Peter J. HAMMOND

# CONSEQUENTIALISM AND THE INDEPENDENCE AXIOM

## 1. INTRODUCTION

St.Thomas Aquinas wrote to the effect that no good consequences could make a bad action good, and that no bad consequences could make a good action bad. John Stuart Mill, following a utilitarian tradition, put forward exactly the contradictory hypothesis that, as a necessary condition for rationality, actions should be judged entirely by their consequences. In ethics, this way of judging actions has come to be known as "consequentialism", following Anscombe's (1958) critical discussion. Some of the history of consequentialism is considered further in Hammond (1986).

To date, consequentialism has usually been applied in ethics. Yet it can be adapted to normative single person decision theory as well, where I believe it offers a strong justification for some standard axioms. Indeed, some years ago, I realized that the normative principle that behaviour should be ordinal - i.e. maximize a complete and transitive preference ordering - could be justified by requiring "dynamic choice" in decision trees to depend only upon outcomes and not upon the structure of the decision tree. This requirement was called "metastatic choice" in Hammond (1977). I have since realized that there is a similar justification of the contentious independence axiom (Samuelson, 1952) which is the basis of expected utility maximization. The justification of both ordinality and independence is so close in spirit to the consequentialist approach to ethics that it is no misuse of terminology to call it "consequentialism" as well, rather than the old and clumsy "metastatic".

In single person decision theory, consequentialism is defined informally to mean that behaviour is explicable merely by its consequences. A little more formally, behaviour will be consequentialist if it corresponds to a "revealed consequence choice function" which, for every finite feasible set of possible consequences, specifies a subset of chosen consequences. Stated this way, consequentialism is clearly



B. R. Munier (ed.), Risk, Decision and Rationality, 503–516.

somewhat restrictive, but the restrictions appear reasonable. They are in accord with Savage (1954), for example, who actually *defines* an act as a mapping from states of the world into consequences, thus making consequentialism a tautology in his framework.

The power of the consequentialist hypothesis is much greater than is immediately apparent, however, once it is applied to an unrestricted domain of "consequential" decision trees, as described in Section 2. For such trees, behaviour is naturally described by the "behavour norms" described in Section 3. Moreover - and this is the key to the whole argument - it is natural to restrict attention to behaviour norms which satisfy an obvious property of dynamic consistency, similar to that which I considered for choice functions in Hammond (1976,1977).

Having completed the task of specifying general consistent behaviour norms for consequential decision trees, Section 4 then gives a formal definition of consequentialism. Section 5 shows how consequentialism implies ordinality - i.e. a consequential norm must specify behaviour which chooses consequences to maximize a (complete and transitive) preference ordering on the space of consequence probability distributions. Thereafter, Section 6 demonstrates how the controversial independence axiom must also be satisfied, and that zero probability nodes have to be excluded from the domain of consequential decision trees. The latter result is novel.

Finally, Section 7 has a concluding discussion concerning the domain of relevant consequences.

## 2. CONSEQUENTIAL DECISION TREES

Decision trees, with their chance, decision and terminal nodes, should be familiar from the lucid introductory lectures of Raiffa (1968). Here, therefore, I merely introduce notation and also adapt the definition so that each terminal node $x$ of a tree $T$ has a well defined consequence $\gamma(x)$ in the space $Y$ of possible consequences. Formally, then, a *consequential decision tree* (or "tree" for short) is a collection :

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

whose respective components are described and interpreted as follows :

(i) $N^*$ is a finite set of *decision nodes*;

(ii) $N^o$ is a finite set, disjoint from $N^*$, of *chance nodes;*

(iii) $X$ is a non-empty finite set, disjoint from $N^* \cup N^o$, of *terminal nodes*;

(iv) $N_{+1} : N \twoheadrightarrow N$ is the *immediate successor* correspondence on the domain $N := N^* \cup N^o \cup X$ of all nodes in $T$, but with the proviso that $N_{+1}(x) = \emptyset \Leftrightarrow x \in X$ ;

(v) $n_o \in N$ is the *initial node*, the unique node in $N$ satisfying $(\forall\, n \in N)\ n_o \notin N_{+1}(n)$ ;

(vi) for each pair $n \in N^o$ and $n' \in N_{+1}(n)$ , $\pi(n'|n)$ is the (nonnegative real) *probability* of reaching node $n'$ conditional on having reached $n$ , or, in other words, the probability of making the transition from $n$ to $n'$ ;

(vii) $\gamma : X \to Y$ is the *consequence mapping* describing what consequence in $Y$ results from each terminal node in $X$ .

In order that $T$ be indeed a tree, one imposes the restriction that the successor sets $N_{+1}(n)$ and $N_{+1}(n')$ must be disjoint whenever $n, n' \in N$ are distinct. This, indeed, allows the set $N$ to be constructed recursively, starting with $n_o$ , then $N_{+1}(n_o)$ , which is disjoint from $\{n_o\}$ , then $N_{+1}(n)$ for each $n \in N_{+1}(n_o)$ , etc. until the terminal nodes are reached (after a finite number of steps, because $N^*$, $N^o$, and $X$ are all finite).

One should also add that the probabilities $\pi(\cdot|\cdot)$ must satisfy the obvious restriction that $\Sigma_{n' \in N_{+1}(n)}\ \pi(n'|n) = 1$ for each $n \in N^o$, so that $\pi(\cdot|n)$ is indeed a probability distribution over $N_{+1}(n)$.

Notice that terminal nodes are *not* identified with consequences. It is important to allow different terminal nodes in the same tree to result in the same consequence.

Let $T(Y)$ denote the set of all possible consequential decision trees giving rise to consequences in the set $Y$ . It consists of all possible (finite) decision trees together with suitable consequences mappings.

Where a decision tree $T$ is a variable, it will be included as an argument in expressions like $N^*(T)$ , $N_{+1}(T,n)$ , $n_o(T)$, etc.

## 3. CONSISTENT BEHAVIOUR NORMS

For such consequential decision trees, a *behaviour norm* is a correspondence $\beta(T,n)$ whose domain consists of pairs $(T,n)$ in which $T \in \mathcal{T}(Y)$ is a consequential decision tree and $n \in N^*(T)$ is a decision node of $T$. The correspondence takes values $\beta(T,n)$ which are non-empty subsets of $N_{+1}(T,n)$, for each pair $(T,n)$ in the domain. The norm $\beta$ is allowed to take multiple values, as usual in decision theory, so that indifference is possible. If $\beta$ were restricted to take single values, the preference orderings considered below would have to be strong orderings, excluding indifference. And if $\beta$ were allowed to take values which were probability distributions over $N_{+1}(T,n)$ it would be a *stochastic behaviour norm* which I plan to discuss elsewhere.

Let $T \in \mathcal{T}(Y)$ be any consequential decision tree, and $\bar{n} \in N$ any node of $T = \langle N^*, N^o, X, N_{+1}(\cdot), n_o, \pi(\cdot|\cdot), \gamma(\cdot)\rangle$. Then there is a *continuation tree* $\bar{T}$ – or "subtree" – starting at $\bar{n}$, with :

$$\bar{T} = \langle \bar{N}^*, \bar{N}^o, \bar{X}, \bar{N}_{+1}(\cdot), \bar{n}_o, \bar{\pi}(\cdot|\cdot), \bar{\gamma}(\cdot)\rangle .$$

This continuation can be constructed naturally as follows. First, define the binary relation $>$ on $N$ so that $n' > n$ – i.e. $n'$ is a *successor* of $n$ – if and only if there exists a chain $n_1, n_2, \ldots, n_k$ in $N$, with $n_1 = n$ and $n_k = n'$, for which $n_{j+1} \in N_{+1}(n_j)$ $(j = 1 \text{ to } k-1)$. Second, define $N(n)$ as the set $\{n' \in N \mid n' > n \text{ or } n' = n\}$ of nodes in $N$ which either succeed or are equal to $n$. Now let :

(i) $\bar{N}^* = N^* \cap N(\bar{n})$ ;

(ii) $\bar{N}^o = N^o \cap N(\bar{n})$ ;

(iii) $\bar{X} = X \cap N(\bar{n})$ ;

(iv) $\bar{N}_{+1} : \bar{N} \twoheadrightarrow \bar{N}$ be the restriction to $\bar{N} = N(\bar{n})$ of the correspondence $N_{+1}(\cdot)$, satisfying $\bar{N}_{+1}(n) = N_{+1}(n)$ $(\subset N(n) \subset \bar{N})$ for each $n \in \bar{N}$ ;

(v) $\bar{n}_o = \bar{n}$ ;

(vi) $\bar{\pi}(\cdot|\cdot)$ be the restriction to $\bar{N}$ of the probabilities $\pi(\cdot|\cdot)$, satisfying $\bar{\pi}(n'|n) = \pi(n'|n)$ for each $n \in \bar{N}^o$ and each $n' \in \bar{N}_{+1}(n)$ $(= N_{+1}(n))$ ;

(vii) $\bar{\gamma} : \bar{X} \to Y$ be the restriction to $\bar{X}$ of the consequence mapping $\gamma$, satisfying $\bar{\gamma}(x) = \gamma(x)$ for all $x \in \bar{X}$.

It is then obvious that $\bar{T}$ is itself a consequential decision tree.

From now on, given any node $n$ of $N(T)$, let $T_n$ denote the continuation of tree $T$ from node $n$.

Given $\bar{n} \in N(T)$, because $T_{\bar{n}}$ is a consequential decision tree, the behaviour norm $\beta$ specifies the behaviour set $\beta(T_{\bar{n}}, n)$ at any $n \in N^*(T_{\bar{n}})$. But $N^*(T_{\bar{n}}) \subset N^*(T)$ and so the behaviour norm $\beta$ also specifies the behaviour set $\beta(T,n)$ at $n$.

Having behaviour sets differ at $n$, depending upon whether the decision tree is regarded as $T$ or as $T_{\bar{n}}$, is clearly inconsistent. Indeed, when $n$ is reached, the agent's decision problem is really described by $T_n$; the decisions which may have been taken before reaching $n$ have become irrelevant. Thus, regardless of what $\beta(T,n)$ may be, the agent's behaviour is really described by $\beta(T_n,n)$. In other words, behaviour is entirely described by what happens only at the initial nodes of decision trees in $\mathcal{T}(Y)$.

This leads to the following requirement. A behaviour norm $\beta$ is *consistent* if, for every $T \in \mathcal{T}(Y)$, $\bar{n} \in N(T)$ and $n \in N^*(T_{\bar{n}})$, one has $\beta(T_{\bar{n}}, n) = \beta(T,n)$.

From now on, only consistent behaviour norms will be considered, and they will be called simply *norms*.

## 4. CONSEQUENTIALIST NORMS

An *act* in a consequential decision tree $T$ is a decision rule or pure behaviour strategy $a : N^* \to N$ which satisfies $a(n) \in N_{+1}(n)$ at every decision node $n \in N^*$.

A *simple probability distribution of consequences* (or, for short, a "risky consequence", or even just "consequence") is a function $\mu$, taking non-negative real values on its domain $Y$, for which there exists a finite *support* $S \subset Y$ such that $\mu(y) > 0 \Leftrightarrow y \in S$, and $\Sigma_{y \in S}\, \mu(y) = 1$. The set of all such simple probability distributions will be denoted by $M_o(Y)$ or just $\tilde{Y}$.

Given any act $a$ in $T$, the probabilities $\xi(a,n)$ $(n \in N)$ of reaching each node in $T$ can be calculated by forward recursion as follows :

(i) $\xi(a,n_o) = 1$ ;

(ii) if $\xi(a,n)$ has already been calculated, and if $n' \in N_{+1}(n)$, then :

$$\xi(a,n') = \begin{cases} \pi(n'|n)\,\xi(a,n) & (\text{if } n \in N^o) \\ \xi(a,n) & (\text{if } n \in N^*,\ n' = a(n)) \\ 0 & (\text{if } n \in N^*,\ n' \neq a(n)) \end{cases}$$

Ultimately, this forward recursion ends as the terminal nodes of $T$ are reached, and one has $\xi(a,x) \geqslant 0$ (all $x \in X$), $\Sigma_{x \in X}\, \xi(a,x) = 1$, so that $\xi(a,x)$ $(x \in X)$ is itself a probability distribution on the set $X$ of terminal nodes of $T$. Given the consequence mapping $\gamma$, the act $a$ then results in the risky consequence $\rho(a) \in \tilde{Y}$ satisfying :

$$\rho(a)(y) = \Sigma_x\{\xi(a,x) \mid \gamma(x) = y\} \quad (\text{all } y \in Y).$$

The risky consequence $\mu \in \tilde{Y}$ is *feasible* in the consequential decision tree $T$ if there exists an act $a$ in $T$ for which $\mu = \rho(a)$ – i.e. $\mu$ is the result of $a$. Write $F(T)$ for the set of all risky consequences which are feasible in $T$ ; it is a non-empty finite subset of $\tilde{Y}$.

The risky consequence $\mu \in \tilde{Y}$ is *revealed chosen* by $\beta$ in $T$ if there exists an act $a$ in $T$ with $a(n) \in \beta(T,n)$ (all $n \in N^*(T)$) for which $\mu = \rho(a)$ – i.e. $\mu$ is the result of an act $a$ which, at each decision node of $T$, selects a member of the behaviour set specified by $\beta$. Write $\Phi_\beta(T)$ for the set of all risky consequences which are revealed chosen by $\beta$ in $T$ ; it is a non-empty finite subset of $F(T)$.

Two trees $T, T' \in \mathcal{T}(Y)$ are *consequentially equivalent* if $F(T) = F(T')$. Consequential equivalence says that only the feasible sets of risky consequences in a tree are relevant; one can ignore the structure of the trees which makes those consequence sets feasible.

The norm $\beta$ specifies *consequentially equivalent behaviour* in the two trees $T, T' \in \mathcal{T}(Y)$ if $\Phi_\beta(T) = \Phi_\beta(T')$.

Now comes the crucial definition of the paper. The norm $\beta$ is *consequentialist* if it is defined at all decision nodes of all decision trees in $\mathcal{T}(Y)$ and specifies consequentially equivalent behaviour in any pair of consequentially equivalent decision trees. Thus does behaviour become explicable merely by its consequences.

Let $\mathcal{F}$ denote the set of non-empty finite subsets of the risky consequence space $\tilde{Y}$. A *consequence choice function* is a mapping $C : \mathcal{F} \to \mathcal{F}$ satisfying $C(F) \subset F$ for all $F \in \mathcal{F}$.

A consequence choice function $C$ is *revealed* by the norm $\beta$ if $\Phi_\beta(T) = C(F(T))$ for all $T \in \mathcal{T}(Y)$. Evidently $\beta$ is consequentialist only if it reveals a unique consequence choice function which will be written as $C_\beta$.

## 5. CONSEQUENTIALISM IMPLIES ORDINALITY

Were consequentialist norms allowed to specify inconsistent behaviour, they would be completely characterized by the existence of a (revealed) consequence choice function. Consistency, however, imposes serious restrictions. The first is now described.

A *preference ordering* $R$ on $\tilde{Y}$ is a complete transitive binary relation. A choice function $C$ on $F$ is *ordinal* if there is a preference ordering $R$ on $\tilde{Y}$ such that, for all $Z \in F$:

$$C(Z) = \{y^* \in Z \mid y \in Z \Rightarrow y^* \mathcal{R}\, y\}$$

i.e. if $C(Z)$ consists of the members of $Z$ which maximize $\mathcal{R}$. Arrow (1959) proves that the choice function $C$ on $F$ is ordinal if and only if, for all $Z, Z' \in F$ such that $Z' \subset Z$ and $C(Z) \cap Z'$ is non-empty, one has $C(Z) \cap Z' = C(Z')$ (see also Hammond, 1977).

I shall now show that consequentialism implies that Arrow's condition for ordinality is satisfied by the revealed consequence choice function $C_\beta$. Consider, indeed, any $Z, Z' \in F$ and $\bar{\mu} \in \tilde{Y}$ such that $Z' \subset Z$ and $\bar{\mu} \in C(Z) \cap Z'$. Construct the tree:

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

to satisfy:

(i) $N^* = \{n_o, n_1\}$;

(ii) $N^o = \{n(\mu) \mid \mu \in Z\}$;

(iii) $X = \cup \{X(\mu) \mid \mu \in Z\}$ where, for all $\mu \in Z$, $X(\mu) = \{x(\mu,y) \mid y \in Y,\ \mu(y) > 0\}$;

(iv) $N_{+1}(n_o) = \{n_1\} \cup \{n(\mu) \mid \mu \in Z \setminus Z'\}$, $N_{+1}(n_1) = \{n(\mu) \mid \mu \in Z'\}$, and $N_{+1}(n(\mu)) = X(\mu)$, $(\forall \mu \in Z)$;

and, for all $\mu \in Z$, all $x(\mu,y) \in X(\mu)$:

(v) $\pi(x(\mu,y) \mid n(\mu)) = \mu(y)$;

(vi) $\gamma(x(\mu,y)) = y$.

With this construction of $T$, notice that:

$$\rho(a) = \begin{cases} \mu \in Z' & (\text{if } a(n_o) = n_1 \text{ and } \\ & \quad a(n_1) = n(\mu) \in N_{+1}(n_1)) \\ \mu \in Z \setminus Z' & (\text{if } a(n_o) = n(\mu) \in N_{+1}(n_o) \setminus \{n_1\}) \end{cases}$$

In particular, for each $\mu \in Z$, there is a unique act $a$ with $\rho(a) = \mu$, and $a(n_o) = n_1 \Leftrightarrow \rho(a) \in Z'$.

So, with this construction of $T$, $F(T) = Z$. But $\bar{\mu} \in C(Z)$ by hypothesis, so $\bar{\mu}$ is revealed chosen by $\beta$ in $T$. There is therefore an act $\bar{a}$ with $\rho(\bar{a}) = \bar{\mu}$ which is a selection from $\beta(T,\cdot)$. But $\bar{\mu} \in Z'$ by hypothesis, so $\rho(\bar{a}) = \bar{\mu}$ implies $\bar{a}(n_o) = n_1$, by the previous paragraph. Thus $n_1 \in \beta(T,n_o)$.

Consider now *any* risky consequence $\mu \in C_\beta(Z) \cap Z'$. Arguing as in the previous paragraph, with $\bar{\mu}$ replaced by $\mu$, there is an act $a$ with $\rho(a) = \mu$ which is a selection from $\beta(T,\cdot)$. Because $\mu \in Z'$, this implies that $a(n_o) = n_1$ and that $a(n_1) = n(\mu) \in \beta(T,n_1)$. By consistency of the norm $\beta$, it follows that $n(\mu) \in \beta(T_{n_1},n_1)$. This implies that $\mu \in \Phi_\beta(T_{n_1})$ because $\mu$ is the consequence of the act which selects node $n(\mu)$ at the only decision node $n_1$ of $T_{n_1}$. But $F(T_{n_1}) = Z'$ by construction of $T$, and so :

$$\mu \in \Phi_\beta(T_{n_1}) = C_\beta(F(T_{n_1})) = C_\beta(Z') .$$

Conversely, consider *any* risky consequence $\mu \in C_\beta(Z')$. Then $\mu \in \Phi_\beta(T_{n_1})$ and so $n(\mu) \in \beta(T_{n_1},n_1)$. By consistency of the norm $\beta$, it follows that $n(\mu) \in \beta(T,n_1)$. But $n_1 \in \beta(T,n_o)$ as shown two paragraphs ago. So the act $a$ with $a(n_o) = n_1$ and $a(n_1) = n(\mu)$ is a selection from $\beta(T,\cdot)$ which results in the consequence $\mu$. It follows that :

$$\mu \in \Phi_\beta(T) = C_\beta(F(T)) = C_\beta(Z) .$$

But of course $C_\beta(Z') \subset Z'$ so in fact $\mu \in C_\beta(Z')$ implies that $\mu \in C_\beta(Z) \cap Z'$.

So, under the hypothesis that $Z, Z' \in F$, $Z' \subset Z$, and $\bar{y} \in C_\beta(Z) \cap Z'$, we have proved that $C_\beta(Z) \cap Z' = C_\beta(Z')$. Arrow's characterization of ordinality has been verified, and $C_\beta$ must be ordinal. There must exist a (unique) *revealed preference ordering* $\mathcal{R}_\beta$ on $\tilde{Y}$ which is maximized by the consequence choice function $C_\beta$ revealed by the consequentialist norm $\beta$.

## 6. CONSEQUENTIALISM IMPLIES INDEPENDENCE

If risky consequences are excluded, with chance nodes not permitted in decision trees, ordinality characterizes consequentialist norms completely. With risk, however, the (revealed) preference ordering $\mathcal{R}_\beta$ corresponding to a consequentialist norm $\beta$ also has to satisfy the following *independence* property for every real number $\alpha$ with $0 < \alpha < 1$, and every set $\mu_1, \mu_2, \mu_3 \in M_o(Y)$ :

$$\alpha\mu_1 + (1-\alpha)\mu_3 \ \mathcal{R}_\beta \ \alpha\mu_2 + (1-\alpha)\mu_3 \ \Leftrightarrow \ \mu_1 \ \mathcal{R}_\beta \ \mu_2 .$$

Here, of course, $\alpha\mu_j + (1-\alpha)\mu_3$ $(j = 1, 2)$ denotes the risky consequence $\mu$ for which $\mu(y) = \alpha\mu_j(y) + (1-\alpha)\mu_3(y)$ (all $y \in Y$). This independence property is precisely the axiom first formulated by Samuelson (1952) which has since excited so much controversy.

To demonstrate this property, take any real number $\alpha$ satisfying $0 \leqslant \alpha \leqslant 1$ and any $\mu_1, \mu_2, \mu_3 \in M_o(Y)$. Construct the consequential decision tree :

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

to satisfy :

(i) $N^* = \{n_1\}$ ;

(ii) $N^o = \{n_o\} \cup \{\tilde{n}_j \mid j = 1, 2, 3\}$ ;

(iii) $X = X_1 \cup X_2 \cup X_3$ where, for $j = 1, 2, 3$, $X_j = \{x_j(y) \mid y \in Y, \ \mu_j(y) > 0\}$ ;

(iv) $N_{+1}(n_o) = \{n_1, \tilde{n}_3\}$, $N_{+1}(n_1) = \{\tilde{n}_1, \tilde{n}_2\}$, $N_{+1}(\tilde{n}_j) = X_j$ $(j = 1, 2, 3)$ ;

(v) $\pi(n_1 \mid n_o) = \alpha$, $\pi(\tilde{n}_3 \mid n_o) = 1-\alpha$, and, for $j = 1, 2, 3$ and all $x_j(y) \in X_j$, $\pi(x_j(y) \mid \tilde{n}_j) = \mu_j(y)$ ;

(vi) for $j = 1, 2, 3$ and for all $x_j(y) \in X_j$, $\gamma(x_j(y)) = y$ .

This tree has just two acts $a_1, a_2$ given by $a_j(n_1) = \tilde{n}_j$ $(j = 1, 2)$ whose consequences are $\alpha\mu_j + (1-\alpha)\mu_3$ respectively. Thus :

$$\alpha\mu_1 + (1-\alpha)\mu_3 \ \mathcal{R}_\beta \ \alpha\mu_2 + (1-\alpha)\mu_3 \ \Leftrightarrow \ \tilde{n}_1 \in \beta(T, n_1) .$$

In the continuation tree $T_{n_1}$, there are just two acts whose consequences are $\mu_1$ and $\mu_2$ respectively. So :

$$\mu_1 \ \mathcal{R}_\beta \ \mu_2 \ \Leftrightarrow \ \tilde{n}_1 \in \beta(T_{n_1}, n_1) .$$

Consistency of the norm $\beta$, however, requires that $\beta(T,n_1) = \beta(T_{n_1},n_1)$. Thus $\alpha\mu_1 + (1-\alpha)\mu_3 \; \mathcal{R}_\beta \; \alpha\mu_2 + (1-\alpha)\mu_3 \Leftrightarrow \mu_1 \; \mathcal{R}_\beta \; \mu_2$ and so independence is certainly satisfied whenever $0 < \alpha < 1$.

In addition the above equivalence has been proved for the extreme points $\alpha = 0$ and $\alpha = 1$. For $\alpha = 1$, of course, it is a tautology, but for $\alpha = 0$, we get the equivalence $\mu_3 \; \mathcal{R}_\beta \; \mu_3 \Leftrightarrow \mu_1 \; \mathcal{R}_\beta \; \mu_2$ which is a contradiction when $\mu_2 \; P_\beta \; \mu_1$ (i.e. $\mu_2$ is strictly preferred to $\mu_1$).

Does this contradiction mean that consequentialist norms do not exist ? Evidently, yes, if one insists on including in the domain $\mathcal{T}(Y)$ of the consequentialist norm *all* consequential decision trees. Yet the tree which produces the contradiction, with $\alpha = 0$ and so $\pi(n_1 \mid n_0) = 0$, is perverse in the sense that a chance node $n_0$ is included with a probability distribution attaching zero probability to one of the immediate successors of $n_0$. If risk is described by probabilities, the domain should be restricted to $\mathcal{T}^o(Y)$, the set of all consequential decision trees $T = \langle N^*, N^o, X, N_{+1}(\cdot), n_o, \pi(\cdot|\cdot), \pi(\cdot) \rangle$ with the property that $\pi(n' \mid n) > 0$ whenever $n \in N^o$ and $n' \in N_{+1}(n)$. This is hardly an important restriction; if $\pi(n' \mid n) = 0$, then $n'$ and all its successors in $N(n')$ should be excluded from the tree $T$ anyway, so that the modified tree is in $\mathcal{T}^o(Y)$.

Notice that this exclusion of zero probability nodes applies only to *finite* decision trees. The implications for infinite decision trees of consequentialism also remain largely unexplored.

If the domain is restricted to $\mathcal{T}^o(Y)$, do ordinality and independence of the revealed preference ordering characterize completely the set of consequentialist norms ? Here the answer is yes, because given any ordering $\mathcal{R}$ on $\tilde{Y}$ satisfying the independence property stated earlier in this section (with $0 < \alpha < 1$), one can construct a consequentialist norm $\beta$ whose revealed preference ordering $\mathcal{R}_\beta = \mathcal{R}$. Space does not permit more than the merest sketch of a proof. One constructs, by backward recursion, an indifference class $[\mu](n)$ of members of $\tilde{Y}$ at each node $n$ of any tree $T \in \mathcal{T}^o(Y)$. The recursion starts with terminal nodes $x$, at which $[\mu](x)$ is taken to be the set of risky consequences indifferent to $\gamma(x)$. At any chance node $n$, one takes :

$$[\mu](n) = \Sigma_{n' \in N_{+1}(n)} \pi(n' \mid n) \; [\mu](n')$$

in an obvious notation which is appropriate because of the independence property. And at any decision node $n$, one takes $[\mu](n)$ to maximize $\mathcal{R}$ over the set $\{[\mu](n') \mid n' \in N_{+1}(n)\}$, and the behaviour set $\beta(T,n)$ equal to the set of maximizing nodes $n'$ in $N_{+1}(n)$. Checking the details of this construction and that it yields a consequentialist consistent behaviour norm is then routine but tedious.

Thus, one has the following characterization :

*THEOREM.* For any space of consequences $Y$, and the domain $T^o(Y)$ of all consequential finite decision trees with only positive probabilities at every chance node, the consistent behaviour norm $\beta$ is consequentialist only if there is a (unique) revealed preference ordering $\mathcal{R}_\beta$ satisfying the independence property. Conversely, given any preference ordering $\mathcal{R}$ which satisfies independence on the space $\tilde{Y}$ of all risky consequences, there is a (unique) consequentialist norm $\beta$ on $T^o(Y)$ which reveals a preference ordering $\mathcal{R}_\beta$ equal to $\mathcal{R}$.

Notice that consequentialism does not imply maximizing the expected value of a von Neumann-Morgenstern utility function unless one imposes an additional continuity axiom; only two of the three axioms of Herstein and Milnor (1953) are implications of consequentialism.

## 7. WHAT CONSEQUENCES ARE RELEVANT ?

If consequentialism implies the independence axiom, what are we to make of the Allais (1953) paradox and the kind of preferences discussed by Allais (1979), Machina (1982) and others which violate this axiom ? In Machina's notation, such preferences are usually assumed to be represented by a utility functional $U(F)$ whose argument $F$ is a cumulative distribution function of money income $y$. It is also usually assume that $U(F_1) > U(F_2)$ when $F_1$ stochastically dominates $F_2$. Generally, of course, the independence axiom is violated, so there is no von Neumann-Morgenstern utility function (NMUF) of income $y$ for which $U$ can be written in the expected utility form $U(F) \equiv \int v(y)\, dF(y)$. But $U$ can be written in the expected utility form $U(F) \equiv \int w(y,F)dF(y)$ provided $w(y,F) \equiv U(F)$ ; indeed, making $w$ depend only on $F$ and not on $y$

is really a special case. Thus, preferences which appear to violate the independence axiom will not do so if the domain of relevant consequences is expanded from y's to pairs (y,F) . Of course, the parsimony of the original representation U(F) has been lost.

This is the nature of consequentialism. It says that, for a *sufficiently rich domain of consequences*, behaviour *should* depend *only* on those consequences, and then certain mathematical properties such as the independence axiom emerge as logical implications. As Broome (1986) also discusses, consequences have to be "individuated" finely enough, otherwise something relevant for a normative principle of behaviour is being left out of account. Indeed, if consequences are too coarsely individuated, it may appear that no single preference order gives an acceptable normative principle of behaviour. So a *necessary* condition for the existence of preferences, as well as for the independence axiom, is that consequences are defined finely enough.

This paper shows that exactly the same condition is *sufficient* not only for the existence of preferences, but also for the controversial independence axiom. The conclusion is that, in objecting to the independence axiom, one also objects to consequentialism, which is the best argument I know for the existence of preferences. One is led to ask why those who would reject the independence axiom sill cling to the assumption of a preference ordering.

Over the years I have heard many objections to the arguments of this paper. The most serious concerns the consistency condition of Section 3. It has been claimed that continuation trees should be treated differently from entire decision trees. Perhaps history matters. But if it does, it is a relevant consequence. One may ask then whether historical consequences can be disentangled from the structure of each decision tree so that every tree of the domain $T^o(Y)$ represents a meaningful decision problem, as I have assumed. But that cannot be discussed here.

*Department of Economics*
*Stanford University*

ACKNOWLEDGEMENTS

This and related work over the years have been supported by the U.S. National Science Foundation, the John Simon Guggenheim Memorial Foundation, and C.O.R.E. at the Université Catholique de Louvain. I recall helpful discussions with Ken Arrow, Eddie Dekel, Jacques Drèze, Yossi Greenberg, Mark Machina, Ned McClennen, Hervé Moulin, Prasanta Pattanaik, Amartya Sen and Bob Sugden. I owe much to all of these and no doubt several others too, especially some interesting opponents.

## REFERENCES

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.

Allais, M. (1979), 'The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School', in *The Expected Utility Hypothesis and the Allais Paradox*, M. Allais and O. Hagen (eds.), Dordrecht : D. Reidel, pp. 27-145.

Anscombe, G.E.M. (1958), "Modern Moral Philosophy', *Philosophy, 33*, 1-19.

Arrow, K.J. (1959), 'Rational Choice Functions and Orderings', *Economica, 26* : 121-127.

Broome, J. (1986), 'Should a Rational Person Conform to the Sure-Thing Principle ?', presented at the 3rd International Conference on the Foundations and Applications of Utility, Risk and Decision Theories, Aix-en-Provence.

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. (1986), 'Consequentialist Social Norms for Public Decisions', in *Social Choice and Public Decision Making : Essays in Honor of Kenneth J. Arrow, Vol. I*, W.P. Heller, R.M. Starr and D.A. Starrett (eds.), Cambridge University Press, Chapter 1, pp. 3-27.

Herstein I.H. and J. Milnor (1953),' An Axiomatic Approach to Measurable Utility', *Econometrica, 21* : 291-297.

Machina, M. (1982), ' "Expected Utility"Analysis without the Independence Axiom', *Econometrica, 50* : 277-323.

Raiffa, H. (1968), *Decision Analysis : Introductory Lectures on Choices under Uncertainty*, Reading, Mass : Addison-Wesley.

Samuelson, P.A. (1952), 'Probability, Utility, and the Independence Axiom', *Econometrica, 20* : 670-678.

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