriate ends and the holding of reasonal beliefs, neither of which is entailed by structural rationality, or by consequentialism.

[11] The relationship between forward induction and "conditional rationalizability" is discussed in Hammond (1993b).

Less evident is the fact that structural rationality is not necessary for rational behaviour either. I am prepared to claim that it would be necessary if there were no bounds at all on an agent's rationality. If all our decision problems were straightforward, rationality could well entail behaviour satisfying Axioms 1–5, for a suitable domain of consequences. In practice, however, reality confronts us with enormously complex decision problems. Modelling these as decision trees, and then constructing the sets $F(T)$ and $\Phi_\beta(T)$ as in Section II, is a task which poses horrendous difficulties for the decision analyst, quite apart from any human agent. After all, as has often been noted, it is impossible to model completely even a problem as well structured as how to play chess. Simplification is inevitable.

I shall assume that simplification gives rise to a rather small finite decision tree which models the true decision problem only imperfectly. Many nodes may be omitted, as may many of the possibilities at each node. In this sense there is a *bounded model.* And *bounded rationality* would seem to involve the use of such bounded models. Then there seems no particular reason why an agent's rational behaviour should exhibit structural rationality. After all, behaviour cannot then be explained only by its consequences, but will also depend upon what happens to be included in the agent's decision model.

Nevertheless consequentialism, together with the structural rationality it entails, still makes sense within whatever bounded model lies behind the agent's decision. In tournament chess it is rational not to analyse so deeply that one loses on time. But rationality does demand appropriate patterns of thought in whatever analysis does get conducted within a limited time. The incomplete plan (or "variation") chosen by a player should be the best given the possibilities that could be considered within the time which that player felt able to devote to making that plan. A good player will not deliberately make an inferior move; nor probably will a bad player! Bad moves and even blunders may occur, but only because important considerations and possibilities have been "overlooked" — i.e., omitted from the player's bounded model.

What bounded rationality forces us to consider are *vaguely formulated decision trees.* These will be just like ordinary decision trees, except that both consequences $y \in Y$ and states $s \in E$ will be vaguely formulated — imprecisely described, in effect. It is impossible to distinguish all the consequences of the true decision tree, or all the possible states. Instead, the agent can only analyse a coarsening Y_0 of the consequence domain Y, and a coarsening E_0 of the state space E. The different elements $y_0 \in Y_0$ will determine a partition $\cup_{y_0 \in Y_0} Y^0(y_0)$ of Y into sets of consequences that the agent fails to distinguish. There is a similar partition $\cup_{s_0 \in E_0} E^0(s_0)$ of E. The latter partition is equivalent to Savage's (1954) model of "small worlds."

Such coarsenings already imply that the agent can consider only decision trees in the domain $\mathcal{T}(E_0, Y_0) = \cup_{\emptyset \neq S_0 \subset E_0} \mathcal{T}^{S_0}(Y_0)$ rather than in the true domain $\mathcal{T}(E, Y) = \cup_{\emptyset \neq S \subset E} \mathcal{T}^S(Y)$. But the game of chess, for example, already has very simple consequences — win, lose or draw. Nor is there room for any intrinsic uncertainty. In chess, the necessary coarsening affects the decision tree itself. The domain $\mathcal{T}(E_0, Y_0)$ gets coarsened to $\mathcal{T}_0(E_0, Y_0)$, each of whose members T_0 corresponds to an enormous class $\mathcal{T}^0(E_0, Y_0; T_0)$ of undistinguished complex decision trees. Such coarsening of the set of possible trees can also occur in any other decision problem, of course.

This approach to bounded rationality differs from the well know one due to Simon (1972, 1982, 1986, 1987a, b), because he appears to allow an agent to "satisfice" even within his own model. By contrast, I am asking for optimality within the agent's own model, though of course satisficing is allowed and even required in the choice of that model. This seems much more reasonable as a standard for normative behaviour.

IX. EXPANDING SMALL WORLDS AND IMPROVING BOUNDED RATIONALITY

The discussion in Section VII on rationalizability in games suggests, however, that one should go beyond such fixed bounded models. After all, there is a sense in which rationalizable strategies emerge from increasingly complex models of a game. Initially other players' strategies are treated as entirely exogenous, then as best responses to exogenous beliefs, then as best responses to best responses to exogenous beliefs, etc. A strategy is rationalizable if and only if enriched models of the other players and of their models can never demonstrate the irrationality of that strategy.

Similarly, consider a particular bounded model based on the particular coarsenings Y_0 of Y, E_0 of E, and $\mathcal{T}_0(E_0, Y_0)$ of $\mathcal{T}(E_0, Y_0)$. One feels that a rational agent who uses this model should have some reason for not using a richer bounded model. Such a richer model would be based on coarsenings Y_1 of Y, E_1 of E, and $\mathcal{T}_1(E_1, Y_1)$ of $\mathcal{T}(E_1, Y_1)$ which are refinements of Y_0, E_0 and $\mathcal{T}_0$ respectively. That is, for each $y_1 \in Y_1$, the class $Y^1(y_1)$ of consequences which the agents fails to distinguish from y_1 should be a subset of $Y^0(y_0)$ for the unique $y_0 \in Y_0$ such that $y_1 \in Y^0(y_0)$. Similarly, for each $e_1 \in E_1$ and $T_1 \in \mathcal{T}_1(E_1, Y_1)$, it should be true that $E^1(e_1) \subset E^0(e_0)$ where $e_1 \in E^0(e_0)$, and also that $\mathcal{T}^1(E_1, Y_1; T_1) \subset \mathcal{T}^0(T_0)$ where $T_1 \in \mathcal{T}^0(T_0)$.

In fact, as Behn and Vaupel (1982) and Vaupel (1986) have suggested, what one really expects of a *rational* boundedly rational agent are some some reasonable beliefs about whether a more refined model would change the decision being made at the current node of life's decision tree, as well as reasonable beliefs about what the consequences of a revised decision are likely to be. It is as though the agent were involved in a complicated game with some other very imperfectly known players — namely, versions of the same agent who use more complicated decision models. Should the agent take the trouble to enrich the decision model and so become one of these other players, or is it better to remain with the existing bounded model?

Note that an agent who does choose a more complicated model then has the same kind of decision problem to face once again — namely, the decision whether to complicate the model still further. In this way the agent can be viewed as facing an uncertain potentially infinite hierarchy of vaguely formulated trees, with increasing refined consequence spaces Y_n, state spaces E_n, and tree spaces $\mathcal{T}_n(E_n, Y_n)$ $(n = 1, 2, \ldots)$.

There should then be an associated sequence β_n $(n = 1, 2, \ldots)$ of behaviour correspondences, each satisfying Axioms 1–5 for the appropriate domain $\mathcal{T}_n(E_n, Y_n)$ of decision trees. So there will exist increasingly refined NMUFs $v_n : Y_n \to \Re$ and subjective probabilities $P_n(s_n|E_n)$ $(s_n \in E_n)$.

At this stage one should look for consistency conditions which it is reasonable to impose on different members of this hierarchy. And consider in more detail the decision when to stop analysing the current decision problem more deeply. Thus, a model of rational behaviour becomes hierarchical, and even self-referential, as in the theory of games. There also seem to be some links, of which I am only dimly aware, with recent developments in the theory of the mind and of self-awareness. This seems to be an inevitable implication of bounds on rationality. Nor should such links be at all surprising. However, this is still largely an unexplored topic, as far as I am aware. In fact, I may be proposing going at least one step further in the hierarchy of more and more complicated collective decision problems concerning how best to describe rational behaviour. If so, who knows what the next enriched models are likely to be?

Notes

1. Of course, it is restrictive in general to consider only finite decision trees. But not when discussing the implications of consequentialism, as is done here. The issue with infinite decision trees is whether behaviour can be well defined in a way that naturally extends consequentialist behaviour in finite trees. Obviously, this will require a technical analysis of compactness and continuity conditions.

2. This formulation does exclude stochastic behaviour, according to which $\beta(T, n)$ is a probability distribution over $N_{+1}(n)$. Under the consequentialist axioms set out below, it turns out that a very strong transitivity axiom — or what Luce (1958, 1959) describes as a "choice axiom" — must be satisfied. In fact, there must exist a preference ordering over consequences which is maximized by all the consequences that can occur with positive probability given the agent's stochastic behaviour. Thus randomization occurs only over consequences in the highest indifference class of feasible consequences available to the agent. Furthermore, each consequence must have a positive real number attached indicating the relative likelihood of that consequence occurring, in case it belongs to the highest indifference class that the agent can reach.

3. Actually, Prasanta Pattanaik and Robert Sugden tried to convince me of essentially this point many years ago. Jean-Michel Grandmont and Bertrand Munier were finally more successful during the conference. I am grateful to all four, and would like to apologize to the first two.

4. See Hammond (1976) for more discussion of naïve behaviour in decision trees, especially the "potential addict" example. Also, Amos Tversky has suggested that Axiom 2 be called "subtree consistency." I am sympathetic, but have not followed this suggestion for two reasons: (i) earlier papers used the term "dynamic consistency," and I would like to be consistent myself; (ii) subtrees do represent dynamic choice possibilities when the tree models decisions which will be made in real time.

5. For obvious reasons, one does not speak of (logical) "consequences."

6. Following the terminology of social choice theory, a “preference ordering” means a complete and transitive binary weak preference relation.

7. Some critics of consequentialism, notably Machina (1989) and McClennen (1990), have offered such explanations, especially for behaviour that maximizes a preference ordering and yet violates the independence condition which, as will be discussed in Section V, is one other implication of consequentialism.

8. An important generalization allows the consequence domain Y_s to depend on the state s. As one might expect, it then becomes harder to ensure that behaviour reveals subjective probabilities. However, both the sure thing principle and independence do still follow from the consequentialist axioms.

9. Evidently Axiom 6 cannot be an implication of expected utility maximization if there are only one or two consequences.

10. A similar point is the main concern of the paper by Mariotti presented to the same conference.

11. The relationship between forward induction and “conditional rationalizability” is discussed in Hammond (1993b).

References

Anscombe, F.J. and Aumann, R.J. (1963) “A Definition of Subjective Probability”, *Annals of Mathematical Statistics*, vol. 34, pp. 199–205.

Aumann, R.J. (1987) “Correlated Equilibrium as an Expression of Bayesian Rationality”, *Econometrica*, vol. 55, pp. 1–18.

Behn, R.D. and Vaupel, J.W. (1982) *Quick Analysis for Busy Decision Makers* (New York: Basic Books).

Bernheim, B.D. (1984) “Rationalizable Strategic Behavior”, *Econometrica*, vol. 52, pp. 1007–1028.

Bernheim, B.D. (1986) “Axiomatic Characterizations of Rational Choice in Strategic Environments”, *Scandinavian Journal of Economics*, vol. 88, pp. 473–488.

Ellsberg, D. (1961) "Risk, Ambiguity, and the Savage Axioms", *Quarterly Journal of Economics*, vol. 75, pp. 643–669.

Hammond, P.J. (1976) "Changing Tastes and Coherent Dynamic Choice", *Review of Economic Studies*, vol. 43, pp. 159–173.

Hammond, P.J. (1988) "Consequentialist Foundations for Expected Utility", *Theory and Decision*, vol. 25, pp. 25–78.

Hammond, P.J. (1993a) "Elementary Non-Archimedean Representations of Probability for Decision Theory and Games", to appear in Humphreys, P. (ed.) *Patrick Suppes: Scientific Philosopher* (Dordrecht: Kluwer Academic Publishers).

Hammond, P.J. (1993b) "Aspects of Rationalizable Behavior", to appear in Binmore, K., Kirman, A.P. and Tani, P. (eds.) *Frontiers of Game Theory* (Cambridge, Mass.: M.I.T Press).

Hammond, P.J. (1994) "Consequentialism, Non-Archimedean Probabilities, and Lexicographic Expected Utility", presented to the Second Workshop on "Knowledge, Belief and Strategic Interaction: The Problem of Learning" at Castiglioncello (Livorno), Italy in June 1992; to appear in the conference proceedings published by Cambridge University Press.

Herstein, I.N. and Milnor, J. (1953) "An Axiomatic Approach to Measurable Utility", *Econometrica*, vol. 21, pp. 291–297.

Johansen, L. (1982) "On the Status of the Nash Type of Noncooperative Equilibrium in Economic Theory", *Scandinavian Journal of Economics*, vol. 84, pp. 421–441.

Luce, R.D. (1958) "An Axiom for Probabilistic Choice Behavior and an Application to Choices among Gambles (abstract)", *Econometrica*, vol. 26, pp. 318–319.

Luce, R.D. (1959) *Individual Choice Behavior* (New York: John Wiley).

Machina, M.J. (1989) "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty", *Journal of Economic Literature*, vol. 27, pp. 1622–1668.

McLennen, E.F. (1990) *Rationality and Dynamic Choice* (Cambridge: Cambridge University Press).

Morgenstern, O. (1928) *Wirtschaftsprognose: Eine Untersuchung ihrer Voraussetzungen und Möglichkeiten* [Economic Forecasting: An Investigation of its Presuppositions and Possibilities] (Vienna: Julius Springer).

Morgenstern, O. (1935) "Vollkommene Voraussicht und wirtschaftliches Gleichgewicht [Perfect Foresight and Economic Equilibrium]", *Zeitschrift für Nationalökonomie*, vol. 6, pp. 337–357.

von Neumann, J. (1928) "Zur Theorie der Gesellschaftsspiele [On the Theory of Party Games]", *Mathematische Annalen*, vol. 100, pp. 295–320.

Pearce, D. (1984) "Rationalizable Strategic Behavior and the Problem of Perfection", *Econometrica*, vol. 52, pp. 1029–1050.

Savage, L.J. (1954) *The Foundations of Statistics* (New York: John Wiley).

Simon, H.A. (1972, 1986) "Theories of Bounded Rationality", in McGuire, C.B. and Radner, R. (eds.) *Decision and Organization (2nd edn.)* (Minnesota: University of Minnesota Press) ch. 8, pp. 161–176.

Simon, H.A. (1982) *Models of Bounded Rationality* (2 vols.) (Cambridge, Mass.: M.I.T. Press).

Simon, H.A. (1986) "Rationality in Psychology and Economics", *Journal of Business*, vol. 59 (Supplement), pp. 25–40.

Simon, H.A. (1987a) "Bounded Rationality", in Eatwell, J., Milgate, M. and Newman, P. (eds.) *The New Palgrave: A Dictionary of Economics* (London: Macmillan).

Simon, H.A. (1987b) "Satisficing", in Eatwell, J., Milgate, M. and Newman, P. (eds.) *The New Palgrave: A Dictionary of Economics* (London: Macmillan).

Tan, T.C.-C. and Werlang, S.R.da C. (1988) "The Bayesian Foundations of Solution Concepts of Games", *Journal of Economic Theory*, vol. 45, pp. 370–391.

VAUPEL, J.W. (1986) “Un pensiero analitico per decisori indaffarati [An Analytical Thought for Busy Decision-makers]”, in Sacconi, L. (ed.) *La decisione: Razionalità collettiva e strategia nell'amministrazione e nelle organizzazioni* (Milano: Franco Angeli) ch. 13, pp. 226–238.

WELLER, P.A. (1978) “Consistent intertemporal decision making under uncertainty”, *Review of Economic Studies*, vol. 45, pp. 263–266.