

Rationality in Economics

PETER J. HAMMOND

Department of Economics, Stanford University, CA 94305-6072, U.S.A.
e-mail: hammond@leland.stanford.edu

1 Introduction and Outline

Rationality is one of the most over-used words in economics. Behaviour can be rational, or irrational. So can decisions, preferences, beliefs, expectations, decision procedures, and knowledge. There may also be bounded rationality. And recent work in game theory has considered strategies and beliefs or expectations that are “rationalizable”.

Here I propose to assess how economists use and mis-use the term “rationality.” Most of the discussion will concern the normative approach to decision theory. First, I shall consider single person decision theory. Then I shall move on to interactive or multi-person decision theory, customarily called game theory. I shall argue that, in normative decision theory, rationality has become little more than a structural consistency criterion. At the least, it needs supplementing with other criteria that reflect reality. Also, though there is no reason to reject rationality hypotheses as normative criteria just because people do not behave rationally, even so rationality as consistency seems so demanding that it may not be very useful for practicable normative models either.

Towards the end, I shall offer a possible explanation of how the economics profession has arrived where it is. In particular, I shall offer some possible reasons why the rationality hypothesis persists even in economic models which purport to be descriptive. I shall conclude with tentative suggestions for future research — about where we might do well to go in future.

2 Decision Theory with Measurable Objectives

In a few cases, a decision-making agent may seem to have clear and measurable objectives. A football team, regarded as a single agent, wants to score more goals than the opposition, to win the most matches in the league, etc. A private corporation seeks to make profits and so increase the value to its owners. A publicly owned municipal transport company wants to provide citizens with adequate mobility at reasonable fares while not requiring too heavy a subsidy out of general tax revenue. A non-profit organization like a university tends to have more complex objectives, like educating students, doing good research, etc. These conflicting aims all have to be met within a limited budget.

Measurable objectives can be measured, of course. This is not always as easily as keeping score in a football match or even a tennis, basketball or cricket

match. After all, accountants often earn high incomes, ostensibly by measuring corporate profits and/or earnings.

For a firm whose profits are risky, shareholders with well diversified portfolios will want that firm to maximize the expectation of its stock market value. If there is uncertainty about states of the world with unknown probabilities, each diversified shareholder will want the firm to maximize subjective expected values, using the shareholder's subjective probabilities. Of course, it is then hard to satisfy all shareholders simultaneously. And, as various recent spectacular bank failures show, it is much harder to measure the extent to which profits are being made when there is uncertainly.

In biology, modern evolutionary theory ascribes objectives to genes — so the biologist Richard Dawkins has written evocatively of the *Selfish Gene*. The measurable objective of a gene is the extent to which the gene survives because future organisms inherit it. Thus, gene survival is an objective that biologists can attempt to measure, even if the genes themselves and the organisms that carry them remain entirely unaware of why they do what they do in order to promote gene survival.

Early utility theories up to about the time of Edgeworth also tried to treat utility as objectively measurable. The Age of the Enlightenment had suggested worthy goals like “life, liberty, and the pursuit of happiness,” as mentioned in the constitution of the U.S.A. Jeremy Bentham wrote of maximizing pleasure minus pain, adding both over all individuals. In dealing with risk, especially that posed by the St. Petersburg Paradox, in the early 1700s first Gabriel Cramer (1728) and then Daniel Bernoulli (1738) suggested maximizing expected utility; most previous writers had apparently considered only maximizing expected wealth.

3 Ordinal Utility and Revealed Preference

Over time, it became increasingly clear to economists that any behaviour as interesting and complex as consumers' responses to price and wealth changes could not be explained as the maximization of some objective measure of utility. Instead, it was postulated that consumers maximize unobservable subjective utility functions. These utility functions were called “ordinal” because all that mattered was the ordering between utilities of different consumption bundles. It would have been mathematically more precise and perhaps less confusing as well if we had learned to speak of an *ordinal equivalence class* of utility functions. The idea is to regard two utility functions as equivalent if and only if they both represent the same *preference ordering* — that is, the same reflexive, complete, and transitive binary relation. Then, of course, all that matters is the preference ordering — the choice of utility function from the ordinal equivalence class that represents the preference ordering is irrelevant. Indeed, provided that a preference ordering exists, it does not even matter whether it can be represented by any utility function at all.

This realization underlies the “revolution” in demand theory achieved when Hicks and Allen (1934) purged it of any notion of measurable utility. Only a

preference relation was needed or, at most, a ratio of marginal utilities at each consumption bundle. Moreover, this ratio, or marginal rate of substitution, would be preserved by any increasing differentiable transformation to a new ordinal equivalent utility function for which marginal utilities were also well defined.

At about the same time, Samuelson (1938) was advancing his theory of “revealed preference”. This was an important step toward being able to infer preferences from demand behaviour. But it also emphasized how demand theory could start with demand functions rather than utility functions, and could go on to ask when those demand functions were consistent with the maximization of a preference ordering. There were some important difficulties in the case of three or more goods which Samuelson had overlooked. These Houthakker (1950) helped to resolve with a stronger revealed preference axiom. Later Uzawa (1960) and others, following Houthakker’s important lead, conducted the necessary study of the differential equations that enable preferences to be inferred from demand functions. At roughly the same time, Arrow (1959) showed how it was simpler to state revealed preference axioms for an agent who could be confronted with general finite feasible sets, instead of only convex budget sets. Also, Afriat (1967) devised a procedure allowing one to test whether discrete observations of a consumer’s demands at different combinations of price and wealth are consistent with the existence of a preference ordering.

The effect of this revealed preference revolution was to create an entirely different and much weaker concept of rationality. Classical utilitarians had expected consumers to maximize a specific utility function. Later neo-classical theorists considered an ordinal equivalence class of utility functions, but left open the question whether rationality required severe restrictions on the associated preference ordering. Revealed preference theorists placed almost no restrictions on the preference ordering beyond monotonicity, and possibly continuity or convexity. Even these restrictions, it was agreed, should be abandoned in some contexts. It is possible to have too much of some usually good things, like good wine. So monotonicity is not always satisfied. Also, to adapt a similar French example due to Malinvaud (1972), either one dinner tonight in Milan, or else a nearly equivalent dinner tonight in Rome, would be better than half tonight’s dinner in Milan combined with half tonight’s dinner in Rome. So convexity may also be violated.

In the end, therefore, rationality in demand theory has been largely reduced to consistency conditions. Obviously, it requires behaviour to be consistent with some preference ordering. But behaviour must also be self-consistent in order to allow some revealed preference ordering to be inferred. Also, *any* ordering seems to be allowed. Clearly, then, consistency is insufficient for true rationality, since the latter surely restricts what it is reasonable to maximize. In particular, one should not minimize what it is rational to maximize!

4 Why Be Consistent?

Is consistency even necessary for rationality, however? Also, why should economic agents have any preferences, let alone a complete and transitive preference ordering? And if they do have preferences, why should behaviour maximize a preference relation?

On the whole, economists have seemed remarkably reluctant to face these questions, even though they do seem rather fundamental, given the extent to which the maximization hypothesis has dominated the recent economic literature. Later on I shall have more to say about possible explanations of this reluctance.

The existence of a preference ordering for society is, if anything, even more questionable than for an individual. So, since Arrow's impossibility theorem in social choice theory requires there to be a social ordering, it is not surprising that, on the last pages of the second (1963) edition of *Social Choice and Individual Values*, he sketched a justification for this key assumption. His defence was based on a notion of "path independence" requiring that, when a social decision is taken in several stages, the outcome should not depend on the path leading to the final decision. This idea was formalized a few years later by Plott (1973). But except in a rather strong form, path independence was shown not to imply the existence of a (transitive) preference ordering.

A different and more successful version of path independence was proposed by Campbell (1978). However, rather than the computational processes he considers, or the equivalent sequential procedure due to Bandyopadhyay (1988), it seems better to recognize that often an agent faces a decision tree, in which a sequence of real decisions has to be made. This idea is taken up later in the next and subsequent sections.

An alternative attempt to justify transitivity relies on the "money pump" argument due to Davidson, McKinsey and Suppes (1955). Suppose that a consumer has intransitive, even cyclic, preferences for consumption bundles a, b, c , with $a \succ b$, $b \succ c$, and $c \succ a$. In fact, consider the individual's preferences for combinations (x, m) of consumption bundles x and amounts of money m . Suppose that there exists a small amount of money $\epsilon > 0$ such that, for all m , the agent's preferences satisfy not only $(a, m) \succ (b, m)$, $(b, m) \succ (c, m)$ and $(c, m) \succ (a, m)$, but also:

$$(a, m - \epsilon) \succ (b, m); \quad (b, m - \epsilon) \succ (c, m); \quad (c, m - \epsilon) \succ (a, m)$$

That is, the individual is always willing to pay somewhat more than ϵ in order to get a instead of b , or b instead of c , or c instead of a . This is in part a form of "uniform continuity" of preferences, though it also requires that the willingness to pay for each change is bounded away from 0. With these preferences one has

$$(a, m - 3\epsilon) \succ (b, m - 2\epsilon); \quad (b, m - 2\epsilon) \succ (c, m - \epsilon); \quad (c, m - \epsilon) \succ (a, m)$$

Hence, for $n = 0, 1, 2, \dots$, it follows by induction on n that

$$\begin{aligned} (a, m - (3n + 3)\epsilon) &\succ (b, m - (3n + 2)\epsilon), \\ (b, m - (3n + 2)\epsilon) &\succ (c, m - (3n + 1)\epsilon), \\ \text{and } (c, m - (3n + 1)\epsilon) &\succ (a, m - 3n\epsilon) \end{aligned}$$

So, by allowing the agent to cycle between the successive consumption bundles c, b, a, c, b, a, \dots provided that ϵ units of money are paid every time there is a change, it is possible to “pump” an indefinite amount of money out of the consumer.

As already remarked, the above argument relies on the consumer’s willingness to pay for each change being bounded away from zero. Or if not, that each successive willingness to pay at least gives an infinite series whose sum is infinite. Otherwise, to paraphrase Mongin (1994), the agent may lose ϵ several times over, but is unlikely to reach bankruptcy. In fact, the money pump really only makes sense in a partial equilibrium context where “money” is separated from all other commodities, and one ignores income effects which could alter the consumer’s willingness to pay for a instead of b , or b instead of c , or c instead of a . For this reason, the argument fails to establish that all intransitivities should be removed from a consumer’s preference relation. Moreover, the money pump argument does not explain why a consumer’s choice from an arbitrary feasible set should be determined by preferences, which only purport to explain choices from pairs.

To conclude, it seems that neither path independence nor the money pump argument really succeeds in justifying the existence of a preference ordering. Instead, I shall turn to decision trees.

5 Inconsistent Behaviour in Decision Trees

An important paper by Strotz (1957) considered how changing tastes would affect savings behaviour. Strotz noticed the related problem faced by Odysseus when passing the Sirens and their enchantingly seductive singing. Odysseus (or Ulysses) faced an interesting decision tree. So does the teenager wondering whether to start smoking. Or anybody facing a choice which is liable to change one’s tastes. It is captured by the *potential addict* problem, with a decision tree illustrated in Figure 1.




Figure 1: The potential addict’s decision tree

At initial node n_0 , the agent can start some addictive activity, like smoking or sailing within earshot of the Sirens. That leads to a second decision node n_1 . Or the agent can decline and reach consequence c , which is more or less equivalent to the status quo. At node n_1 , the agent has a choice between a , which is to continue the addictive activity, or b , which is to desist. Option b allows Odysseus to hear the Sirens without sailing onto the rocks, and teenagers to enjoy a few years of smoking while limiting the damage to their own health. Option a involves Odysseus sailing onto the rocks, or the teenager continuing to smoke indefinitely regardless of risks to health. Because the activity is addictive, typically one has $a P_1 b$ according to the preference relation P_1 that applies at node n_1 . But at node n_0 , option b seems best and a seems worst, so $b P_0 c P_0 a$.

Faced with this possible change of tastes, a naive agent ignores the issue entirely, and plans to reach what is initially the best option b by beginning the addictive activity and moving first to n_1 . But once n_1 has been reached, addiction has set in and so the awful option a is the result. On the other hand, a sophisticated agent foresees that a will be chosen at n_1 , so b is not really feasible. Left only with the choice between a and c , the agent chooses c and avoids any risk of addiction.

Odysseus dealt with the issue in a different way, by precommitment. He found a different option d which reproduced for him the favourable consequences of b without the chance to choose a after addiction had set in. That is, he could hear the Sirens without being able to direct his ship onto the rocks. In effect, he collapsed the decision tree so that there was only one decision node. His poor crew, of course, were left with consequences similar to c , since they were not allowed to hear the Sirens as they sailed past their island. But at least most of them got a much greater chance to reach home in the end, one supposes.

In this potential addict example, in effect there is one agent at n_0 , and an entirely different agent at n_1 who has different tastes, especially as regards the pair $\{a, b\}$. The naive agent at n_0 who ignores this duality seems obviously irrational. The sophisticated agent at n_0 is rational, in a sense, but achieves rationality only by realizing the truth that there is a two-person game in extensive form instead of a single-person decision tree. On the other hand, precommitment is an instance of intelligent “lateral thinking,” involving finding a new option to add to the decision tree. Note that the structure of the decision tree is crucial; if there were only one decision node, then for behaviour to be reasonable, it must lead to the consequence b .

The potential addict example is a striking instance of the *essential inconsistencies* analysed in Hammond (1976). See also Elster (1979) and many succeeding works. Another less striking example of inconsistency arises when the Pareto criterion or some other incomplete preference relation is applied in a decision tree. For example, suppose that there is a society of two individuals i and j whose preferences or ordinal utilities for the four social states $\{a, b, c, d\}$ satisfy

$$u_i(b) > u_i(d) > u_i(a) > u_i(c) \quad \text{and} \quad u_j(a) > u_j(c) > u_j(b) > u_j(d)$$

as illustrated in Figure 2. Then the strict preference relation P generated by the Pareto condition satisfies both $a P c$ and $b P d$, but all four other pairs which are subsets of $\{a, b, c, d\}$ are incomparable. In particular, a and b are efficient or acceptable choices, whereas c and d are inefficient or unacceptable.




Figure 2: A utility possibility set and the decision tree T

Consider next the decision tree illustrated in Figure 2. At the initial node n_0 , the agent may move to n_1 , intending to continue to a , or to n'_1 , intending to continue to b . However, at n_1 the option d is no longer inferior because b , the only consequence preferred to d , is no longer feasible. Similarly, at n'_1 the option c is no longer inferior because a , the only consequence preferred to c , is no longer feasible. So, in order to avoid reaching one of the two inferior consequences, the agent at n_1 must remember the plan to choose a because d is inferior to the foregone option b , and the agent at n'_1 must remember the plan to choose b because c is inferior to the foregone option a .

This need to remember foregone options, however, raises the following question: Was there another option e which is Pareto superior to either a or b , but was passed over before reaching n_0 with the intention of proceeding to whichever of a or b is not inferior to e ? In fact, the decision tree may no longer be an adequate description of the agent's decision problem; it is possible that a larger decision tree, of which the current tree is only a subtree, may also be relevant.

As with the potential addict example, the outcome that results from applying the Pareto criterion or other incomplete preference relation to a decision tree can depend on the structure of the tree. To see this, note that if there were only one decision node, there would be no difficulty in choosing either a or b and in rejecting both c and d once and for all.

6 Consequentialism and Normal Form Consistency

The two examples of the previous section show how the structure of the decision tree can influence the set of chosen consequences. This, however, violates what seems to be a fundamental principle in most normative models of rational behaviour. In these normative models, a single feasible set F of different options is specified. Then the principle requires that some choice subset $C(F) \subset F$ must describe what options should result from normatively appropriate behaviour. No attention should be paid to the structure of the decision tree.

The underlying assumption here is a special case of the *normal form invariance* hypothesis for games that was propounded by von Neumann and Morgenstern (1944, 1953). They took the view that it did not matter whether a game was played in extensive form, with players choosing moves at a succession of information sets, or in normal form, with players choosing a single strategic plan of what to do at each of their information sets, and then committing to that plan by asking the “umpire” of the game to execute it on their behalf. Applied to decision trees, which in effect are single-person games, their invariance hypothesis states that it is irrelevant whether the agent makes a sequence of moves at successive decision nodes, or instead commits in advance to one strategy specifying what to do at each decision node.

A similar idea in decision theory is that all that matters about an act is its consequence — or, more generally, what probability distribution over different consequences emerges in each state of the world. After all, Savage (1954) defines an act as a mapping from states of the world to consequences. A few years earlier, Arrow (1951) had written of valuing actions by their consequences. Somewhat later, Elizabeth Anscombe (1958) chose the phrase “consequentialism” as a pejorative label for a doctrine in moral philosophy she wished to attack — namely, that acts should be judged by their consequences. This doctrine has antecedents in the works of Aristotle, and is rather clearly enunciated by Mill, Moore and others. It was vehemently denied by St. Thomas Aquinas, however.

In 1974 I was unaware of much of this philosophical background. But while in Australia, working on dynamic models of social choice, it did occur to me that an important variation of Arrow’s path independence condition, which was mentioned in Section 4, might be to require that the consequences of behaviour depend only on the feasible set of consequences. In particular, they should not depend on the structure of the decision tree. This idea is explored in Hammond (1977), which addressed the issue of why social choice should maximize a preference ordering (and also why Arrow’s condition of independence of irrelevant alternatives should be satisfied by an Arrow social welfare function).

Expressed in a way which I now prefer, the 1977 paper was effectively based on three postulates. Of these, the first is an unrestricted domain condition, requiring behaviour to be well defined at every decision node of every finite decision tree with consequences in a given arbitrary domain Y . The second postulate requires the “continuation subtree” that emerges from any decision node to be regarded as a decision tree in its own right; moreover, behaviour should be the same at any decision node of that subtree as in the same node of the full tree. This postulate I call “dynamic consistency,” though others use this term differently. It is satisfied even by the naive potential addict of Section 5, who at node n_1 chooses the addiction outcome a over the best outcome b . This behaviour can be regarded as the same in the subtree at n_1 that arises after addiction has set in, as well as in the whole tree. Unlike the sophisticated potential addict, however, the naive potential addict fails to take this future behaviour into account at an earlier stage. There is an inconsistency between the naive agent’s plans and behaviour, but no inconsistency in behaviour *per se*. Dynamic consistency, in fact, is almost a tautology. And it can even be

dispensed with, at little cost, by defining behaviour at any decision node to be behaviour in the subtree starting at that particular node.

That brings me to the third and most important postulate, which I now call “consequentialism.” This is like von Neumann and Morgenstern’s concept of normal form invariance. It requires that there should be a consequence choice function C defined on the domain of non-empty (and finite) subsets of Y with the property that, whenever a decision tree T confronts the agent with the feasible set of consequences F , behaviour in the tree T should limit the range of resultant consequences to the *choice set* $C(F) \subset F$. In particular, whenever two trees T and T' have the same feasible set F , behaviour in those trees should give rise to the same consequence choice set $C(F)$. It is this condition that is violated by both naive and sophisticated choice in the potential addict example. To see this, note that in a different decision tree allowing the potential addict to choose b in one stage — perhaps by means of an infallible commitment device — rational behaviour would lead to the *ex ante* best consequence b instead of to a or c . The same consequentialist condition is violated in the second example of Section 5. This illustrates the difficulty with the Pareto criterion or any other incomplete preference relation.

7 Consequentialism Implies Ordinality

Recall that the three “consequentialist” conditions of Section 6 are unrestricted domain, dynamic consistency, and consequentialism itself. Together, these three conditions imply that behaviour in different feasible sets must generate consequences that reveal a (complete and transitive) preference ordering over the consequence domain Y . In other words, “consequentialist” rationality implies “ordinal” rationality. In Hammond (1977), there is a proof based on Arrow’s (1959) characterization of ordinal choice functions. But in more recent work I have preferred to use simple direct proofs of this fact. Here is an outline of the argument.

Consider any non-empty finite feasible set $F \subset Y$. Now, there certainly exists a one-stage decision tree $T(F)$ in which the only decision node is the initial node, and the agent’s behaviour gives rise immediately to just one of the consequences in F . It follows that for every non-empty finite $F \subset Y$, the choice set $C(F)$ is non-empty. Next, suppose that a belongs to the choice set $C(F)$. Then a must be a possible consequence of behaviour in $T(F)$. But also, given any other consequence $b \in F$, there is a different decision tree T with two stages, as illustrated in Figure 3. The first stage of the tree T occurs at initial node n_0 , where the agent chooses either a consequence $y \in F \setminus \{a, b\}$, or else goes on to a second stage. This occurs at decision node n_1 , where the choice is between the consequences a and b .

In tree T , consequentialism requires option a to be one possible result of the agent’s behaviour. Therefore, that behaviour must include the possibility of moving from node n_0 to n_1 . Also, a must be a possible consequence of




Figure 3: A two-stage decision tree T to demonstrate ordinality

behaviour in the subtree $T(n_1)$ beginning at node n_1 , in which the feasible set of consequences is $\{a, b\}$. Therefore $a \in C(\{a, b\})$.

On the other hand, if $a \in C(F)$ and $b \in C(\{a, b\})$, then b is also a possible consequence of behaviour in the subtree $T(n_1)$. It follows that the move from n_1 to b is possible behaviour in $T(n_1)$ and so, by dynamic consistency, in T as well. But it has been seen already that the agent may move from node n_0 to n_1 in tree T . Hence, b is a possible consequence of behaviour in the whole tree T , with feasible set F . That is, $b \in C(F)$.

Let us now define the revealed preference relation R on Y by $x R y$ iff $x \in C(\{x, y\})$. Because $C(\{x, y\})$ is non-empty for every pair set $\{x, y\}$, it follows that the relation R is complete. Now, in the previous two paragraphs it was proved that $a R b$ whenever both $a \in C(F)$ and $b \in F$, and also that $b \in C(F)$ whenever $a \in C(F)$, $b \in F$, and $b R a$. It is then very easy to show that, for any finite set $F \subset Y$, one has

$$C(F) = \{a \in F \mid b \in F \implies a R b\}$$

So behaviour must reveal a choice function C that maximizes R over each finite feasible set $F \subset Y$.

Finally, to show that R is transitive, consider what happens when $F = \{a, b, c\}$ and $a R b, b R c$. The three possibilities are:

1. $a \in C(F)$, in which case $a R c$ because $c \in F$;
2. $b \in C(F)$, in which case $a \in C(F)$ because $a R b$, so case 1 applies;
3. $c \in C(F)$, in which case $b \in C(F)$ because $b R c$, so case 2 applies.

Thus, $a R c$ in every case, so R is indeed transitive.

Equally important is the fact that the ordinality property gives a complete characterization of consequentialist behaviour. That is, given any preference ordering on Y , there exists a behavioural rule defined for all finite decision trees which satisfies dynamic consistency and consequentialism, while revealing that particular preference ordering. This can be proved by a suitable backward

induction argument, like that involved in applying the principle of optimality in dynamic programming.¹ So consequentialism, just like ordinality, fails to limit behaviour beyond the existence of some preference ordering. Behaviour which maximizes a very unreasonable preference ordering will still be consequentialist, even though it seems clearly irrational.

8 Consequentialism and Objectively Expected Utility

As remarked already, the three consequentialist axioms of Section 7 imply nothing more than the existence of a preference ordering in riskless decision trees. But suppose that risk is introduced through chance nodes at which a move occurs according to a specified or “objective” probability law. The consequences of possible behaviour become simple probability distributions in the space $\Delta(Y)$ of “lotteries” λ with the property that $\lambda(y) > 0$ only for y in a finite support. That is, consequences become risky.

The classical model of decision-making under risk requires behaviour to result in risky consequences $\lambda \in \Delta(Y)$ which maximize the expected value $\mathbb{E}_\lambda v := \sum_{y \in Y} \lambda(y) v(y)$ w.r.t. λ of what has come to be called a *von Neumann–Morgenstern utility function* (or NMUF) $v : Y \rightarrow \mathfrak{R}$. Or, to be more exact, a unique cardinal equivalence class of such functions, with all possible ratios $\frac{v(a) - v(c)}{v(b) - v(c)}$ ($a, b, c \in Y$) of non-zero utility differences uniquely defined and equal to constant marginal rates of substitution between shifts in probability from c to a and shifts in probability from c to b . Nevertheless, in practice there appear to be many systematic deviations from expected utility maximizing behaviour, such as those which Allais (1953, 1979a, b, 1987) and many successors have noticed in various experiments. In response to these widely observed empirical regularities, much attention has recently been paid to non-expected utility models of decision-making under risk.

As I have pointed out elsewhere (Hammond, 1988b), non-expected utility maximizers are liable to be essentially inconsistent in the same way as the potential addict of Section 5. Indeed, there is even a risk of abandoning a project that was initially judged to be beneficial. To see this, consider the decision tree illustrated in Figure 4, with four consequence lotteries $\lambda, \mu, \nu, \rho \in \Delta(Y)$. At the initial decision node n_0 , the agent is confronted with the choice between obtaining consequence ρ for sure, or of going on to the chance node n_1 . There nature selects consequence ν with probability $1 - \alpha$, but with probability α the agent is taken to decision node n_2 where the choice is between λ and μ .

Let v be any NMUF defined on Y . Then for any $0 < \alpha \leq 1$, the respective expected utilities of lotteries λ, μ, ν must satisfy

$$\alpha \mathbb{E}_\lambda v + (1 - \alpha) \mathbb{E}_\nu v \geq \alpha \mathbb{E}_\mu v + (1 - \alpha) \mathbb{E}_\nu v \iff \mathbb{E}_\lambda v \geq \mathbb{E}_\mu v$$

¹For details, see my forthcoming chapter “Objective Expected Utility: A Consequentialist Perspective” to appear in the *Handbook of Utility Theory* (Kluwer Academic Publishers).




Figure 4: A decision tree to demonstrate essential inconsistency

But the expressions in the left hand inequality are the expected utilities of the compound lotteries $\alpha \lambda + (1 - \alpha) \nu$ and $\alpha \mu + (1 - \alpha) \nu$ respectively. So the expected utility model implies the *independence axiom* (I) requiring that

$$\alpha \lambda + (1 - \alpha) \nu R \alpha \mu + (1 - \alpha) \nu \iff \lambda R \mu$$

As is well known, any continuous preference ordering on $\Delta(Y)$ satisfying (I) can be represented by the expected value of each NMUF in a unique cardinal equivalence class. So any non-expected utility maximizing agent with continuous preferences must violate axiom (I).

Now, one way of violating the axiom (I) strictly would be to have

$$\alpha \mu + (1 - \alpha) \nu P \alpha \lambda + (1 - \alpha) \nu \quad \text{but} \quad \lambda P \mu$$

Then, given continuous preferences on $\Delta(Y)$, there must exist $\rho \in \Delta(Y)$ for which

$$\alpha \mu + (1 - \alpha) \nu P \rho P \alpha \lambda + (1 - \alpha) \nu \quad \text{and} \quad \lambda P \mu$$

Yet these preferences almost reproduce those of the potential addict example. At node n_0 , the best option is $\alpha \mu + (1 - \alpha) \nu$. So a naive agent is likely to choose node n_1 , intending to continue on to node μ if given the opportunity. But if node n_2 is reached, the agent prefers λ to μ and so chooses λ . At node n_0 , the anticipated outcome of this choice is the lottery $\alpha \lambda + (1 - \alpha) \nu$, which is worse than either of the two alternative lotteries $\alpha \mu + (1 - \alpha) \nu$ and ρ faced by the agent when at n_0 . Like the sophisticated agent in the potential addict example, the agent would do better to understand that μ is going to be made unavailable by the agent's own future behaviour. For this reason, ρ should be chosen instead of facing the risk of choosing λ over μ and so achieving $\alpha \lambda + (1 - \alpha) \nu$. The agent would do better still, of course, by finding some precommitment device that collapses the decision tree and allows $\alpha \mu + (1 - \alpha) \nu$ to be chosen without any chance of replacing μ by λ .

Any such essential inconsistency will be avoided if and only if prescribed behaviour at n_0 in the subtree beginning at n_1 is the same as in the subtree beginning at n_2 . But this requires precisely that axiom (I) be satisfied.

In this extended framework, what are the implications of the three “consequentialist” axioms introduced in Section 7 — namely, unrestricted domain,

dynamic consistency, and consequentialism itself? Somewhat surprisingly, they lead to the absurd conclusion that there should be universal indifference, with behaviour at each decision node left entirely arbitrary. Obviously, as before the axioms imply the existence of a preference ordering. The trouble is that axiom (I) has to be satisfied by this ordering even when $\alpha = 0$. So, for all $\lambda, \mu, \nu \in \Delta(Y)$, one has $\nu R \nu \iff \lambda R \mu$. Since the left hand side of this logical equivalence is a tautology, the right hand side must be true for all pairs $\lambda, \mu \in \Delta(Y)$, so there is indeed universal indifference.

There are three possible escapes from this ridiculous impasse. The first is to disregard those parts of a decision tree that can be reached only with probability zero by pruning them off. The unrestricted domain of finite decision trees is replaced by an “almost” unrestricted domain, excluding trees having any zero probability move at some chance node. This restriction implies that axiom (I) is only required to hold when $0 < \alpha < 1$. Pruning off zero probability branches in this way is eminently sensible in single-person decision theory. But in multi-person games it leads to problems because it excludes consideration of what would happen in the zero probability event that one player were to stray off a presumed equilibrium path. In other words, it excludes all of the important considerations that were captured so elegantly in Selten’s (1965) concept of subgame perfection.

A second escape is to weaken dynamic consistency and require it to hold only in subtrees that are reached with positive probability. This is “almost” dynamic consistency and allows $\alpha \lambda + (1 - \alpha)\nu I \alpha \mu + (1 - \alpha)\nu$ to be true when $\alpha = 0$ even though $\lambda P \mu$ or $\mu P \lambda$.

Elsewhere (Hammond 1994, 1997) I have begun to explore a third escape. This involves arbitrarily small or “infinitesimal” probabilities which are positive yet less than any positive real number — see also McLennan (1989) and Blume, Brandenburger and Dekel (1991a, b). Any such probability can be regarded as a suitable vanishing sequence of positive probabilities, in effect.

Given any of these three escapes, the consequentialist axioms imply that behaviour should have consequences which maximize a preference ordering R on the space of consequence lotteries. Moreover, the ordering R should satisfy the independence axiom (I). This is a complete characterization of consequentialist behaviour because, given any such ordering on $\Delta(Y)$ that satisfies independence, there exists behaviour satisfying the three consequentialist axioms which generates R as the revealed preference ordering.

In particular, consequentialism does not imply expected utility maximization because the preference ordering R could still embody discontinuities. To rule these out, it is enough to invoke a requirement that, as probabilities at chance nodes vary but the rest of the decision tree remains fixed, including the consequences at its terminal nodes, so the correspondence (or multi-valued mapping) from probabilities at chance nodes to behaviour at decision nodes must have a relatively closed graph. (Note that the graph cannot be closed because zero probabilities are excluded.) With this fourth axiom of continuous behaviour, one can prove that behaviour must have consequences in the form of lotteries

that maximize expected utility. Moreover, there is a unique cardinal equivalence class of von Neumann–Morgenstern utility functions.

9 Consequentialism and Unordered Events

Not all lack of certainty can be described by means of objective probabilities. There can be uncertainty about unknown states of nature, instead of about risky probabilistic prospects or lotteries. The deservedly famous decision theory of Savage (1954) in particular sought to deal with this issue. He laid out seven important postulates that are sufficient to determine uniquely both an agent’s subjective probabilities and also a cardinal equivalence class of von Neumann–Morgenstern utility functions.

Some of Savage’s most important postulates are logical implications of the consequentialist axioms, when they are applied to finite decision trees with, instead of chance nodes, “natural nodes” at which nature’s move helps determine the “state of the world” or “state of nature”. Specifically, suppose there is not only a consequence domain Y but also a (finite) domain S of possible states of the world affecting what consequence results from a given act. Then an *event* E takes the form of any non-empty subset $E \subset S$. For each such event E , there must exist a *contingent* preference ordering R^E defined on the Cartesian product space $Y^E := \prod_{s \in E} Y_s$ of *state contingent consequence functions* $y^E = \langle y_s \rangle_{s \in E}$ mapping E into Y . Furthermore, the different conditional orderings must satisfy Savage’s *sure-thing principle* requiring that, whenever $E = E_1 \cup E_2$ where E_1 and E_2 are disjoint, and whenever $a^{E_1}, b^{E_1} \in Y^{E_1}$, $c^{E_2} \in Y^{E_2}$, then

$$(a^{E_1}, c^{E_2}) R^E (b^{E_1}, c^{E_2}) \iff a^{E_1} R^{E_1} b^{E_1}$$

This is very like the independence axiom for lotteries with objective probabilities, and it can be proved in essentially the same way that it is a logical implication of the three consequentialist axioms.

Another of Savage’s postulates is that preferences are *state independent*. This means that there must exist an ordering R^* on Y with the property that, whenever $s \in S$ and $a, b \in Y$, then $a R^{\{s\}} b \iff a R^* b$. In other words, $R^{\{s\}} = R^*$ for all $s \in S$. This is not an implication of consequentialism, strictly speaking. However, it makes sense to regard a decision tree in which only one state of the world is possible as effectively identical to a corresponding decision tree without any uncertainty. After all, the state of the world is known, as also is the consequence of every strategy in the decision tree. If this accepted, it makes sense to postulate that the consequences of behaviour in any such decision tree should depend only on the feasible set of sure and certain consequences in Y , independent of the state of the world. By repeated application of the sure thing principle enunciated above, it is then easy to show that, whenever $E \subset S$ is an event and $a1^E, b1^E$ denote the two constant state contingent consequence functions satisfying $a_s = a \in Y$ and $b_s = b \in Y$ for all $s \in E$, then $a1^E R^E b1^E \iff a R^{\{s'\}} b$ for all $s' \in S$. This fact is useful in the argument of the next paragraph.

Now, one of Savage’s most important postulates is that there be an *ordering of events* from most to least likely. In particular, whenever $E_1, E_2 \subset S$ are disjoint events, it is natural to say that E_1 is *at least as likely* as E_2 provided that, whenever $a P^* b$, then $(a1^{E_1}, b1^{E_2}) R^{E_1 \cup E_2} (b1^{E_1}, a1^{E_2})$. In other words, the agent weakly prefers the relatively favourable outcome a to occur in E_1 and the relatively unfavourable outcome b to occur in E_2 , rather than the other way around. What Savage assumed, following earlier work on “qualitative probability” by Keynes (1921), Ramsey (1926), de Finetti (1937) and others, was that this “at least as likely” relation is well defined. This is clearly necessary for the existence of subjective probabilities, which must attach higher probability to a more likely event. Yet the ordering of events condition is not an implication of consequentialism, even when state independence is imposed and there is an abundance of possible states. In fact, the consequentialist axioms imply no more than the existence of different contingent preference orderings R^E satisfying the sure thing principle.

To confirm this, it is enough to exhibit a family of contingent preference orderings that fails to induce an ordering of events despite satisfying the sure thing principle. To do so, let $S = \{s_1, s_2\}$ and $Y = \{a, b, c\}$. Then define the state independent utility function $v : Y \rightarrow \mathfrak{R}$ so that:

$$v(a) = 1; \quad v(b) = 0; \quad v(c) = -1. \quad (1)$$

Now consider the preference ordering on Y^S induced by the specific additive utility function

$$U^S(y^S) = \phi_1(v(y_{s_1})) + \phi_2(v(y_{s_2})) \quad (2)$$

where ϕ_1 and ϕ_2 are increasing functions satisfying

$$\begin{array}{lll} \phi_1(1) & = & 2, & \phi_1(0) & = & 0, & \phi_1(-1) & = & -1 \\ \phi_2(1) & = & 1, & \phi_2(0) & = & 0, & \phi_2(-1) & = & -2 \end{array} \quad (3)$$

Suppose now that the two contingent orderings on Y_{s_1} and Y_{s_2} are represented by the utility functions $\phi_1(v(y_{s_1}))$ and $\phi_2(v(y_{s_2}))$ respectively. Because (2) has an additive form, the sure thing principle is evidently satisfied. Moreover, the preferences on Y_{s_1} , Y_{s_2} are even state independent, as are those on the set $Y1^S := \{(y_{s_1}, y_{s_2}) \in Y_{s_1} \times Y_{s_2} \mid y_{s_1} = y_{s_2}\}$, since all are resrepresented by the same utility function $v(y)$. Nevertheless

$$\begin{array}{lll} U^S(a, b) & = & 2, & U^S(b, a) & = & 1 \\ U^S(b, c) & = & -2, & U^S(c, b) & = & -1 \end{array}$$

So the agent’s behaviour reveals a preference for winning a in state s_1 to winning it in state s_2 , when the alternative losing outcome is b . On the other hand, it also reveals a preference for winning b in state s_2 to winning it in state s_1 , when the alternative losing outcome is c . Hence, there is no induced ordering of the events $\{s_1\}$ and $\{s_2\}$.

Savage, of course, introduced other postulates whose effect is to ensure a rather rich set of states. Adding such postulates, however, in general will not

induce an ordering of events. To see this, suppose that S is the entire interval $[0, 1]$ of the real line instead of just the doubleton $\{s_1, s_2\}$. Instead of the additive utility function (2), consider the integral

$$\bar{U}^S(y^S) = \int_0^{1/2} \phi_1(v(y(s)))ds + \int_{1/2}^1 \phi_2(v(y(s)))ds$$

with v given by (1) and ϕ_1, ϕ_2 by (3). Also, so that the integrals are well defined, y^S should be a measurable function from S to Y , in the sense that the set $\{s \in S \mid y(s) = y\}$ is measurable for all $y \in Y$. Then the particular CCF $y^S = \left(a 1_{[0, \frac{1}{2}]}, b 1_{(\frac{1}{2}, 1]}\right)$ with

$$y(s) = \begin{cases} a & \text{if } s \in [0, \frac{1}{2}] \\ b & \text{if } s \in (\frac{1}{2}, 1] \end{cases}$$

is preferred to the lottery represented by $\left(b 1_{[0, \frac{1}{2}]}, a 1_{(\frac{1}{2}, 1]}\right)$ in the same notation. But $\left(c 1_{[0, \frac{1}{2}]}, b 1_{(\frac{1}{2}, 1]}\right)$ is preferred to $\left(b 1_{[0, \frac{1}{2}]}, c 1_{(\frac{1}{2}, 1]}\right)$. So there is no induced likelihood ordering of the two events $[0, \frac{1}{2}]$ and $(\frac{1}{2}, 1]$. In fact, it is easy to confirm that this example satisfies Savage's postulates P1–P3 and P5–P7; only the ordering of events postulate P4 is violated.

10 Subjective Expected Utility

The subjective expected utility model can be given a consequentialist justification, however. Doing so requires extending the domain of finite decision trees so that, as well as decision nodes and terminal nodes, one can have both the chance nodes of Section 8, and also the natural nodes of Section 9. Then the consequences of behaviour become lotteries over state contingent consequence functions. On the space of such lotteries the three consequentialist axioms imply the existence of contingent preference orderings satisfying both the independence axiom and the sure thing principle. Even so, and even if one also imposes the continuous behaviour axiom of Section 8 and the state independence condition discussed in Section 9, there still remain several possible deviations from subjective expected utility. For details, see Hammond (1988a).

However, a decision tree which has both chance and natural nodes is almost like a three-person game in extensive form. The players are the agent, of course, but also chance, who moves according to “objectively” specified probability laws, and nature. Applying the idea of consequentialist normal form invariance to such a tree helps justify a key axiom of an alternative approach to subjective probability due to Anscombe and Aumann (1963). This key axiom is called *Reversal of Order*, or (RO). To state this condition, let $\rho = \langle \rho_i \rangle_{i=1}^k$ be any probability distribution or “roulette lottery” over the set of integers $i \in \{1, 2, \dots, k\}$, and let $\lambda_i^S \in \Delta(Y^S)$ ($i = 1, 2, \dots, k$) be any collection of k lotteries, with $\lambda_{is}(y)$ as the marginal probability of obtaining consequence y in

each state $s \in E$. Let $\sum_{i=1}^k \rho_i \lambda_i^S(y^S)$ be the compound lottery in which the roulette lottery ρ is resolved before the “horse lottery” that determines which $s \in E$ occurs. Then condition (RO) requires $\sum_{i=1}^k \rho_i \lambda_i^S(y^S)$ to be indifferent to the alternative compound lottery in which the horse lottery is resolved first, and its outcome $s \in S$ determines which of the list $(\sum_{i=1}^k \rho_i \lambda_{is}(y_s))_{s \in S}$ of marginal roulette lotteries occurs next.




Figure 5: Decision trees T and T' (when $k = 2$ and $E = \{s, \tilde{s}\}$)

This condition can also be given a consequentialist justification by showing how it is implied by a hypothesis I call “consequentialist normal form invariance.” To illustrate this, consider two decision trees T and T' described as follows. Tree T begins with the chance node n_0 , which is succeeded by the set of natural nodes $N_{+1}(n_0) = \{n_i \mid i = 1, 2, \dots, k\}$. The transition probabilities are $\pi(n_i|n_0) = \rho_i$ ($i = 1, 2, \dots, k$). Each n_i is succeeded by the set of terminal nodes $N_{+1}(n_i) = \{x_{is} \mid s \in E\}$. The consequences are assumed to be given by $\gamma(x_{is}) = \lambda_{is} \in \Delta(Y)$ ($s \in E$). Tree T is illustrated in the left half of Figure 5 for the case when $k = 2$ and $E = \{s, \tilde{s}\}$.

On the other hand, tree T' begins with the natural node n'_0 , whose successors form the set $N'_{+1}(n'_0) = \{n'_s \mid s \in E\}$. Then each n'_s is a chance node whose successors form the set $N'_{+1}(n'_s) = \{x'_{is} \mid i = 1, 2, \dots, k\}$ of terminal nodes. The transition probabilities are $\pi'(x'_{is}|n'_s) = \rho_i$ ($i = 1, 2, \dots, k$). The consequences are assumed to be given by $\gamma'(x'_{is}) = \lambda_{is} \in \Delta(Y)$ ($s \in E$). Tree T' is illustrated in the right half of Figure 5, again for the case when $k = 2$ and $E = \{s, \tilde{s}\}$.

Both trees represent a three-person extensive game between chance, nature, and the agent, who actually has no decision to make. In tree T it is natural to assume that $N_{+1}(n_0)$ is a single information set for nature. Similarly, in tree T' it is natural to assume that $N'_{+1}(n'_0)$ is a single information set for chance. Then the extensive form games represented by the two trees will have identical normal forms, in which the agent has only one strategy, whereas chance’s strategies are indexed by $i \in \{1, 2, \dots, k\}$ and nature’s strategies are indexed by $s \in E$. Furthermore, in either extensive form game, when chance chooses i and nature chooses s , the consequence is the lottery $\lambda_{is} \in \Delta(Y)$. In fact, the *consequentialist normal form invariance* condition is that trees like T and T' with identical three-person normal forms should be regarded as giving rise to equivalent feasible sets, and that behaviour should generate equivalent choice sets of consequences in each case. This is a natural extension of the normal form invariance hypothesis mentioned previously.

Evidently, in tree T the sole feasible consequence available to the decision maker is $\sum_{i=1}^k \rho_i \lambda_i^E(y^E)$, whereas in tree T' it is $\langle \sum_{i=1}^k \rho_i \lambda_{is}(y_s) \rangle_{s \in E}$. Consequentialist normal form invariance requires these to be regarded as equivalent. But this is precisely condition (RO).

With this reversal of order condition added to the three consequentialist axioms, plus state independence and continuity, it follows that behaviour must maximize the expected value of each NMUF in a unique cardinal equivalence class, with subjective probabilities applied to the different possible moves by nature. Moreover, all these subjective probabilities must be positive. In particular, no event can be “null” according to the criterion specified by Savage. Furthermore, as the agent moves through the decision tree, these subjective probabilities must be revised according to Bayes’ Rule. This is a strengthening of the standard subjective expected utility (or SEU) model; accordingly, it will be called the SEU* model.

As in previous sections, it can also be shown that the SEU* model is the only implication of these axioms. Nothing yet rules out absurd utility functions. Or, in the case of subjective probabilities, absurd beliefs reflected in these probabilities. We are still not past rationality as structural consistency. But at least consequentialism, with its focus on what is reasonable behaviour in non-trivial decision trees, seems to offer a stronger defence of the standard structural consistency criterion and the associated SEU model.

That concludes the discussion of normative single-person decision theory in this paper. It is time to move on to two branches of multi-person decision theory. The first is social choice theory, which is concerned with finding a common objective for a group of heterogeneous individuals. And the second is game theory, which studies how different individuals interact in their behaviour.

11 Social Choice Theory

Most irrational behaviour by ordinary people has little effect outside a small circle of family, friends, workmates, etc. Mistakes by professionals such as doctors can be somewhat more devastating, and lead to unnecessary loss of property, life or limb. So can reckless driving by any of us. Of course, Akerlof and Yellen (1985a, b) argued that even small mistakes, when made by many people simultaneously, can have surprisingly large adverse consequences. Nevertheless, the irrational behaviour that has the largest effects on the most people comes from political leaders, public officials, etc. Even in contemporary Europe, their mistakes and/or criminal behaviour can still destroy nation states. It seems obvious that we should expect those granted substantial power over their colleagues or fellow citizens to meet higher standards of rationality.

Now, in my view social choice theory should be about specifying suitable objectives for public officials and others responsible for major decisions affecting large numbers of individuals. In particular, it should specify objectives for economic policy. Accordingly, the structural consistency conditions discussed in previous sections seem even more appropriate for social choice than they are

for individual behaviour. This suggests that the structure of the social decision tree should not affect how the consequences of public decisions emerge from the set of feasible consequences. So the consequentialist axioms apply, and imply the existence of a social ordering. When there is risk, it is appropriate to maximize the expected value of a cardinal von Neumann–Morgenstern social welfare function (or NMSWF). And to attach positive subjective probabilities to states of nature which do not have specified objective probabilities, so reaching the SEU* model.

However, the NMSWF W and its expected value $\mathbb{E}W$ should reflect the interests of the individuals who are affected by the decisions under consideration. In principle, this could even be the entire world population, together with future generations yet to be born. In order to model these interests, it helps to adapt a device used by Foley (1970) and Milleron (1972) to treat public goods, and imagine for a moment that each individual could have their own separate world to live in.

Let I denote the set of all individuals. For each $i \in I$, the choice of world for individual i , taking account of only i 's personal interests, is a degenerate social choice problem. Applying the consequentialist axioms and their relatives to the class of all degenerate decision trees involved in choosing a world for i , we are able to infer that the interests of each individual i should be represented by the expected value of their own von Neumann–Morgenstern individual welfare or utility function w^i . That is, if we were choosing a world for i alone, it would be appropriate to maximize $\mathbb{E}w^i$.

Now consider any social decision tree which happens to be one individual i 's personal decision tree in the sense that, although one is choosing a separate world for each individual, in fact any decision that is made affects only i 's world, leaving all other individuals' worlds completely unaffected. For such decisions, it seems reasonable to insist that the social objective $\mathbb{E}W$ should collapse to the individual objective $\mathbb{E}w^i$, in the sense of representing the same social preference ordering. But then it is easy to prove that, whenever two lotteries over possible worlds for all individuals are equally attractive for all those individuals, society must be indifferent between them. From these arguments, the strict Pareto condition follows. That is, if all individuals are no worse off, then so is society; also, if any individual is better off, then so is society unless some other individual is worse off. In particular, if all individuals are equally well off, then so is society; this is the Pareto indifference condition.

Given these Pareto properties, and the fact that each individual's risky world can be chosen separately and independently, it follows that the social welfare function is some positively weighted sum of all individuals' welfare functions. In fact, this was first proved by Harsanyi (1955). By suitable normalization of the cardinal individual welfare functions, these weights can all be made equal to 1. Also, where decisions affect the set of individuals who come into existence, one can presume that individuals who will never come into existence have no interest in what world is created for them, and so have constant utility in all such worlds of permanent non-existence. Then we can normalize individual utilities so that

this constant utility is zero. Formally, then, the social welfare function takes the classical utilitarian form, being equal to the sum of all individual utilities.

This objective is for an artificial problem where all individuals can have their own separate and independent risky worlds. But the natural constraint that we must all share the same world, and so that all individuals' risks must be perfectly correlated, is merely a restriction on what decision trees are relevant. The form of the objective is not affected.

The social welfare objective specified above is formally the same as for classical utilitarianism. But the "utilities" being added have quite a different interpretation. Expected individual utility is defined in a way that makes it an appropriate objective in decision trees affecting only one individual. And it can be shown rather easily that each utility ratio amounts to a constant marginal rate of substitution between the numbers of individuals in two particular situations. In particular, it is better to change society in order to have more individuals with higher utility, and fewer with lower utility. It is also better to have more individuals with positive utility, but this should be interpreted carefully because the zero level of utility is defined so that society is better off with more individuals who have positive utility, but worse off with more individuals who have negative utility. For more details, see Hammond (1991, 1996).

Once again, by thinking carefully about what is appropriate behaviour in decision trees, some surprisingly strong conclusions have emerged.

12 Some Special Two-Person Games

Even two-person games are much more complicated than single-person decision trees. Rational behaviour for each player usually depends on what behaviour is rational for the other player. The only exception appears to be in games like Prisoner's Dilemma, or finite repetitions of it, in which each player has a unique strategy that strictly dominates all others — i.e., is uniquely best regardless of what strategy the other player chooses. Yet even for Prisoner's Dilemma, some people have argued that, at least in special circumstances, it is rational to play the dominated strategy of cooperating instead of the dominant strategy of defecting. For example, suppose that Prisoner's Dilemma is going to be played between two computer programs. Then Howard (1988) showed how to write a program which checks to see if its opponent is an exact copy of itself. Such a program could then go on to cooperate if and only if it is playing against such a copy.

Another class of relatively simple two-person games are those in which there is perfect and complete information, and the game is bounded in length. An example of such a game is chess, and indeed the first modern mathematical work on game theory was Zermelo's (1912) pioneering analysis of chess. He proved that, with best play by both White and Black, chess has a definite result — a win for White, a draw, or a win for Black. Which of these is right depends on a finite calculation, but one which is so complicated that it may not be possible ever to determine what is the result of a perfectly played game of chess. Clearly,

normative models of behaviour lose their practical significance when the decision problem is too complicated ever to be analysed fully.

The procedure for resolving chess, in theory but not in practice, is backward induction. This works backwards from terminal positions in which the game is already over because there is checkmate, stalemate, a draw by repetition or perpetual check, or a draw because each player has made 50 moves without capturing a piece or moving a pawn. Then the backward induction procedure evaluates all positions where one player can force checkmate in one move, or else can only avoid being checkmated by moving immediately to a terminal drawn position. The procedure goes on to evaluate positions which, given best play by both sides, require at most two moves to reach a terminal position, then at most three, etc. In principle, all possible positions get evaluated in this way, including the starting position.

The rules of chess give the two players strictly opposing objectives, so the normative relevance of backward induction is uncontroversial. Indeed, it is applied in practice in order to evaluate simple chess endgames. But backward induction is controversial in other important games like the ‘‘Centipede’’ considered by Rosenthal (1981) and many later writers. This game is illustrated in Figure 6 (see also Osborne and Rubinstein, 1994).




Figure 6: Rosenthal’s Centipede Game, Modified

There are two players labelled P_1 and P_2 with respective strategy sets

$$S_1 = \{ a_i \mid i = 0, 1, \dots, n \}, \quad S_2 = \{ b_j \mid j = 0, 1, \dots, n \}$$

If P_1 chooses $a_i \in S_1$, then unless P_2 ends the game beforehand, P_1 moves across in the tree exactly i successive times before moving down. Similarly, if P_2 chooses $b_j \in S_2$, then unless P_1 ends the game beforehand, P_2 moves across exactly j successive times before moving down. The game ends immediately if either player moves down. If $i = j = n$, it ends anyway after P_2 has chosen b_n . Outside this case, if the two players choose (a_i, b_j) where $i = k \leq j$, then the game ends after they have both moved across k times before P_1 moves down. But if $i > k = j$, then P_1 moves across $k + 1$ times but P_2 only k times before moving down to end the game. Hence, the two players’ respective payoffs are

$$v_1(a_i, b_j) = \begin{cases} i + 1 & \text{if } i \leq j \\ j & \text{if } i > j \end{cases} \quad \text{and} \quad v_2(a_i, b_j) = \begin{cases} i & \text{if } i \leq j \\ j + 2 & \text{if } i > j \end{cases}$$

It is now easy to see that whenever P_2 is given the opportunity to choose b_n , it would be unreasonable to do so because b_{n-1} is available and yields a higher payoff. But once b_n has been eliminated, whenever P_1 is given the opportunity

to choose a_n , it would be unreasonable to do so because a_{n-1} is available and yields a higher payoff. Then, for $k = n - 1, \dots, 1$, once a_n, \dots, a_{k+1} have all been eliminated, each successive b_k is worse than the alternative b_{k-1} . Similarly, for $k = n - 1, \dots, 1$, once b_n, \dots, b_k have all been eliminated, each successive a_k is worse than the alternative a_{k-1} . Thus, the backward induction procedure successively deletes all strategies until the only ones remaining are b_0 for P_2 and a_0 for P_1 . The strategy profile (a_0, b_0) is one Nash equilibrium, of course, though there are others as well. The tree collapses to the single branch a_0 . Backward induction suggests that this is the only part of the tree which needs to be analysed.

Starting with Rosenthal (1981) and Binmore (1987), several game theorists have found this backward induction argument to be unconvincing, for the following reason. Suppose P_1 were unexpectedly faced with the opportunity to play a_k after all, because neither player has yet played down, and in fact each player has played across k times already. Backward induction applied to the remaining subtree leads to the conclusion that P_2 , if given the move, will play b_k next time, so P_1 should play a_k . Yet P_2 has already played across k times, whereas backward induction implies that P_2 should move across whenever there is a move to make. So, as Binmore in particular argues most persuasively, if k is large enough, P_1 has every reason to doubt whether the backward induction argument applies to P_2 's behaviour after all. Furthermore, if $n - k$ is also large, there may be much to gain, and at most 1 unit of payoff to lose, from allowing the game to continue by moving across instead of playing a_k .

Of course, P_2 can then apply a similar argument when faced with the choice between b_k and continuing the game. Also, P_1 should understand how moving across once more instead of playing a_k will reinforce P_2 's doubt about whether the backward induction argument applies to P_1 , and so make it more likely that P_2 will decline to play b_k . This strengthens P_1 's reasons for not playing a_k . Similar reasoning then suggests that P_2 should not play b_{k-1} , that P_1 should not play a_{k-1} , etc. In the end, it may be sensible for P_1 not to play a_0 , for P_2 not to play b_0 , etc. Indeed, there are some obvious close parallels between this argument and that of Kreps, Milgrom, Roberts, and Wilson (1982) for players first to cooperate and then play tit-for-tat in the early stages of a finitely repeated Prisoner's Dilemma. Or for a chain-store to play "tough" in the paradox due to Selten (1978).

13 Perfect Recall

One major rationality assumption in orthodox game theory is perfect recall. This requires that, in an extensive form game tree, all the branches passing through a particular information set of any player must pass through an identical sequence of that player's previous information sets; in addition, at each preceding information set of that player, there must be exactly one strategy allowing the later information set to be reached. The interpretation is that players remember all the information that was available previously, and also whatever

moves they chose at their own previous information sets. Because agents with finer information partitions are able to make better decisions, it seems obvious that players in a game should strive for perfect recall as a rather obvious normative standard.

Actually, for some years, I thought that the assumption of perfect recall could be made without loss of generality. After all, a player who forgets something can be treated as a different player. So there seems no harm in treating each information set as belonging to a different “local” player, and replacing each original player by a team of local players who share a common objective but may not be able to share their information. For example, it seems not to matter much whether one regards the game of bridge like von Neumann and Morgenstern did, as a game between two players, North–South and East–West, each with imperfect recall. Or instead as a four person game, with North, South, East and West treated as separate players who form two opposing teams.

Piccione and Rubinstein (1994) point to an interesting counter example, however. When a player has imperfect recall, the same information set can be encountered more than once along some branches of the game tree. In this case, imperfect recall plays an essential role, and even in a single person game it is possible for a mixed strategy to be better than any pure strategy. However, this is a rather exceptional situation in which the player loses count of how many times a particular information set has been visited. Now, players who can keep count of how often they visit each of their information sets must actually have information sets which can only be visited at most once. For such players, it *is* possible to treat imperfect recall by introducing enough local players.

14 Nash Equilibrium

Two-person games of perfect and complete information like chess or the Centipede can be “solved” by backward induction. They are not typical, however. Generally, players must move simultaneously, or in ignorance of the other’s previous moves. They would like to know these moves so that they could react appropriately, but they do not and cannot know.

An important class of games are those like chess with two players who have strictly opposing objectives. What one wins the other loses. Hence the name “zero-sum,” though this terminology does not fit in well with modern ordinal or even cardinal concepts of utility which cannot be added across people. For such games, Borel (1924) suggested that mixed or randomized strategies could be important, as when a poker player chooses at random whether to bluff or not. A little later, von Neumann (1928) proved that allowing two players in a zero-sum game to choose mixed strategies would give each player a determinate *maximin* expected payoff. This can be found by maximizing with respect to his own strategy the payoff that is minimized with respect to the opponent’s strategy.

Meanwhile, Morgenstern (1935) had realized that economic forecasting was difficult or impossible unless one could somehow determine agents’ expectations

concerning the future, including their expectations of other agents' behaviour, which would depend on these other agents' expectations, etc. The two-person zero sum game theory which von Neumann had developed must have seemed at the time a way of resolving the resulting infinite regress. So they got together to write von Neumann and Morgenstern (1944, 1953) — one of the most important and influential scientific books of all time. However, only their analysis of two-person zero-sum games is really satisfactory; indeed, the last two-thirds of their book (apart from the general concept of a stable solution set, and the appendix on expected utility) receives little attention these days, and is probably only of historical interest.

The next major step in game theory is due to Nash (1950, 1951) who was able to show that any n -person game would have an equilibrium in mixed strategies, with each player's mixed strategy maximizing their own expected utility given all the other players' mixed strategies. The concept of Nash equilibrium is one of the most widely used in economics. Yet it is not without its conceptual difficulties, as Luce and Raiffa (1957) in particular pointed out rather early on.

If a game has a unique Nash equilibrium, both players might be expected to find and play it. If two players 1 and 2 have sets of equilibrium strategies E_1 and E_2 with the property that any pair $(s_1, s_2) \in E_1 \times E_2$ is a Nash equilibrium, then the equilibria are said to be *exchangeable*. In this case, any equilibrium $s_1 \in E_1$ is a best response to any equilibrium $s_2 \in E_2$, and vice versa. In particular, whenever player 2 chooses any $s_2 \in E_2$, all equilibrium strategies $s_1 \in E_1$ give player 1 exactly the same expected payoff, and vice versa. So each player i should be content to play an arbitrary strategy $s_i \in E_i$, and thereby reach one of the multiple equilibria. But there remain many games like Battle of the Sexes with non-exchangeable multiple Nash equilibria, even in pure strategies. That is, the equilibrium set might well contain at least two strategy pairs $(s_1, s_2), (s'_1, s'_2)$ for which s_1 must be a best response to s_2 , by definition, but may not be a best response to s'_2 . Then, to be sure of reaching a Nash equilibrium, player 1 needs to know which equilibrium strategy player 2 is likely to choose, but cannot. Similarly for player 2.

15 Rationalizability

This problem of non-exchangeable strategies went unresolved for about 25 years, until Bernheim (1984) and Pearce (1984) included independent work on rationalizable strategies in their respective Ph.D. theses, and had papers published in *Econometrica*. One of the key steps in the Bernheim/Pearce revolution was their replacement of the objective probabilities that are specified in mixed strategies by subjective probabilities which each player attaches to the other's strategies. Previously, Harsanyi (1967–8) had discussed how players might attach subjective probabilities to each others' "types" in a game of incomplete information. Also, related ideas can be found in the somewhat obscure work of Armbruster and Böge (1979) and of Böge and Eisele (1979). But it was Bernheim and

Pearce, followed shortly by Tan and Werlang (1988), who developed a useable new game theory.

In fact, rationalizable strategies are those that are optimal given rationalizable beliefs in the form of subjective probability distributions attached to other players' strategies. And rationalizable beliefs are those attaching probability 1 to the event that all other players will choose some rationalizable strategies. Bernheim and Pearce showed how one could apply this obviously circular definition rather easily. Indeed, rationalizable strategies in two-person games could be characterized as all those that survive iterative deletion of strictly dominated strategies for both players. There is a similar characterization for n -person games provided that rationalizable beliefs allow other players' strategies to be correlated whenever necessary.

Though any strategy appearing in Nash equilibrium is certainly rationalizable, there can be many other rationalizable strategies besides. So, in an era when many game theorists were applying ideas due to Selten (1965) or Kreps and Wilson (1982) in order to refine the set of Nash equilibria, it is perhaps surprising that the Bernheim/Pearce coarsening of the equilibrium set received as much attention as it did. This may be because of the persuasiveness of their approach, which simply and conclusively answers the question of what minimal restrictions standard decision theory imposes on possible behaviour in a normal form game.

Nevertheless, as Mariotti (1996) points out, there does remain a rather serious flaw in the arguments used by Bernheim and Pearce, as well as by Tan and Werlang and others. After all, they assume that players have beliefs about each others' strategies that are represented by subjective probabilities, and that players take decisions according to the SEU model described in Section 10.² In Sections 6–10, these assumptions were justified by axioms concerning behaviour in an almost unrestricted domain of decision trees. In fact subjective probabilities relate to the willingness of the agent to trade off good consequences against bad in different states of the world — in other words, to make particular bets. Instead of variations in single-person decision trees, one might assume directly that players have preferences over consequences. But then such preferences really have meaning only when players are faced with choices between pairs or other subsets of consequences. In other words, in order to apply the axioms of decision theory to games, and so give preferences decision-theoretic content, one must allow variations in the set of strategies and their consequences which are available to each player.

Now, in single-person decision theory, it makes sense to assume that nature determines the uncertain state of the world unsystematically or haphazardly. Then it also makes sense to assume that nature's choice is unaffected by variations in the agent's decision tree or other representation of the decision problem.

²In his work on correlated equilibrium, Aumann (1987) is more careful. He postulates a space of extrinsic states, each of which is assumed to determine a unique strategy profile that is suggested to all players before the game starts. Equilibrium occurs when individuals have identical beliefs about the probabilities of all the different extrinsic states, and when for each player it is optimal given these beliefs to follow the suggested strategy.

But in games, all players are assumed to choose strategies in their own best interests. Then variations in one player's strategy set affect the whole game in ways which are likely to affect what strategies it is in these other players' interests to choose. For this reason, as Mariotti argues, it is far from evident that one can legitimately treat other players' purposefully chosen strategies in the same way as moves determined haphazardly by nature. Above all, it is not obvious that the SEU criterion can be applied to bets contingent on other players' strategies in the same way as to bets contingent on states of nature. Nor is it obvious that, even if players are willing to accept bets in such a way, they should treat the game itself as effectively a similar bet to which the SEU criterion can be applied.

To overcome this serious difficulty, Battigalli (1996) suggests that instead of having the players themselves accept bets on strategy choices in the game, one could consider a clone of each player whose behaviour in identical decision problems will be exactly the same. This clone is regarded as entirely outside the game, in the sense that the clone's decisions have no effect whatsoever on the consequences of the strategies chosen by any of the players participating in the game. Then the clone is in a position to treat the game in the same way as a small gambler who is not a jockey treats a horse race. That is, the clone can regard the strategy choices of all the players in the game as states of nature, with subjective probabilities attached. Also, the clone can be faced with a decision problem that is effectively a copy of that faced in the game itself by the corresponding original player. Because the clone will maximize subjective expected utility in this copy of the game, the original player will do the same in the original game.

In this way, by applying the axioms presented earlier in this essay in a suitable context, one *can* justify the claim that each player should indeed use the SEU model, just as most game theorists have always asserted. The same axioms also justify the standard specification of a game in normal form, with "payoff" functions that are really NMUFs. There is the obvious qualification that each player's payoff function or NMUF is only determined up to a unique cardinal equivalence class. Clearly, this is unimportant. Note, however, that the framework used here differs from that of Börgers (1993), whose assumptions do not allow decision problems with objective probabilities to be considered, and so yield only a unique ordinal equivalence class of utility functions.

Another requirement is that each player i 's SEU model include a unique specification of appropriate subjective probabilities for other players' strategy profiles in the game, as well as for any exogenous states of nature which may be relevant. Failure to specify these probabilities leaves the description of the players' decision models fundamentally incomplete. Yet specifying them arbitrarily ignores the fact that, in the end, other players' strategies are really not like states of nature, because other players face their own individual decision problems in the game, which they try to resolve rationally, at least to some extent. This tension is precisely the challenge that non-cooperative game theory must meet.

A more serious worry is that, as explained in Section 10, consequentialism actually implies the SEU* model, in which each player i 's subjective probability attached to *any* profile of other players' strategies must be strictly positive. This contradicts much of orthodox game theory, where each player is required to attach zero probability to the event that one or more other players choose strategies which are not their best responses. In particular, the probability of any player choosing a strictly dominated strategy must be zero.

As mentioned in Section 8, one remedy might be to weaken dynamic consistency to almost dynamic consistency. In extensive form games, however, this approach forces separate consideration of every "subgame of incomplete information" that a solution is presumed to reach only with zero probability. This effectively violates the normal form invariance hypothesis, as enunciated by von Neumann and Morgenstern (1944, 1953), which claims that it is enough to analyse the normal or strategic form of any game. This invariance hypothesis, moreover, especially applied in its "consequentialist" form, is really what justifies the subjective expected utility approach to decision making, which represents the key assumption of modern game theory. The only satisfactory remedy appears to be using infinitesimal probabilities of the kind mentioned in Section 8.

16 Restrictions on Beliefs

The SEU* model, even within a general multi-person game, actually imposes very few restrictions on agents' beliefs. After all, it is a structural consistency model of rational behaviour. About the only way in which an agent who uses the SEU* model properly could ever be inconsistent is to reach some unforeseen situation. So the only advice it seems to offer is to take every possibility into account and accord everything a positive small probability. Be broad-minded, in other words.

Excessive broad-mindedness seems clearly irrational, however. One should not believe the Earth is flat. One should believe the sun will rise tomorrow (unless you are in either the Arctic or the Antarctic during winter). One should expect most people who are young enough to learn from their obvious mistakes most of the time. Earthquakes will occur in California, avalanches in the Alps, accidents on the roads, and volcanic eruptions on Mount Vesuvius or Etna. These are just instances of physical or psychological reality which we ignore at our peril. The question of what it is reasonable to believe in such instances represents the hard philosophical problem of induction. It has exercised many great minds over several millenia without reaching any very good resolution that I am aware of.

Nevertheless, several distinguished game theorists such as Harsanyi and Aumann have argued that players would have common probabilistic beliefs if they pooled all their information. Equivalently, it as though all players started with a common prior distribution which then became updated based on information acquired by each player just before the game starts. In the case when this common prior also concerns what strategies the different players in a game will

choose, they will then reach a Nash equilibrium — or, if correlation is allowed, a correlated equilibrium (see especially Bernheim, 1986). Similar ideas underline the concept of rational expectations equilibrium that was made famous by Muth and Lucas.

Where a game is played only once, it seems obvious that players will not have enough time to learn what other players are likely to do. After all, induction can only apply to situations that are repeated sufficiently often. Thus, the hypothesis of rational expectations seems implausible except in games that repeated many times, or where players in a “one-shot” game have reason to believe that they all play strategies appropriate to a particular “focal” equilibrium (Schelling, 1960). In fact, though I am certainly willing to admit that players can and should learn from their past experiences if they are in a repeated game, I have never been persuaded that most of the games economists should be interested in really are repeated in the relevant technical sense, nor that one should limit attention to “rational” expectations or common prior beliefs. So the idea that agents should make use of all available information is a good one, but too much is likely to remain hidden to allow equilibrium to be reached in all but trivial games which, for instance, happen to have unique rationalizable strategies.

Of course, abandoning equilibrium does much to weaken the predictive power of game theory. After all, if a game has a unique equilibrium, that can be used to determine each player’s strategy and also each player’s expectations of the other players’ strategies. The shift from equilibrium to rationalizability typically expands the set of strategies and of associated expectations which rationalize those strategies. Then one may be forced to look for considerations beyond the scope of traditional game theory in order to help predict players’ expectations and their strategy choices.

Actually, even in traditional game theory, considerations that are excluded from the traditional description of a game may be relevant to the choice of equilibrium. An important example is forward induction, as discussed by Kohlberg and Mertens (1986). This works most clearly when the first player to move has the chance to end the game immediately with a well defined consequence or outside option, or alternatively of entering a subgame to be played with another player. Typically, this subgame, like Battle of the Sexes, has multiple non-exchangeable equilibria. Then the claim is that, if the first player chooses to enter this subgame, it can only be in the expectation of achieving an outcome that is no worse than the outside option. In several interesting games, this is enough to restrict the relevant equilibrium set in the subgame to a single outcome. For further discussion, see van Damme (1989), Osborne (1990), Hillas (1990), Hammond (1993), and especially Gul (1996).

Notice however that, if forward induction can be used to refine the equilibrium set within some subgame, perhaps it could also be used for an entire game. After all, one or other players may have given up past outside options to enter the game that is about to be played. So, in addition to the usual description of a game in normal or extensive form, in some cases it is also important to know

what happened before the game started. Traditional game theory ignores all such considerations.

17 Beyond Structural Consistency

The previous section was concerned with possible restrictions on beliefs that transcend the usual SEU* model, which recommends only that everything which might be relevant should be taken into account. It is also reasonable to impose restrictions on preferences. Up to a point, more of the good things that life has to offer should be preferred to less. Excessive risk of severely adverse outcomes like death or mutilation should clearly be avoided. People who have any funds in savings accounts should use them to pay off high interest credit card balances as rapidly as possible. Many people should probably buy fewer lottery tickets than they do, unless serious financial distress is so imminent that the only hope of escape is to win a large prize. Schoolchildren should be taught how to handle financial affairs, including how to manage pension accounts, mortgages, etc. Not only schoolchildren, in fact. Those who are about to retire and go on a pension often need advice from a source other than agents seeking a fat commission.

These rather vague thoughts illustrate the point that, even when choosing on their own behalf, individuals may well benefit from guidance in establishing what their preferences should be, as well as their beliefs about the likely consequences of different possible actions. Full rationality is much more than mere structural consistency. This is especially true if one believes that full rationality also requires suitable attention to ethical values and respect for other people, as well as animals, etc. Such considerations, and the judgement that some preferences are rational while others are not, goes beyond traditional economics and into the realm of philosophy. At this point, economists should recognize their limitations and admit that they know little about many fundamentally important questions concerning the rationality of different preferences. So, counting myself as an economist instead of a philosopher, I shall pass on to the next topic.

18 Procedural Rationality versus Satisficing

The SEU* model recommends taking all possibilities into account, examining every available decision for its likely consequences, and then maximizing. The implicit methodology regards problems like playing chess perfectly as relatively trivial. Such unbounded rationality is clearly far too exacting.

A little over 40 years ago, Simon (1955, 1957) realized this and went on to introduce an alternative concept which he called *procedural rationality*. Trying to maximize subjective expected utility, even when it is feasible, is often unnecessarily burdensome. Far better to “satisfice” by finding an action that meets an appropriate “aspiration level”.

Part of Simon’s theory, even its main message, makes eminent sense. The unbounded rationality of the SEU model does need replacing with some less

onerous procedure for making “boundedly” rational decisions. However, the satisficing approach does not seem satisfactory, especially as a normative standard. Satisficing might make sense if the only difficulty were to find the best decision among a large number of possibilities being given due consideration. Instead, it seems to me that a much more common problem arises because it is very hard to analyse more than a very small number of options in any detail, whereas it is rather easy in principle to rank those which have been analysed thoroughly and so find the best from this restricted set. In other words, agents who face complicated decision or game trees are forced to include only very few possibilities within any conceptual model they use to analyse the decision they have to make. In other words, they bound their models.

Now, a bounded model is likely to be worth revising later on, when the agent is closer to the time when an important decision has to be made, and history has rendered irrelevant many of the contingencies which had been possible *ex ante*. Of course changes in the agent’s model of the decision tree may affect the consequences of behaviour. But if the agent never sees any need to revise his bounded model, there is no reason why its structure should influence behaviour and its consequences. And even if the bounded model is likely to change, for behaviour within whatever bounded model the agent chooses to construct, all the standard axioms retain their normative appeal. In particular, they recommend that, within the bounded model, the agent should choose the best possible action. This seems quite different from Simon’s recommendation that the agent satisfice in choosing an action.

So, following similar ideas by Behn and Vaupel (1982), Vaupel (1986), and others, in my view a better normative theory of bounded rationality recognizes that the agent should satisfice in the construction of a bounded model. But within that model the theory recommends the agent to choose the best possible action. This seems quite different from Simon’s recommendation that the agent satisfice in choosing an action.

19 Stochastic Descriptive Models

So far, I have been almost exclusively concerned with normative standards of rational behaviour. But as the discussion of bounded rationality makes clear, the SEU model pays little or no attention to psychological reality. Of course, psychological reality seems more relevant for describing actual behaviour which is generally not rational than it is for thinking about normative standards of rational behaviour. Nevertheless, normative standards that pay no attention at all to psychological reality can easily be dismissed as irrelevant. This is especially so in game theory, where one can imagine offering advice to one player on how to behave rationally when some or all of the other players are not necessarily rational. Such advice should be based on a model which describes realistically the behaviour of all but the one player who is receiving the advice.

Constructing a realistic descriptive model of behaviour is perhaps more of a task for psychologists than for economists. And I know even less about psy-

chology than I do about philosophy. Nevertheless, it seems that only stochastic models offer any serious chance of describing behaviour at all well. Moreover, the following slight variation of a well known example due to McFadden (1974) illustrates how, in any realistic stochastic model of behaviour, the structure of the decision tree is very likely to be important.

Suppose somebody is contemplating going by public transport from Milan to, say, Como. Suppose too that this person can go either to the railway station and catch a train, or else to the bus station and catch one of two buses — red or blue. Now, if the railway and bus stations are very close, the traveller may well have time to examine the timetables for the trains and for both red and blue buses before making a single decision determining how to travel to Como. In this case, lacking further information, we might think that there is a $1/3$ chance of the agent taking the train, or either kind of bus, depending on when the agent reached the station, the respective fares, which train or bus is scheduled to reach Como first, etc.

Suppose on the other hand that the railway and bus stations are sufficiently far apart that the traveller will certainly depart from whichever station he goes to first. Then the decision is made in two stages. First the traveller decides whether to head for the railway or bus station. At the railway station there is no further decision to make, but at the bus station there is the second decision between the red and blue buses. In this case the probability of the traveller going to the railway station and taking a train may well rise to $1/2$, whereas with probability $1/2$ he goes over to the bus station. If the probabilities of catching either kind of bus remain equal, both drop to $1/4$. For more details, see Amemiya (1981, 1985).

Such behaviour has consequences which depend on the structure of the decision tree. So the consequentialist axiom is violated. The example brings out important differences between, on the one hand, normative models of rational behaviour and, on the other hand, descriptively accurate models that may not be rational at all.

20 Confusing Normative and Descriptive

Most theoretical work in economics, even when it is concerned with description or prediction rather than prescription, is actually based on the standard normative model. That is, it postulates preference maximizing consumers, and profit maximizing firms. Often consumers are assumed to maximize expected utility and firms to maximize expected profit, though there has recently been extensive interest in models where consumers maximize non-expected utility.

This confusion of normative and descriptive seems to me a fundamental methodological error. I shall not try to defend it, though I will venture four positive explanations. Of these, the first is rather obvious: inertia or laziness. Demand theory in economics has always been based on utility maximization, though utility was first refined to ordinal utility and then virtually dispensed with through the revealed preference approach. Nevertheless, the “neoclassical”

utility maximization approach to the theory of consumer demand gets taught to succeeding generations of students, and so passed on.

A second possible explanation is mathematical sophistication and clarity. The remarkable intellectual revolution in economics which occurred during the 1950s owed much to the use of mathematical techniques. It will be hard for anything as simple and messy as most stochastic choice models to supplant the familiar maximization models of modern neoclassical theory.

A third possible explanation is one I owe to Joan Robinson, and her book *Economic Philosophy*. It is the suggestion that many economists have been little more than apologists for the *laissez faire* free market system. This is perhaps closer to the truth in the U.S.A., where many economists teach at business schools, and even more receive direct or indirect financial support from business people. Other economists outside business schools may nevertheless consult, teach people who are planning business careers, teach at universities that rely on business profits to endow their chairs, or simply teach at universities where students have their high tuition fees paid by parents whose incomes are derived from business. Given all these financial pressures, it is perhaps somewhat surprising that there are nevertheless many critics of the free enterprise system who flourish in U.S. universities.

Anyway, whatever their motivation, whenever defenders of free enterprise care to think at all about the performance of the economic system as a whole, they love to invoke the fundamental efficiency theorems of welfare economics. Under important qualifications which are often conveniently forgotten, these theorems demonstrate that free markets are associated with Pareto efficient outcomes in which it is impossible to make any one consumer better off without making some others worse off. Yet even when they are valid, these efficiency theorems are totally devoid of normative content unless one postulates that consumers choose what is right for themselves, by rationally maximizing an appropriate preference ordering (or ordinal utility function). Another favourite result concerns the (potential) gains from free trade, enhanced competition, privatization, or other forms of economic liberalization. Yet this result also relies on the same kind of maximization postulate. Defenders of free enterprise would have the force of their arguments considerably reduced if consumers were known to be behaving irrationally. Yet then economists might have a new role to play advising consumer organizations instead of business. Perhaps, however, they see that business can afford to pay better.

A fourth reason for using the model of rational behaviour descriptively could be that it has been tested experimentally. This would be the best reason, of course, if it were valid. However, the experimental evidence to date seems to be rather mixed. Where experiments face subjects with relatively simple decision problems with clear monetary objectives, and also give subjects an opportunity to learn from experience, there is some evidence that behaviour converges to what would be predicted in the (often unique) equilibrium of a game theoretic model. This seems to be especially true of the experimental markets studied by Plott and others — see Sunder's (1995) survey, for instance. Outside some rather special laboratory situations, however, there seems to be

little evidence of rationality in actual behavior. This is true even for special classes of experimental subjects like economics or business students, who are supposed to understand something of what it means to be rational.

21 Problems Solved and Unsolved

I feel that my own understanding of what constitutes rationality in economics has progressed far since I first became interested in the topic about 30 years ago. Then, I remember being somewhat horrified by the first economics textbook I read, which happened to be Samuelson's *Economics*, because of the way in which consumer behaviour was reduced to analysing an indifference map involving only that one consumer's personal consumption.

The theory of rational behavior itself had made huge conceptual advances in the period from the mid 1940s, when von Neumann and Morgenstern published their book, until 1963, when Anscombe and Aumann's important work appeared. However, the standard axioms were often heavily criticized, even as a basis for a normative theory. So the challenge had become how to investigate the foundations more carefully, by considering what might lie behind the standard axioms.

With this in mind, the three consequentialist axioms and their ancillaries, applied to behaviour in single person decision trees, do appear to offer a more secure foundation for standard normative decision theory. This is especially true of the SEU model in the form proposed by Anscombe and Aumann (1963). But there are problems — especially with the unrestricted domain assumption when the agent may have reason to favour some decision trees over others because they allow choices to be deferred and so offer more flexibility. Nevertheless, this consequentialist approach has been flattered by the attention of several good humoured and sympathetic critics such as Machina (1989), McClennen (1990) and Munier (1996). It has been an interesting challenge of late, because of the significant developments in non-expected utility models of choice under risk and in alternatives to subjective probability. But on the whole the non-standard models purport to be descriptive, whereas I claim that the SEU model is suitable as a normative standard only. Saying that often enough takes much of the sting out of the criticisms.

For me at least, the exploration of normative single-person decision theory seems almost over. There may be interesting technical problems involved in extending the theory to infinite decision trees, and/or those with an infinite horizon. But these are mathematical niceties. More pressing issues arise in dealing with games. Even for normal form games, the work on justifying the SEU* approach remains too novel for me to be confident yet of its ultimate acceptability. But for games in extensive form there are more serious issues. For instance, there is recent unpublished work by Asheim and Dufwenberg (1996) claiming that players may have more than one set of rationalizable strategies. One should also reconcile the need for positive probabilities with the desire to

attach zero probabilities to strategies that are not rationalizable. Serious open questions like these remain incompletely resolved.

Last but not least is the urgent need to develop better concepts of rationality than the structural consistency conditions which currently dominate the economics literature. There is certainly a need for alternative normative models that allow us to say more about which beliefs and preferences are rational. And for normative models which make fewer psychological demands on agents, as well as recognizing psychological reality better than current models seem to. Although I do maintain that the fundamental separation between normative and descriptive models of behaviour remains necessary.

Acknowledgements

November 1997 final version of material prepared for a lecture presented at the Università Cattolica del Sacro Cuore, Milan, on 28th March 1996. Several parts are based on draft chapters to appear in Salvador Barberà, Peter J. Hammond and Christian Seidl (eds.) *Handbook of Utility Theory* (in preparation for Kluwer Academic Publishers). Section 15 in particular has benefited from the comments of Marco Mariotti on a draft chapter for the *Handbook*.

References

- Afriat, S.N. (1967) "The Construction of a Utility Function from Demand Data" *International Economic Review*, **8**, 67–77.
- Akerlof, G.A. and J.L. Yellen (1985a) "Can Small Deviations from Rationality Make Significant Differences to Economic Behavior?" *American Economic Review*, **75**, 708–720.
- Akerlof, G.A. and J.L. Yellen (1985a) "A Near-Rational Model of the Business Cycle, with Wage and Price Inertia" *Quarterly Journal of Economics*, **100**, 823–838.
- Allais, M. (1953) "Le comportement de l'homme rationnel devant le risque: critique des postulats et des axiomes de l'école américaine," *Econometrica*, **21**, 503–546; summary version of paper whose translation appeared as Allais (1979a).
- Allais, M. (1979a) "The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School," in Allais and Hagen (1979), pp. 27–145.
- Allais, M. (1979b) "The So-called Allais Paradox and Rational Decisions under Uncertainty," in Allais and Hagen (1979), pp. 437–681.
- Allais, M. (1987) "Allais Paradox," in Eatwell, J., Milgate, M. and P. Newman (eds.) *The New Palgrave: A Dictionary of Economics* (London: Macmillan, 1987); reprinted in Eatwell, J., Milgate, M. and P. Newman (eds.) *The New Palgrave: Utility and Probability* (London: Macmillan, 1990), pp. 3–9.
- Allais, M. and O. Hagen (eds.) (1979) *The Expected Utility Hypothesis and the Allais Paradox* (Dordrecht: D. Reidel).

- Amemiya, T. (1981) "Qualitative Response Models: A Survey" *Journal of Economic Literature*, **19**, 1483–1536.
- Amemiya, T. (1985) *Advanced Econometrics* (Oxford: Basil Blackwell).
- Anscombe, F.J. and R.J. Aumann (1963) "A Definition of Subjective Probability" *Annals of Mathematical Statistics*, **34**, 199–205.
- Anscombe, G.E.M. (1958) "Modern Moral Philosophy" *Philosophy*, **33**, 1–19.
- Armbruster, W. and W. Böge (1979) "Bayesian Game Theory," in O. Moeschlin and D. Pallaschke (eds.) *Game Theory and Related Topics* (Amsterdam: North-Holland), pp. 17–28.
- Arrow, K.J. (1951) "Alternative Approaches to the Theory of Decision in Risk-Taking Situations," *Econometrica*, **19**, 404–437; reprinted in Arrow (1971), ch. 1 and in Arrow (1984), ch. 2.
- Arrow, K.J. (1959) "Rational Choice Functions and Orderings" *Economica*, **26**, 121–127; reprinted in Arrow (1984), ch. 6.
- Arrow, K.J. (1963) *Social Choice and Individual Values* (2nd. edn.) (New Haven: Yale University Press).
- Arrow, K.J. (1971) *Essays in the Theory of Risk-Bearing* (Chicago: Markham; Amsterdam: North-Holland).
- Arrow, K.J. (1984) *Collected Papers of Kenneth J. Arrow, Vol. 3: Individual Choice under Certainty and Uncertainty* (Cambridge, Mass.: Belknap Press of Harvard University Press).
- Arrow, K.J., E. Colombatto, M. Perlman, and C. Schmidt (eds.) (1996) *The Rational Foundations of Economic Behaviour* (London: Macmillan).
- Asheim, G. and M. Dufwenberg (1996) "Admissibility and Common Knowledge," preprint, Department of Economics, University of Oslo.
- Aumann, R.J. (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, **55**, 1–18.
- Bandyopadhyay, T. (1978) "Revealed Preference Theory, Ordering and the Axiom of Sequential Path Independence" *Review of Economic Studies*, **55**, 343–351.
- Battigalli, P. (1996) "Comment," on Mariotti (1996), in Arrow *et al.* (eds.) (1996), pp. 149–154.
- Behn, R.D. and J. W. Vaupel (1982) *Quick Analysis for Busy Decision Makers* (New York: Basic Books).
- Bernheim, B.D. (1984) "Rationalizable Strategic Behavior" *Econometrica*, **52**, 1007–1028.
- Bernheim, B.D. (1986) "Axiomatic Characterizations of Rational Choice in Strategic Environments," *Scandinavian Journal of Economics*, **88**, 473–488.
- Bernoulli, D. (1738) "Specimen theoriae novae de mensura sortis" in *Commentarii Academiae Scientiarum Imperialis Petropolitanae*; translated by Sommer (1954).

- Binmore, K. (1987) “Modelling Rational Players, I,” *Economics and Philosophy*, **3**, 179–214; reprinted (1990) in *Essays on the Foundations of Game Theory* (Cambridge, Mass.: Basil Blackwell), ch. 5, pp. 151–185.
- Blume, L., A. Brandenburger and E. Dekel (1991a) “Lexicographic Probabilities and Choice Under Uncertainty” *Econometrica*, **59**, 61–79.
- Blume, L., A. Brandenburger and E. Dekel (1991b) “Lexicographic Probabilities and Equilibrium Refinements” *Econometrica*, **59**, 81–98.
- Böge, W. and T. Eisele (1979) “On Solutions of Bayesian Games,” *International Journal of Game Theory*, **8**, 193–215.
- Borel, E. (1924) “Sur les jeux où interviennent l’hasard et l’habileté des joueurs” in J. Hermann (ed.) *Eléments de la théorie des probabilités* (Paris: Librairie Scientifique), 204–224.
- Börgers, T. (1993) “Pure Strategy Dominance,” *Econometrica*, **61**, 423–430.
- Campbell, D.E. (1978) “Realization of Choice Functions” *Econometrica*, **46**, 171–180.
- Cramer, G. (1728) Letter to Nicholas Bernoulli; extracts printed in translation by Sommer (1954).
- Davidson, D., J.C.C. McKinsey and P. Suppes (1955) “Outlines of a Formal Theory of Value, I” *Philosophy of Science*, **22**, 140–160.
- De Finetti, B. (1937) “La prévision: ses lois logiques, ses sources subjectives,” *Annales de l’Institut Henri Poincaré*, **7**; translated as “Foresight: Its Logical Laws, Its Subjective Sources,” in Kyburg and Smokler (1964), pp. 93–158.
- Elster, J. (1979) *Ulysses and the Sirens* (Cambridge: Cambridge University Press).
- Foley, D.K. (1970) “Lindahl’s Solution and the Core of an Economy with Public Goods” *Econometrica*, **38**, 66–72.
- Gul, F. (1996) “Rationality and Coherent Theories of Strategic Behavior” *Journal of Economic Theory*, **70**, 1–31.
- Hammond, P.J. (1976) “Changing Tastes and Coherent Dynamic Choice,” *Review of Economic Studies*, **43**, 159–173.
- Hammond, P.J. (1977) “Dynamic Restrictions on Metastatic Choice,” *Economica*, **44**, 337–350.
- Hammond, P.J. (1988a) “Consequentialist Foundations for Expected Utility,” *Theory and Decision*, **25**, 25–78.
- Hammond, P.J. (1988b) “Orderly Decision Theory: A Comment on Professor Seidenfeld,” *Economics and Philosophy*, **4**, 292–297.
- Hammond, P.J. (1991) “Interpersonal Comparisons of Utility: Why and How They Are and Should Be Made,” in J. Elster and J.E. Roemer (eds.), *Interpersonal Comparisons of Well-Being* (Cambridge: Cambridge University Press), ch. 7, pp. 200–254.

- Hammond, P.J. (1993) "Aspects of Rationalizable Behavior," in Binmore, K., Kirman, A.P. and P. Tani (eds.) *Frontiers of Game Theory* (Cambridge, Mass.: M.I.T Press), ch. 14, pp. 277–305.
- Hammond, P.J. (1994) "Elementary Non-Archimedean Representations of Probability for Decision Theory and Games," in Humphreys, P. (ed.) *Patrick Suppes: Scientific Philosopher, Vol. I: Probability and Probabilistic Causality* (Dordrecht: Kluwer Academic Publishers), ch. 2, pp. 25–59.
- Hammond, P.J. (1996) "Consequentialist Decision Theory and Utilitarian Ethics," in F. Farina, F. Hahn, and S. Vannucci (eds.) *Ethics, Rationality, and Economic Behaviour* (Oxford: Clarendon Press, 1996), pp. 92–118.
- Hammond, P.J. (1997) "Consequentialism, Non-Archimedean Probabilities, and Lexicographic Expected Utility," to appear in C. Bicchieri, R. Jeffrey and B. Skyrms (eds.) *The Logic of Strategy* (Oxford University Press).
- Harsanyi, J.C. (1955) "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy*, **63**, 309–321.
- Harsanyi, J.C. (1967–8) "Games with Incomplete Information Played by 'Bayesian' Players, I–III," *Management Science*, **14**, pp. 159–182, 320–334, 486–502.
- Hicks, J.R. and R.G.D. Allen (1934) "A Reconsideration of the Theory of Value, I and II" *Economica*, **1**, 52–76 and 196–219.
- Hillas, J. (1990) "On the Definition of the Strategic Stability of Equilibria" *Econometrica*, **58**, 1365–1390.
- Houthakker, H.S. (1950) "Revealed Preference and the Utility Function" *Economica*, **17**, 159–174.
- Howard, J. (1988) "Cooperation in the Prisoner's Dilemma" *Theory and Decision*, **24**, 203–213.
- Keynes, J.M. (1921) *A Treatise on Probability* (London: Macmillan).
- Kohlberg, E. and J.-F. Mertens (1986) "On the Strategic Stability of Equilibria," *Econometrica*, **54**, 1003–1037.
- Kreps, D., Milgrom, P., Roberts, J. and R. Wilson (1982) "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, **27**, 245–252.
- Kreps, D. and R. Wilson (1982) "Sequential Equilibrium," *Econometrica*, **50**, 863–894.
- Kyburg, H.E. and H.E. Smokler (1964) *Studies in Subjective Probability* (New York: John Wiley).
- Luce, R.D. and H. Raiffa (1957) *Games and Decisions: Introduction and Critical Survey* (New York: John Wiley).
- Machina, M.J. (1989) "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty," *Journal of Economic Literature*, **28**, 1622–1668.

- Malinvaud, E. (1972), *Lectures on Microeconomic Theory* (Amsterdam: North Holland).
- Mariotti, M. (1996) “The Decision Theoretic Foundations of Game Theory,” in Arrow *et al.* (eds.) (1996), ch. 6, pp. 133–148.
- McClennen, E.F. (1990) *Rationality and Dynamic Choice: Foundational Explorations* (Cambridge: Cambridge University Press).
- McFadden, D. (1974) “Conditional Logit Analysis of Qualitative Choice Behavior,” in P. Zarembka (ed.) *Frontiers in Econometrics* (New York: Academic Press), pp. 105–142.
- McLennan, A. (1989) “Consistent Conditional Systems in Noncooperative Game Theory” *International Journal of Game Theory*, **18**, 141–174.
- Milleron, J.-C. (1972) “Theory of Value with Public Goods: A Survey Article” *Journal of Economic Theory*, **5**, 419–477.
- Mongin, P. (1994) “L’optimisation est-elle un critère de rationalité individuelle?” *Dialogue*, **33**, 191–222.
- Morgenstern, O. (1935) “Vollkommene Voraussicht und wirtschaftliches Gleichgewicht” *Zeitschrift für Nationalökonomie*, **6**, pp. 337–357; translated by F.H. Knight as “Perfect Foresight and Economic Equilibrium” in A. Schotter (ed.) *Selected Writings of Oskar Morgenstern* (New York: New York University Press, 1976), ch. 10, pp. 169–183.
- Munier, B. (1996) “Comment,” in Arrow *et al.* (eds.) (1996), pp. 43–47.
- Myerson, R.B. (1978) “Refinements of the Nash Equilibrium Concept,” *International Journal of Game Theory*, **7**, 73–80.
- Nash, J.F. (1950) “Equilibrium Points in n -Person Games,” *Proceedings of the National Academy of Sciences of the USA*, **36**, 48–49.
- Nash, J.F. (1951) “Non-Cooperative Games,” *Annals of Mathematics*, **54**, 286–295.
- Nau, R.F. and K.F. McCardle (1990) “Coherent Behavior in Noncooperative Games” *Journal of Economic Theory*, **50**, 424–444.
- Osborne, M.J. (1990) “Signaling, Forward Induction, and Stability in Finite Repeated Games” *Journal of Economic Theory*, **50**, 22–36.
- Osborne, M.J. and A. Rubinstein (1994) *A Course in Game Theory* (Cambridge, Mass.: MIT Press).
- Pearce, D. (1984) “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, **52**, 1029–1050.
- Piccione, M. and A. Rubinstein (1994) “On the Interpretation of Decision Problems with Imperfect Recall” University of Tel Aviv, Sackler Institute for Economic Studies, Working Paper No. 24/94.
- Plott, C.R. (1973) “Path Independence, Rationality and Social Choice” *Econometrica*, **41**, 1075–1091.

- Ramsey, F.P. (1926) "Truth and Probability," in *The Foundations of Mathematics and Other Logical Essays* (London: Kegan Paul, 1931), pp. 156–198; reprinted in Kyburg and Smokler (1964) and in Gärdenfors, P. and N.-E. Sahlin (eds.) (1988) *Decision, Probability, and Utility* (Cambridge: Cambridge University Press).
- Rosenthal, R.W. (1981) "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox," *Journal of Economic Theory*, **25**, 92–100.
- Samuelson, P.A. (1938) "A Note on the Pure Theory of Consumer's Behavior" and "— : An Addendum" *Economica*, **5**, 61–71 and 353–4.
- Savage, L.J. (1954) *The Foundations of Statistics* (New York: John Wiley).
- Schelling, T.C. (1960) *The Strategy of Conflict* (Cambridge, MA: Harvard University Press).
- Selten, R. (1965) "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit" *Zeitschrift für die gesamte Staatswissenschaft*, **121**, 301–324 and 667–689.
- Selten, R. (1978) "The Chain Store Paradox," *Theory and Decision*, **9**, 127–159.
- Simon, H.A. (1955) "A Behavioral Model of Rational Choice" *Quarterly Journal of Economics*, **69**, 99–118.
- Simon, H.A. (1957) *Models of Man* (New York: John Wiley).
- Sommer, L. (1954) "Exposition of a New Theory on the Measurement of Risk," (translation) *Econometrica*, **22**, 23–36.
- Sunder, S. (1995) "Experimental Asset Markets: A Survey" in J.H. Kagel and A.E. Roth (eds.) *Handbook of Experimental Economics* (Princeton: Princeton University Press), ch. 6, pp. 445–500.
- Tan, T.C.-C. and S.R. da C. Werlang (1988) "The Bayesian Foundations of Solution Concepts of Games" *Journal of Economic Theory*, **45**, 370–391.
- Uzawa, H. (1960) "Preference and Rational Choice in the Theory of Consumption" in K.J. Arrow, S. Karlin and P. Suppes (eds.) *Mathematical Methods in the Social Sciences, 1959*, (Stanford: Stanford University Press), ch. 9.
- van Damme, E. (1989) "Stable Equilibria and Forward Induction" *Journal of Economic Theory*, **48**, 476–496.
- Vaupel, J.W. (1986) "Un pensiero analitico per decisori indaffarati" in L. Sacconi (ed.) *La decisione: Razionalità collettiva e strategie nell'amministrazione e nelle organizzazioni* (Milano: Franco Angeli, 1986), ch. 13, pp. 226–238.
- von Neumann, J. (1928) "Zur Theorie der Gesellschaftsspiele," *Mathematische Annalen*, **100**, 295–320; reprinted in A.H. Taub (ed.) *Collected Works of John von Neumann, Vol. VI* (Oxford: Pergamon Press, 1963), pp. 1–26.
- von Neumann, J. and O. Morgenstern (1944; 3rd edn. 1953) *Theory of Games and Economic Behavior* (Princeton: Princeton University Press).
- Zermelo, E. (1912) "Über eine Anwendung der Mengenlehre auf die Theorie der Schachspiels" *Proceedings of the Fifth International Congress of Mathematicians*, Vol. II, 501–504.