

Discussion Paper No. 994

CASE-BASED DECISION THEORY*

by

Itzhak Gilboa**

and

David Schmeidler***

May 1992

*We are grateful to many people for conversations and discussions which influenced this work. In particular we would like to mention the faculty and students of the Institute for the Learning Sciences at Northwestern University, as well as faculty and guests at the Sante Fe Institute. We would also like to thank Max Bazerman, Eva Gilboa, Ehud Kalai, Mort Kamien, Jim Peck, Mike Sang and Karl Schlag for comments and references. NSF Grants Nos. SES-9113108 and SES-9111873 are gratefully acknowledged.

**Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208.

***Department of Economics, Ohio State University, Columbus, Ohio 43210; Departments of Economics and Statistics, Tel Aviv University, Tel Aviv 69978, Israel.

Abstract

This paper suggests the view that decision-making under uncertainty is, at least partly, case-based. We propose a model in which cases are assumed as primitives, and provide a simple axiomatization of a decision rule which chooses a "best" act based on its past performance in "similar" cases. Each act is evaluated by the sum—over cases in which it was chosen—of the product of the similarity of the past case to the problem at hand and the utility level that resulted from this act in the past.

As in expected utility theory, both the utility and the similarity functions may be derived from preferences and the latter are represented by (the maximization of) a sum of products. However, there are some crucial differences between case-based decision theory and expected utility theory. In the former:

- every two acts are evaluated over completely different (and disjoint) histories of cases;
- neither probabilities nor states of the world are assumed as primitives. Moreover, the theory does not distinguish between certain and uncertain acts;
- the notions of "satisficing" decisions and aspiration levels pop up naturally from the axiomatic derivation of case-based decisions.

The paper also discusses various aspects, variations and applications of the basic model.

1. Introduction

Within the areas of economic, decision and game theory, expected utility theory enjoys the status of an almost-unrivaled dominant paradigm for decision-making in face of uncertainty. Relying on such sound foundations as the classical works of de Finetti (1937), von Neumann-Morgenstern (1944), Savage (1954) and Anscombe-Aumann (1963), the theory has formidable power and elegance, whether interpreted as positive or normative, for situations of given probabilities ("risk") or unknown ones ("uncertainty") alike.

The strength and appeal of the expected utility paradigm are only attested to by its various generalizations. While evidence has been accumulating that, if taken in its classical form, the theory is too restrictive (at least as a descriptive one), almost all the alternative models suggested in the literature attempt to cling to it as much as possible. With few exceptions, the generalized models relax some of the more "demanding" axioms (such as linearity of preferences with respect to probabilities, or additivity of the probability measure)—while retaining the more "basic" ones (such as transitivity). With almost no exceptions at all, these models retain the framework of whatever classical model they generalize. (See Machina (1987), Harless-Camerer (1991), and Camerer-Weber (1991) for extensive surveys.)

Yet it seems—and often also heard—that in many situations of choice under uncertainty, the very language of expected utility models is inappropriate. For instance, it is hardly controversial that most of the problems of interest are not formulated in terms of given probabilities ("risk"). Excluding some examples such as gambling and planned experiments, probabilities can seldom be argued to be a part of the "observable" environment of the decision maker.

Indeed, one of the fundamental contributions of de Finetti (1937), Savage (1954) and Anscombe-Aumann (1963) is the derivation of "probability"

from a more basic model, whose primitives include states of the world but not a measure on them. While there is no denial that the applicability of such a model by far exceeds that of the "risk" model, we argue that it still falls short of covering the whole range of interesting phenomena. More specifically, we claim that in many decision problems under uncertainty, the states of the world are neither naturally given, nor can they be simply formulated.

Comparing two examples may be helpful. First, consider Savage's famous omelette problem (Savage (1954, pp. 13-15)): he is about to add a sixth egg to a bowl containing five, and faces uncertainty regarding its freshness. Quite clearly, by introducing two states of the world—"the egg is fresh" and "the egg isn't fresh"—one may reduce uncertainty to the question of which of them obtains.

Next assume that our decision maker (DM) is about to dine in a restaurant but, alas, (s)he does not understand the language in which the menu is written. (Our DM is in a foreign country, attending a conference on EU generalizations.) Here the uncertainty is about the meaning of a certain word or, worse still, many words, expressions and propositions. The corresponding set of states of the world is bound to be humongous and can hardly be assumed as "naturally" given.

From a descriptive viewpoint, we find it highly unrealistic to assume that people tend to solve the menu problem by states-of-the-world type of reasoning, or that they behave in a way consistent with such reasoning. But even with a normative interpretation the classical model seems unsatisfactory: as much as one may find (say) Savage's model appealing, applying it in this case is hardly a practical recommendation for our DM once (s)he decides to dine out.

Naturally, similar (and more vehement) claims can be made regarding more

demanding models; the requirement that our DM would have a prior over the states of the world, for instance, seems even more far-fetched as a descriptive claim and even more useless as a normative one.

Before proceeding to propose an alternative model, a clarification is due. Our DM in the restaurant may formulate a states-of-the-world model in which all things edible are classified as "fish," "meat" and so on, and perhaps even form beliefs about these states. Moreover, we trust that adamant Bayesians will always be able to formulate such models and feel comfortable with them. However, we are not trying to claim that the expected utility paradigm is useless, nor that it is inferior to the proposed alternative. The claim is only that for certain situations, especially those involving ignorance, and especially for descriptive purposes, other paradigms may be more insightful.

Our suggestion is to adopt models, developed in psychology and artificial intelligence, and adapt them to decision making. More specifically, we would like to focus on the theory of case based reasoning ("CBR"; see Riesbeck and Schank (1989)), which says, loosely, that the main reasoning technique people use, especially in novel situations, is based on comparing the situation at hand to past cases, making analogies and drawing corresponding conclusions.

Encouraged by the success of CBR in the case artificial intelligence, we are suggesting the following model: let us assume that a set of "problems" is given as a primitive, and that there is some measure of similarity on it. (We will be more precise in the sequel.) The problems are to be thought of as descriptions of choice situations, as "stories" involving decision problems. For instance, our DM may have already been in a similar situation, facing a cryptic menu, though possibly in a different country. Yet the similarity of the situation brings this memory to mind, and with it, hopefully, the

recollection of the choice made and the outcome that resulted. We will refer to the combination of these three—the problem, the act and the result—as a "case."

Generally, let us assume that all "similar" cases are recalled, and based on them each decision is evaluated. The specific model we propose and axiomatize here will evaluate a decision by the sum, over all cases in which it was chosen, of the product of the similarity of the case to the one at hand and the resulting utility. (The "utility" will be assumed to be scaled such that the value zero will be a default value.)

For example, suppose that our dining DM (or "DDM" for short) has already been three times in a similar situation. In two of them (s)he opted for an item on the menu containing the letters "sala"; in a fast-food restaurant the decision resulted in a salad (which was more or less what DDM expected), while in a regular restaurant (in a different country) it ended up being an interesting but inedible type of fish. The third case recalled was also in a regular restaurant, but here the choice made was the first item in the middle column of the menu. It turned out to be a steak.

Given that the decision problem at hand takes place at a regular restaurant, the two last cases get higher "similarity" or "relevance" coefficient than the first. These coefficients are multiplied by a utility function (measuring the tastiness of the food) and the results are summed separately for the two decisions. Thus, the act "choose an item whose name contains 'sala'" would be numerically evaluated as the sum of two terms; that of "choose the first item in the middle column" would be evaluated by a single term.

The decision between these two acts will be made according to these values. However, if both are negative, DDM may resort to a different act, whose default value is zero. Thus, reasonable ("satisficing") acts may be

chosen again and again, while catastrophic ones will lead DDM to further experiment and try new choices.

A preview of the formal model may be helpful at this point. We assume as primitives a set P of problems, a set A of acts (or decisions) and a set R of outcomes (or results). The set of cases is defined to be $C = P \times A \times R$. A decision problem is represented by a problem $p \in P$ and a memory (or history) $M \subseteq C$. The memory M is interpreted as all past decision situations, together with the acts chosen and the outcomes that happened to result from them. We will assume that if $m = (q, a, r) \in M$ then $p \neq q$, and if $m' = (q', a', r') \in M$ as well, with $m' \neq m$, then $q' \neq q$. From a binary "preference" relation we derive a similarity function $s: P^2 \rightarrow [0,1]$ and (in a trivial sense) a utility function $u: R \rightarrow \mathbf{R}$ such that, given a problem $p \in P$ and a memory M , the acts $a \in A$ are ranked according to

$$(*) \quad U(a) = \sum_{(q,a,r) \in M} s(p,q)u(r).$$

(Where the summation over the empty set is taken to yield zero.)

We will assume that no problem $p \in P$ may be encountered more than once, and the fact that $p, q \in P$ may be "identical" can be reflected by $s(p,q) = s(q,p) = 1$. Notice that we have not required the similarity function to be symmetric. Indeed, this will not be necessary, and, in view of psychological evidence, can be unduly restrictive. (See Tversky (1977).)

Although the formula above cannot fail to remind us of expected utility (which we consider as a very successful "case"), one should note that it has very little in common with it: first, there is no reason for the coefficients $s(p, \bullet)$ to add up to 1 or any other constant. Moreover, while in expected utility theory every act is evaluated at every state, here each act is evaluated over a different set of cases. To be precise, if $a \neq b$, the set of

elements of M summed over in $U(a)$ is disjoint from that corresponding to $U(b)$. In particular, this set may well be empty for some a 's.

Second, in expected utility theory uncertainty is represented by the probabilities. In the theory suggested here—let us dub it case-based decision theory (CBDT)—the uncertainty, to the extent that it has been observed, would be reflected in the fact that for very similar (or even "identical") problems p, q and the same act a , M may contain elements (p, a, r) and (q, a, t) with $r \neq t$.

In other words, beliefs and/or probabilities do not explicitly exist in this model. However, they may be implicitly inferred from the number of summands in (*). That is, if the decision maker happens to choose the same act in many similar cases, the evaluation function (*) may be interpreted as gathering statistical data, or as forming a "frequentist" prior.

Putting this decision rule in a dynamic context one may tell the following story: at the beginning, the memory set M is empty. At every stage, some process introduces a decision problem $p \in P$ to the DM. We do not model this process, and implicitly assume the DM is not aware of it, has no beliefs about it and so forth. When $M = \emptyset$, the DM's choice is bound to be arbitrary (every $a \in A$ maximizes $U(\bullet) = 0$.) At later stages, the DM's choice is biased by M —in favor of acts a with $U(a) > 0$ and against acts a with $U(a) < 0$.

For the sake of illustration, consider the extreme case in which the DM faces the "same" problem over and over again, i.e., $s(p, q) = 1$ for all $p, q \in P$. Further assume that there is no uncertainty, i.e., that for every $a \in A$ there is $r_a \in R$ such that only elements of the form (q, a, r_a) are in M . In this case, the first act $a \in A$ with $u(r_a) > 0$ will be chosen over and over again. Differently put, our DM will never even attempt to maximize the utility u . (S)he is satisfied with the "reasonable" act a (so defined by

$U(a) > 0$) and exhibits, if you will, extreme conservatism or uncertainty aversion.

There are two main reasons that may prod our decision maker to experiment: first, uncertainty is present, and certain acts may result in negative U values. Second, not all cases are similar. Consider, for instance, the following scenario: a_1 may be chosen for the first problem p_1 . If the second problem p_2 happens to be rather different (for instance, making an omelette as opposed to choosing from an unintelligible menu), say $s(p_2, p_1) = 0$, a different act a_2 may be chosen. For yet another problem p_3 , which may be similar to both p_1 and p_2 (say, ordering your meal from a translated menu) both acts a_1 and a_2 may be competing to be chosen. Thus, a_2 may eventually be preferred to a_1 even in cases which are "identical" to p_1 and even if no uncertainty is present: the mere variety of problems may introduce enough "noise" to induce experimentation.

However, the decision rule (*) by no means forces experimentation. Our decision makers, the U -maximizers, are not u -maximizers; they tend to be content with reasonable options, unless they have a "good reason" to believe others are better. (Compare this mode of behavior with the notion of "satisficing" decisions of Simon (1957) and March-Simon (1958).) We will later discuss some extensions of this model.

Further discussion may prove more useful after a formal presentation of our model and results. We devote Section 2 to this purpose. In Section 3 we discuss some generalizations of the basic model. We show that these generalizations may capture additional features such as the adjustment of the "aspiration level" and "frequentist" expected utility maximization. Section 4 deals with the structure of the similarity function and outlines directions for further research. Section 5 presents some economic applications. Section 6 compares CBDT to EUT, while Section 7 concludes with some comments.

2. The Model

Let P be a nonempty set of problems. For simplicity, assume that P is finite. (For a dynamic problem one may need an infinite set, out of which only finitely many elements appear in M at each stage.) Let A be a finite and nonempty set of acts. For notational convenience we will assume that all the acts A are available at all problems $p \in P$. It is straightforward to extend the model to deal with the more general case in which for each $p \in P$ there is a subset $A_p \subseteq A$ of available acts. Let R be a set of outcomes or results. The set of cases is $C = P \times A \times R$.

A memory is a (finite) subset $M \subseteq C$ such that $m_i = (p_i, a_i, r_i) \in M$ ($i = 1, 2$) and $m_1 \neq m_2$ implies $p_1 \neq p_2$. Given a memory M , denote

$$H = H(M) = \{p \in P \mid \exists a \in A, r \in R, \text{ s.t. } (p, a, r) \in M\}.$$

That is, H is the set of problem recalled.

Next, we would like to define a preference order over acts, which will be representable by the functional U above. In principle, for every p and M we should have a separate order $\succeq_{p, M}$ over the finite set A . However, we will assume a much more informative order $\succeq_{p, H}$ which depends only on p and the observed problems H ($p \notin H$), and which may compare any pair of hypothetical acts which are "compatible" in a sense to be explained shortly.

For convenience, let us formally introduce a new outcome r_0 to R to be interpreted as "this act was not chosen." For every memory M , and $q \in H = H(M)$, there will be one act chosen at q —with an outcome defined by M —and the other acts will be assigned r_0 .

It seems innocuous to assume that an act is evaluated based on the outcomes it led to alone. Thus, for a given H , let an act profile be an element of $R^H = X$.

Obviously, every act has a unique "act profile" defined by M . We will assume that $\succeq_{p,H}$ compares acts based only on their act profile, or that $\succeq_{p,H} \subseteq X \times X$.

However, we will not assume that $\succeq_{p,H}$ is a complete order on R^H : consider two act profiles assigning $r_1 \neq r_0$ and $r_2 \neq r_0$, respectively, to some $q \in H$. Naturally, these cannot be compared even hypothetically: for any memory M , at most one act may be chosen in case q , and therefore at most one act may have a value different from r_0 in its act profile for any given q . We therefore define two act profiles $x, y: H \rightarrow R$ to be compatible if for every $q \in H$ either $x(q) = r_0$ or $y(q) = r_0$ (or both).

We can now state our first axiom.

A1: For every $p \in P$ and every history $H = H(M)$, $\succeq_{p,H}$ is reflexive and transitive on $X = R^H$, and for every compatible $x, y \in X$, $x \succeq_{p,H} y$ or $y \succeq_{p,H} x$.

We now wish to formulate some "monotonicity," "continuity" and "independence/separability" assumptions that will guarantee the additively separable representation of $\succeq_{p,H}$ on X .

The state of the art in decision theory is such that one actually faces a non-trivial choice problem here: there is a wide variety of frameworks and axioms guaranteeing such a result. Indeed, some care must be taken since our relation is only a partial one, but it seems safe to conjecture that in almost any framework the crucial axiom may be appropriately modified to guarantee the desired result with no serious loss of elegance.

To simplify the exposition we will henceforth assume (explicitly) that $R = \mathbf{R}$ (the reals) and (implicitly) that it is already measured in "utils." That is, the rest of the axioms should be interpreted as if R were scaled so that the "utility" function be the identity.

Notice that under the assumption $R = \mathbf{R}$, $X = \mathbf{R}^H$ may be identified with \mathbf{R}^n for $n = |H|$. Furthermore, we will assume that $r_0 = 0$, whence $x, y \in X$ are compatible iff $x * y = 0$ (where the product is taken as a pointwise operation on vectors in \mathbf{R}^n .) We can now formulate

A2 Monotonicity: For every p, H , $x \geq y$ and $x * y = 0$ implies $x \succeq_{p, H} y$.

A3 Continuity: For every p, H , and $x \in X$ the sets $\{y \mid y \succeq_{p, H} x\}, \{y \mid x \succeq_{p, H} y\}$ are closed (in the standard topology on \mathbf{R}^n).

A4 Separability: For every p, H and $x, y, z, w \in X$, if $x * y = 0$, $(x + z) * (y + w) = 0$, $x \succeq_{p, H} y$ and $z \succeq_{p, H} w$, then $(x + z) \succeq_{p, H} (y + w)$.

Proposition 1: If A1-A4 hold, then for every $p \in P$ and every H , there exists a function

$$s_{p, H}: H \rightarrow \mathbf{R}_+$$

s.t.

$x \succeq_{p, H} y$ iff $\sum_{q \in H} s_{p, H}(q)x(q) \geq \sum_{q \in H} s_{p, H}(q)y(q)$ for all compatible $x, y \in X$.

Next we would like to express the fact that the similarity measure is independent of the specific memory M .

A5 Similarity Invariance: For every $p, q_1, q_2 \in P$ and every two memories M^1, M^2 with $q_1, q_2 \in H^i = H(M^i)$ ($i = 1, 2$) and $p \notin H^i$ ($i = 1, 2$), the following holds:

$$\text{if } x, y \in \mathbf{R}^{H^1}, z, w \in \mathbf{R}^{H^2}, x \sim_{p, H^1} y, z \sim_{p, H^2} w$$

and $x + \alpha v_1 \sim_{p,H} y + \beta v_2$, then $z + \alpha v_1 \sim_{p,H} y + \beta v_2$,

where v_j stands for the unit vector in \mathbb{R}^H ($i = 1, 2$) corresponding to q_j ($j = 1, 2$) and whenever the compared profiles are compatible.

Proposition 2: Assume that, on top of the conditions of Proposition 1, A5 holds. Then there exists a function $s: P^2 \rightarrow [0, 1]$ such that for all $p \in P$, every memory M with $p \notin H = H(M)$ and every compatible $x, y \in \mathbb{R}^H$

$$x \succeq_{p,H} y \text{ iff } \sum_{q \in H} s(p, q)x(q) \geq \sum_{q \in H} s(p, q)y(q).$$

The proofs of both propositions are given in the Appendix.

3. Variations

3.1 Memory-Dependent Utility

The framework used in Section 2, in which outcomes are identified with utility levels, is rather convenient to convey the main idea, but it may also be misleading: it entails the implicit assumption that the utility function does not depend on the memory M , on time (which may be implicit in M) and so forth.

To illustrate this point, consider an axiomatization which is similar to that given above, but which derives the utility function on an abstract (say, connected topological) space. (Axioms A4 and A5 will obviously have to be rephrased, to express the utility "addition" in more primitive terms.) Corresponding to Proposition 1, one may then derive a representation theorem where both the similarity and the utility functions depend on p, H . Similar to axiom A5, one may impose an additional axiom that will guarantee that for all p, H , $u_{p,H}(\bullet)$ are identical up to a positive linear transformation.

Yet in this context it becomes clear that there may be some interest in a more general model, where the utility is allowed to vary with memory. Recall that the utility is normalized so as to set $u(r_0) = 0$. One may refer to this value as the "aspiration level" of the decision maker: as long as some acts have a U-value exceeding $u(r_0)$, the DM is "satisfied" and does not attempt new choices. Thus we may view $u(r_0)$ as a behavioral definition of the "aspiration level."

Given this interpretation, it is natural to suggest that the aspiration level be adjusted according to past achievements. Strictly speaking, our model (even the generalized one) does not allow for an axiomatic derivation of such utility functions, since $u_{p,H}$ depends only on the problems encountered--and not on the acts and outcomes associated with them in M . Yet one may construct a similar model in which the utility depends only on

$$\bar{u}(M) := \max\{u(r) \mid \exists p \in P, a \in A, \text{ s.t. } (p,a,r) \in M\}.$$

(For such a model, the preference order $\succeq_{p,H,\bar{u}(M)}$ will be defined on all utility profiles with the same maximal utility level $\bar{u}(M)$.)

Then one may model the fact that aspirations are adjusted upwards by setting, say

$$u_{p,H,\bar{u}(M)}(r_0) = \bar{u}(M) + 1,$$

where the utility difference of 1 plays the role of a "just noticeable difference" in aspirations.

We shall not expatiate on this model here and focus on the basic model, with memory-independent utility, or fixed aspiration level. However, for some dynamic applications relaxing this assumption may prove a theoretical

necessity. (See subsection 5.2 below.)

3.2 Act-Dependent Similarity

In Section 2 above we assumed that preference is given between any two compatible utility profiles. Implicitly we therefore assumed that, fixing the set of problems encountered H , the DM has preferences on hypothetical acts which differ from the actual ones not only in terms of the outcome but also in terms of the act chosen. To clarify these two "levels" or counterfactuals, consider the following example: suppose the DM prefers act a to b . We pose to him/her two types of questions:

- I. Remember the case $c = (p, a, r)$ where you chose a and got r ? Well, assume the outcome were t instead of r . Would you still prefer a to b ?
- II. Remember the case $c = (p, a, r)$ where you chose a and got r ? Well, now imagine you actually chose another act a' and received t . Would you still prefer a to b ? How about a' to b ?

One may argue that questions of type II are too hypothetical to serve as foundations of any "behavioral" decision theory. While the DM has no control over the outcome r , he/she may insist that in problem p he/she would never have tried act a' and the preference question is meaningless.

But even if we take a less extreme position there is some theoretical value in an axiomatic derivation of a similar model in which only answers to questions of type I are assumed given. We outline such a model in Appendix 2.

Apart from the less stringent (implicit) rationality requirement that such a model imposes, it also allows us to derive a similarity function which depends on the acts. That is, for a memory M let

$$E = E(M) = \{(p, a) \in P \times A \mid \exists r \in R, \text{ s.t. } (p, a, r) \in M\}.$$

Then the model derives a similarity function $s_{p,E}$ which depends not only on the problems encountered but also on the acts chosen in each.

A special case of this model (which is axiomatized in Appendix 2) is the following: let $s(p,q)$ be the similarity function as in Section 2, but assume that each act is evaluated by its average performance:

$$(**) \quad V(a) = \sum_{(q,a,r) \in M} s'(p,q)u(r)$$

where

$$s'(p,q) = \begin{cases} \frac{s(p,q)}{\sum_{(q',a,r) \in M} s(p,q')} & \text{if well-defined} \\ 0 & \text{otherwise.} \end{cases}$$

Thus, for every act a the similarity coefficients $s'(p,q)$ add up to 1 (or to zero). Notice that the similarity function here depends on E , and not merely on H .

In general, maximization of V does not appear to be a very reasonable description of behavior. For instance, V is discontinuous in the similarity values. Thus, if an act a was chosen in a single case q and resulted in a very desirable outcome, it will be chosen as long as $s(p,q) > 0$ but will be considered a "new" act if $s(p,q) = 0$.

However, consider the special case where $s(p,q) = 1$ for all $p,q \in R$. (See Appendix 2 for an axiomatization.) In this case, V is simply the average utility of each act.

The condition $s(\bullet,\bullet) = 1$ means that (at least as evidenced by the DM's preferences) all problems are basically identical. In a sense, our model reduces to classical decision theory, where each problem is considered in

isolation. For this case, this variant of case-based decision theory reduces to "frequentist expected utility theory": the DM maximizes the expected utility where the outcome distribution for each act is simply assumed to be given by the observed frequencies. (Notice also that in this particular model the discontinuity at $s(\bullet, \bullet) = 0$ disappears, since $s(\bullet, \bullet) = 1$.)

3.3 Experimentation

The model presented above, as well as the "frequentist prior" which is implicitly gathered in the basic model (maximization of U as in $(*)$) tempt one to conjecture that under certain conditions, expected utility maximization may pop up (asymptotically) from case-based decisions.

Unfortunately, this does not seem to be the case in general. Consider the following set-up: $A = \{a, b\}$, $s(\bullet, \bullet) = 1$. We will assume that, unbeknown to the DM, nature chooses the outcomes for each act by given distributions in an independent fashion. That is, there are two random variables R_a, R_b such that whenever the DM chooses $a(b)$, the outcome is chosen according to a realization of $R_a(R_b)$, independently of past choices and realizations.

Further assume the following distributions:

$$R_a = \begin{cases} 1 & .6 \\ -1 & .4 \end{cases}; \quad R_b = \begin{cases} 100 & .7 \\ -2 & .3 \end{cases}$$

First consider a U -maximizer DM. At the beginning, both a and b have identical (empty) histories, and the decision is arbitrary. Suppose that the DM chooses a with probability $.5$, and that R_a results in $+1$. From then on our DM will choose a as long as the random walk generated by these choices is positive. However, given that $.6 > .4$ this happens with positive probability.

Next consider a V -maximizer DM. Suppose that (s)he chose b , which resulted in the outcome (-2) . From now on this DM will always choose a .

In both cases we find that there is a positive probability that our DM will not maximize the "real" expected utility even in cases where there exists an "objective" process which defines these "real" probabilities. Even the frequentist EU maximization does not have to converge to the "objective" EU maximization with ("objective") probability 1.

These examples naturally suggest yet another generalization of the model: the introduction of conscious, intentional experimentation. For either U- or V-maximizers one may alter the decision rule so that every act will be chosen every so often regardless of its U/V value. It seems reasonable that with sufficiently frequent experimentation, frequentist EU maximization will boil down, asymptotically, to EU maximization with probability close to 1 as desired. However, we do not pursue this track in this paper.

3.4 Similar Acts

There are applications in which it is natural to assume that the DM has some information regarding an act without having tried it in the past. For instance, one may consider buying a house in a neighborhood one has lived in before. Thus some information about this act may be gleaned from other acts which have been attempted and are "similar" to it in some sense.

Some of these applications may be embedded in our model by redefining the act (say, "buy a house in this neighborhood" rather than "buy this house") and/or by appropriately defining the similarity over the problems. (See also comment 7.3 in Section 7 below.) However, these may not suffice if we would like our DM to be able to compare acts which were actually tried to acts which are only similar to them. Thus one may consider a generalization of the model in which the DM has two similarity functions: $s_P: P^2 \rightarrow [0,1]$ on problems and $s_A: A^2 \rightarrow [0,1]$ on acts. Given these, one may redefine the

evaluation functional to be

$$U'(a) = \sum_{(p,q,r) \in M} s_P(p,q) s_A(a,b) u(r).$$

Discussions, axiomatizations and applications of this model are beyond the scope of this paper.¹

4. The Structure of Similarity

In the model we present here, the decision problems are some abstract set P , on which the similarity measure is derived axiomatically. Yet much insight into specific problems may be gained from analyzing the structure of a "decision problem" and the corresponding structure of the similarity function. We do not purport to develop here a general theory of similarity, partly because (at least to a large extent) such theories do exist in the psychological literature on analogies. (See Gick and Holyoak (1980, 1983), Falkenhainer, Forbus and Gentner (1989), and others.) However, we would like to draw the reader's attention to some possibilities of specific modeling which, in particular, will also show that case-based decision theory is more general than may seem at first.

4.1 Collective Memory

Thus far it was implicitly assumed that the memory consists of cases in

¹It appears that a preference order over acts does not contain enough information for an axiomatic derivation of unique or even meaningful similarity functions s_P and s_A . These functions leave too much freedom, and the theory that the DM maximizes U' for some s_P, s_A may not say much beyond the claim that the DM's preference over acts is a weak order (this would depend, of course, on the specifics of the model). However, one may assume as a primitive a "more similar than" relation, defined on pairs of acts (i.e., a subset of A^4) on top of the preference relation on acts, and axiomatize a representation of preferences by U' with a function s_A which also represents the act-similarity relation.

which our DM made similar decisions. However, some cases recalled by the DM may involve other decision makers as the "protagonists" of the problems. Especially when it comes to weighty decisions—such as, say, buying a car—one is more likely to learn from other people's experience, which is typically shared among many decision makers, rather than count solely on first-hand experience. Moreover, there are some decisions which, by their temporal nature, can only be made based on others' experience. One's (first) career choice and (first) marriage have no similar cases in one's personal memory, yet are important enough to seek the advice of more experienced DM's.

One would expect that, should a problem $p \in P$ have a "protagonist" as a formal component of it, the similarity of p to q will be higher (other things being equal) if they have the same protagonist than otherwise. Moreover, some measure of proximity between protagonists may be a factor in the similarity of two problems, where one's friend has a higher proximity level than a stranger, yet lower than the self.

4.2 Hypothetical Cases

Suppose you have to drive to the airport in one of two ways. When you get there safely you learn that the other road was closed for construction. A week later you are faced with the same choice. Regardless of your aspiration level $u(r_0)$, it seems obvious that you will choose the same road again. (Road constructions, at least in psychologically-plausible models, never end.)

In other words, the cases recalled from memory may well contain some hypothetical, counterfactual ones. ("Had I taken the other way, I would never have made it.")

As in the case of protagonists, other (counterfactual) "possible worlds" need not be lumped together when similarity is concerned. One may quantify some "degree of belief" in a counterfactual statement as above (with, say, 1

designating an actual case) and let it play a role in the similarity function.

Hypothetical cases may endow a case-based decision maker with reasoning abilities he/she would otherwise lack. It seems that any knowledge the DM possesses and any conclusions he/she deduces from it can, inasmuch as they are relevant to the decision at hand, be reflected by hypothetical cases.

Indeed, one may actually "simulate" an expected-utility maximizer by a case-based decision maker whose memory contains sufficiently rich hypothetical cases: given a set of states of the world Ω and a set of consequences R , let the set of acts be $A = R^\Omega = \{a: \Omega \rightarrow R\}$. Further assume that the DM has a utility function $u: R \rightarrow \mathbf{R}$ and a probability measure p on Ω (which is some measurable space. For simplicity it may be assumed finite.) The corresponding case-based decision maker would have a hypothetical case for each pair of state of the world ω and act a :

$$M = \{((\omega, a), a, a(\omega)) \mid \omega \in \Omega, a \in A\}$$

By setting the similarity of the problem at hand to the "problem" (ω, a) to equal $p(\omega)$, U -maximization reduces to expected utility maximization. (Naturally, if Ω or R are infinite one would have to allow for infinite memory M as well.)

Thus, EUT may be mathematically embedded in CBDT. However, we do not find this construction very appealing and we believe that the more interesting applications of CBDT will be those in which hypothetical cases are restricted to be psychologically plausible. (See Section 5 below.)

4.3 Parameterized Problems

It will often occur that some specific numerical data will be part of the "story" of a problem. For instance, all problems occur at a given time;

purchasing decisions involve prices and quantities, interest rates and so forth. It seems natural to use these data as part of the description of the problem, i.e., to parameterize it. Thus "Should I buy this car at price p ?" may be a parameterized problem, and one may have several such cases—for possibly different values of p —in memory.

Once a parameter is given, it makes sense to consider similarity functions which explicitly depend on this parameter. Thus, the buying problem at price p may be similar to the same one at price q to the extent

$$e^{-\theta|p-q|}$$

for some $\theta > 0$.

Similarly, an analogous decay function for the time parameter may reflect the fact that the older data are deemed less relevant, and maybe also have a lower probability of being recalled. (Recall that our model does not distinguish between the probability of recall and the conscious similarity judgment. The similarity function summarizes both.) However, one may consider more general functions, which allow for primacy as well as recency effects.

4.4 History-Dependent Problems

There are cases, such as repeated games, in which the description of a single decision already contains history of its own. For instance, the decision on a one-shot move in a repeated game may generate a sequence of problems p_t , each of which has a $(t - 1)$ -long history of the play.

In such cases the similarity function may depend on some features of this history. For instance, for every history one may compute the relative frequencies with which each of the opponent's move was chosen, and the

similarity of two problems may depend on the values of these summary statistics. Alternatively, a finite-recall strategy will (implicitly) judge two cases to be similar if their most recent histories are identical. In a sense, every type of repeated-game strategy which uses less information than the complete history implicitly defines some notion of similarity on such histories.

4.5 Rule-Based Decisions

Many decisions seem to be taken in an almost automated way. People often operate using "rules of thumb" such as "Do not invest in totalitarian states," "Keep all receipts," "Do not incur a debt exceeding ten percent of the worth of your assets," and so forth.

Such rules may be viewed as summarizing many cases. (Indeed, this is probably the way most of them came into existence.) Thus one may incorporate them in our framework in (at least) two ways: first, they may be "translated" to many cases, probably those from which they originated, to yield an equivalent decision rule. Alternatively, one may introduce a rule as a single case, where the similarity that is born to it by any relevant decision problem by far exceeds the similarity to "regular" cases. (See also the discussion of "ossified cases" in Riesbeck and Schank (1989).)

In this section we attempted to describe some of the possible directions one may choose when analyzing the concept of "similarity." Of course, none of them is fully explored, let alone axiomatically justified. At this point we would merely like to emphasize the wide applicability of case-based decision theory.

5. Applications

This section is devoted to economic applications of case-based decision theory. All we could hope to provide here are some sketchy illustrations, which certainly fall short of complete models. Our goal is merely to convince the skeptical reader that CBDT may have drastically different implications than EUT, and that it may be able to explain some phenomena better.

5.1 A Market for a Single Good

Consider a market for a single good, in which buyers and sellers meet to trade up to one unit each. Assume that the possible prices at which transactions may take place are $(1, \dots, 2N - 1)$ for some $N \geq 1$. (The choice of an odd integer is only a matter of notational convenience in the sequel.)

For simplicity of exposition we will assume that there are $2N - 1$ buyers and $2N - 1$ sellers whose reservation prices form linear supply and demand curves. Formally, let \bar{B} be a set of buyers and \bar{S} be a set of sellers, with $|\bar{B}| = |\bar{S}| = 2N - 1$, say $\bar{B} = \{1, \dots, 2N - 1\}$ $\bar{S} = \{2N + 1, \dots, 4N - 1\}$. For $i \in \bar{B}$ let $r_i = i$, to be interpreted as buyer i 's evaluation of a unit of the good. Similarly, for $j \in \bar{S}$ let $r_j = j - 2N$ be seller j 's valuation.

Trade takes place in stages. At each stage, each agent (seller or buyer) chooses an act out of (a subset of) $A = \{1, \dots, 2N - 1\}$, which is interpreted as an offer (an asking price) or a bid, accordingly.

It will be convenient to assume that buyer i is restricted to choose prices from

$$A_i = \{1, \dots, i\}$$

and seller j —from

$$A^j = \{j - 2N, \dots, 2N - 1\}$$

These "messages" are assumed to be submitted to a market mechanism, which matches buyers and sellers and attaches to each pair a price which is acceptable to both. Matched pairs leave the market and the process continues to the next stage.

Formally for (stage) $\ell \geq 0$ let $B_\ell \subseteq \bar{B}$ and $S_\ell \subseteq \bar{S}$ denote the buyers/sellers which are still in the marketplace at stage ℓ . ($B_0 = \bar{B}$, $S_0 = \bar{S}$). A market mechanism μ is a function, which maps quadruples of the form

$$(B, S, (b_i)_{i \in B}, (a_j)_{j \in S})$$

--where $B \subseteq \bar{B}$, $S \subseteq \bar{S}$ and $b_i, a_j \in A$ -- to subsets

$$T \subseteq B \times S \times A$$

such that

- (i) If $t_1 = (i_1, j_1, p_1)$, $t_2 = (i_2, j_2, p_2) \in T$ and $t_1 \neq t_2$, then $i_1 \neq i_2$ and $j_1 \neq j_2$;
- (ii) For all $(i, j, p) \in T$

$$a_j \leq p \leq b_i$$

- (iii) If $i \in B \setminus T_B$ and $j \in S \setminus T_S$ (where T_B and T_S stand for the projections of T on B and S , respectively), then

$$b_i < a_j$$

(Note that condition (iii) means that T is a maximal transaction set satisfying (i) and (ii) with respect to set inclusion. We do not necessarily assume that T is maximal in cardinality.)

Thus, if at stage j the buyers make bids $(b_j)_{i \in B_j}$ and the sellers—offers $(a_j)_{j \in S_j}$ —we define

$$B_{t+1} = B_t \setminus T_B$$

$$S_{t+1} = S_t \setminus T_S$$

for $T = \mu(B_t, S_t, (b_j)_{i \in B_j}, (a_j)_{j \in S_j})$.

We finally turn to describe the behavior of agents in this model. It will hardly be a surprise that we assume them to be case-based decision makers. However, we do not assume that they have actually traded in this market before. Rather, we equip them with a little knowledge of the world and with some hopes. Specifically, for each agent there will be a range of prices for which he/she has a "hypothetical case" in mind, reflecting the fact that the agent knows that, if trade occurs, he/she will get at least as good a price as he/she chooses to declare, that is, that trade is voluntary. The range of prices for which hypothetical cases are assumed is bounded by the agent's reservation price on the one hand, and his/her "aspiration level" on the other.

Formally, for all $i \in \bar{B}$ ($j \in \bar{S}$) we assume an (integer) aspiration level h_i (h_j) as given. Suppose that $0 \leq h_i \leq r_i$ and $0 \leq h_j \leq 2N - 1 - r_j$. At stage 0, buyer i 's (seller's j) memory is

$$M_i^0 = \{(bh_{p,p,r_i} - p) \mid r_i - h_i \leq p \leq r_i\}$$

$$(M_j^0 = \{(sh_{p,p,p} - r_j) \mid r_j \leq p \leq r_j + h_j\})$$

(The elements bh_p, sh_p denote some abstract "hypothetical problems".)

Note that the outcomes $(r_i - p; p - r_j)$ are given in terms of net surplus, with no reference to the aspiration levels. The latter, however, are reflected in the definition of the utility functions:

$$\begin{aligned} u_i(x) &= x - h_i \\ (u_j(x) &= x - h_j) \end{aligned}$$

for $x \in \mathbf{R}$.

Thus the aspiration level of each agent could be implicitly defined by the memory M_i (M_j): the set of cases the agent has imagined, or hoped for, contains all prices from his/her reservation value down (up) to a certain "best" price, which we take to be the aspiration level on the price scale. We therefore normalize the utility function so that it be zero for a transaction made at this price. (Similarly, the reservation prices r_i (r_j) can also be implicitly defined by M_i^0 (M_j^0).)

Finally, we have to describe the agents' similarity functions. We will assume that all hypothetical cases are only remotely similar to the real ones, and set

$$\begin{aligned} s_i(\ell, bh_p) &= \epsilon \quad \forall i, \ell, p \\ (s_j(\ell, sh_p) &= \epsilon \quad \forall j, \ell, p) \end{aligned}$$

where s_i (s_j) is the similarity function of buyer i (seller j) ℓ stands for the ℓ -th stage problem and $\epsilon \in (0, 1/2N)$.

Real cases, on the other hand, are similar to each other more than they are to the hypothetical ones:

$$s_i(\ell, \ell') = 1 \quad \forall i, \ell, \ell'$$

$$(s_j(\ell, \ell') = 1 \quad \forall j, \ell, \ell').$$

For simplicity we assume that each agent's memory is supplemented only by this agent's own experience. In particular, whenever the agent has made a bid (offer) of p at some stage ℓ , but failed to trade, a case

$$(\ell, p, 0)$$

will be added to $M_i(M_j)$. That is

$$M_i^{\ell+1} = M_i^\ell \cup \{(\ell, p, 0)\}$$

$$(M_j^{\ell+1} = M_j^\ell \cup \{(\ell, p, 0)\}.)$$

For brevity's sake, we omit some obvious formal definitions. However, we will freely refer to a "decision rule" (or a "strategy") for the agents, to the "process" defined by such rules and a market mechanism μ , and so forth.

We can now formulate the following.

Proposition 3: Let μ be any market mechanism. The unique (stagewise) U-maximizing decision rule for buyer i is the following: start at

$$b_i^0 = r_i - h_i$$

and, as long as no trade occurs (i.e., for ℓ such that $i \in B_\ell$),

$$b_i^\ell = (r_i - h_i) + [\ell \pmod{h_i} + 1]$$

Correspondingly, the unique U-maximizing decision rule for seller j is:

$$a_j^l = (r_j + h_j) - [\ell(\bmod h_j + 1)]$$

for all ℓ s.t. $j \in S_\ell$.

With these decision rules, no trade occurs after stage $4N^2$. (I.e., $B_\ell = B_{4N^2}$ and $S_\ell = S_{4N^2}$ for all $\ell > 4N^2$.)

Furthermore, there is no more possible trade at this point, that is, for all $i \in B_{4N^2}$ and $j \in S_{4N^2}$,

$$r_i < r_j.$$

(The proofs of all propositions appear in Appendix 1.)

It is worth noting that the market mechanism does not assure us that the outcome is Pareto-efficient. It is easy to see that if, say, different buyers have different aspiration levels, those who end up with the good may not be those with the highest reservation prices.

5.2 A Repeated Market for a Single Good

The analysis presented above, and especially the sub-optimality of the market mechanism, naturally call for repetition of the market game, with each round's successful buyers in the role of the successive round sellers.

However, the exercise we choose to do here is slightly different: we will assume that the whole "market game" of sub-section 5.1 is repeated, with the same agents in the same roles, and attempt to study the trading price behavior over time. In particular, we would like to see if these converge to the equilibrium price, and if so—how fast.

There are many ways to model the agents' learning from one stage to

another; each agent may recall his/her own transactions, as well as those of other agents. Furthermore, each agent's aspiration level may change over time as a function of both his/her and others' past performance. To simplify matters, we will assume that the only information that is transferred from one round to the next is the average price of the deals struck in the previous round, and that this price is reflected only in the agents' aspiration levels.

To minimize indices, we omit some of the obvious formal definitions. We assume that the game described in sub-section 5.1 is played repeatedly and denote a generic round by $t \geq 0$. Recall that each round consists of $4N^2$ stages, in each of which the market mechanism μ is operated once.

At stage $t = 0$ the aspiration levels h_i, h_j are arbitrarily given as above. For $t \geq 0$ let p_t be the average of the prices at which the good was traded in round t , rounded off to a closest integer. (This quantity is well-defined by Proposition 3.) Then for $i \in \mathcal{B}$ the aspiration level at round $t + 1$ is

$$h_i = \max\{0, r_i - p_t\}$$

and for $j \in \mathcal{S}$,

$$h_j = \max\{0, p_t - r_j\}$$

Thus, each agent adjusts his/her aspiration level, as if under the assumption that he/she can do just as well as the average buyer/seller.

The following result characterizes the price behavior:

Proposition 4: Let μ be any market mechanism and consider the process defined by μ , the U-maximizing decision rules (at each stage of each round) and the

aspiration levels defined above. Then the following hold:

- (i) If $p_t \geq N$ then $p_t \geq p_{t+1} \geq N$, and if $p_t \leq N$ then $p_t \leq p_{t+1} \leq N$.
- (ii) For $t \geq 1$, let $K = |N - p_{t-1}|$. Then at round t there are $(N - K)$ transactions at stage 0 and price p_{t-1} , and then there follow from K to $2K$ additional stages, at each of which there is one more transaction at consecutively decreasing (if $p_{t-1} > N$) or increasing (if $p_{t-1} < N$) prices.
- (iii) Let $K = |N - p_0|$. Then after at most N rounds p_t is fixed at a price p such that

$$|p - N| < \sqrt{N+1}$$

Furthermore, from that point on no round has more than $(2\sqrt{N+1} + 1)$ stages of trade.

Remarks

a. As will be clear from the proof, the number of rounds after which p_t is fixed can be more tightly bounded. (For instance, $\lfloor N/2 \rfloor + 1$ is also a bound.)

b. One may claim that it is not entirely realistic that the bids and offers are updated by 1 unit at each stage. Indeed, combining suggestions from subsections 3.4 and 4.3 above, we may consider a similarity function over acts, such that offering a price p will be similar to offering a price q to an extent

$$e^{-\theta|p-q|}$$

(For some $\theta > 0$.)

In such a model, a failure to trade at a price p will decrease the U -value of "similar" (close) prices as well, leading to a faster convergence to an almost-equilibrium price.

c. The model presented here may be extended to continuous prices as well. The "unit" by which the agents change their bids/offers should then be thought of as a "just noticeable difference."

d. As will be clear from the proofs, neither Proposition 3 nor 4 depend on the market mechanism μ being deterministic. That is, one may consider any random choice of buyer-seller matchings whose realizations satisfy the conditions defining market mechanisms, and the results would still hold.

e. Finally, note that the market behavior described by Proposition 4 appears to be in agreement with experimental results (see Plott (1982)). In particular, the nature and rate of convergence to equilibrium price may be better explained by this model than by traditional Bayesian equilibrium analysis with expected-utility-maximizing agents.

5.3 To Buy or Not to Buy

Consider a firm which is about to introduce a sequence of new products $(1, \dots, n)$ into a market. For simplicity let us assume that all consumers are identical. With the introduction of product i , they face a decision problem p_i , with two possible acts $\{a, b\}$ where "b" stands for buying the product and "a" for abstaining from purchase. To simplify matters even further, we assume that the prices are fixed (i.e., not considered a decision variable of the firm), and so are the quantities. For instance, $(1, \dots, n)$ may be food products. A consumer's decision to "buy" product i , say, a cereal or a soup, implies consumption on a regular basis in quantities which are (literally) naturally given.

Further assume that the representative consumer has already decided to buy product 0 by the same firm. For each pair $i, j \in \{0, 1, \dots, n\}$ a similarity function $s(i, j) = s(p_i, p_j)$ is given, which presumably reflects the perceived similarity between the products.

It is interesting to note that the order in which the products are introduced may make a difference. For instance, let $n = 2$ with the following similarity matrix:

$s(i, j)$	j	0	1	2
i				
0		1	.5	0
1		.5	1	.5
2		0	.5	1

For simplicity, assume that each consumer will derive a utility level 1 (with certainty) from each product consumed. Then if the firm introduces product 1 and then product 2, both will be purchased. However, if product 2 precedes 1, when it is introduced nothing which resembles it exists in memory. Thus the consumers' decision between a and b will be arbitrary and with (say) probability .5 it will not be consumed.

In this set-up it is quite obvious that an optimal policy for the firm will be the following: consider a directed graph whose nodes are the products, and each arc (i, j) has a weight $s(i, j)$. Find any (Hamiltonian) path in the graph which does not use zero-similarity arcs, and introduce the

products in this order.

More generally, however, it is not obvious what will an optimal policy look like—where each product has a potentially different distribution over utility levels and, more importantly, where competition with other firms is also present.

5.4 Reputation

Case-based consumer decisions give rise to aspects of reputation quite naturally. Consider a similar model to the one outlined above, but let us now assume that any two purchasing decisions are similar to a positive degree in the consumer's mind. On the other hand, let us now assume that in a given ("traditional") market of product 1 only firm A is operating. Product 2, by contrast, is a new product and both firms A and B are competing in it. Other things being equal, firm A will have an edge in market 2 if it satisfies consumers' expectations in market 1 (i.e., $U(A) > 0$). Thus one would expect successful firms to enter new markets even if the technology needed in them is completely different from that used in the traditional ones.

Expected utility theory can, of course, also explain the role of reputation in the context of equilibrium between the firms. We find, however, that CBDT makes much weaker rationality assumptions in explaining this phenomenon.

5.5 Introductory Offers

Another phenomenon which is close in nature is the introduction of new products at discounted rates. Again, one may explain the optimality of such marketing policies with "fully rational" expected utility consumers. For instance, if there is some cost to experimentation and/or risk aversion, a fully rational consumer may tend to buy the product at the regular price after

having bought it at the introductory (lower) price. Yet with case-based decision makers (as consumers) the formation of habits is a natural feature of the model.

6. CBDT and EUT

We devote this section to a few comments on the comparison between expected utility and case-based decision theories.

6.1 It seems worthwhile to emphasize that we do not consider case-based decision theory (CBDT) "better" than or a substitute for expected utility theory (EUT). We simply view them as complementary theories. In problems involving probabilities, for instance, it is neither realistic nor recommended to ignore them. Similarly, in case "states of the world" are naturally defined, it is likely (and certainly desirable) that they be used in a decision maker's reasoning process, even if a (single, additive) prior cannot be easily formed.

However, when neither probabilities nor states of the world are salient features of the problem, we believe that CBDT may capture some aspects EUT tends to ignore.

We may thus refine Knight's distinction between "risk" and "uncertainty" by introducing a new category of "ignorance": "risk" refers to situations where both states of the world and probabilities on them are given; "uncertainty"—to situations in which states are naturally defined, but probabilities are not. Finally, "decision under ignorance" refers to decision problems for which states are not defined, let alone probabilities. EUT is appropriate for decision making under risk. In face of uncertainty one may still use generalizations of EUT such as non-additive probabilities (Schmeidler (1989)) and multiple-priors (Gilboa-Schmeidler (1989) and Bewley

(1986)). However, in cases of ignorance CBDT is a viable alternative to the EUT paradigm.

6.2 The mathematical similarity between the theories—the fact that CBDT may be viewed as a paraphrase of EUT—is also reflected in the ontological status of the terms employed. "Similarity" in CBDT plays a similar role to "probability" in subjective EUT. Both are in principle measurable, though in practice evasive. But both play an important role in capturing some essential features of the decision making process, and can, with a certain amount of arbitrariness, be used in applications.

6.3 The classical derivation of EUT, as well as the derivation of CBDT in this paper, are behavioral in that the theoretical constructs in these models are induced by (in-principle) observable choices. Yet the scope of applicability of these theories may be more accurately delineated if we attempt to judge the psychological plausibility of the various constructs (at least by direct introspection).

Attempting to do so, one may try to classify decision problems according to their novelty and the amount of reasoning they require. On one hand there are almost automated decisions, such as saying "Hi" to a person when (s)he enters the room. On the other extreme one may consider weighty decision in unfamiliar situations, such as getting married or investing in a politically unstable country. Somewhere in between these extremes one finds decisions which do require some deliberation, but for which historical data do exist, such as "solid" investments.

We would like to suggest a tentative classification, according to which CBDT is useful at the extremes of this scale and EUT in the middle: automated decisions are done by CBDT in the guise of "rules"; when deliberation is

required, but enough data exists for the formulation of a prior or, at least, states-of-the-world model, EUT (and its generalizations) seem both appropriate and realistic; finally, when none of those exist, as is often the case in novel situations, CBDT is again a more accurate description of decision making, as well as a more realistic "rationality goal" to aspire to.

6.4 Another distinction between EUT and CBDT, which is related to the previous point, is the following: in EUT, as in decision theory at large, the unit of analysis is typically a single decision problem. That is, a problem which is considered in isolation. In CBDT, by contrast, the decision problem is studied against the background of other problems. The preference order, as well as the whole analysis, is history-/memory-/context-dependent.

We find that a decision problem under uncertainty may indeed be studied "in isolation" if it has been repeatedly encountered in the past in very much the same form. Past experience with the same problem allows one to define states of the world and perhaps even a prior. However, when the problem is "new" to the DM, there is not enough data to formulate all the relevant states of the world, let alone to form a prior over them. (See also the following comment.)

Relating this point to our discussion of experimentation in subsection 3.3 above, one may contend that if a problem repeats itself—i.e., if $s(\bullet, \bullet) = 1$ —a "rational" DM, that is, one who is aware of this fact, may wish to experiment intentionally, and EUT may be a better description of the DM's behavior than CBDT. However, if there is no repetition, or if the DM is only boundedly rational (and is not aware of the long-run benefits of experimentation), CBDT may prove more realistic than EUT.

6.5 The classical EUT maintains that no loss of generality is involved

in assuming that the states of the world are known. Indeed, one may always define the states of the world to be all the functions from all conceivable acts to all conceivable outcomes. Such a definition is tautological and makes no implicit assumptions about the "reality" modeled.

This view is theoretically very appealing, but suffers from several drawbacks:

- (i) the set of conceivable acts (which are all the functions from states to outcomes) is much larger than the actual ones the DM can choose from. Indeed, assume one starts from a set of acts A and a set of outcomes X . The states of the world are X^A , i.e., all the functions from acts to outcomes. The set of conceivable acts will be $\bar{A} = X^{X^A}$, that is, all functions from states of the world to outcomes. Hence the cardinality of the conceivable acts \bar{A} is by two orders of magnitude larger than that of the actual sets A . Yet using a model such as Savage's, one needs to assume a (complete) preference order on \bar{A} , while in principle only a preference order on A is given. Differently put, such a "canonical construction" of the states of the world gives rise to preferences which are intrinsically hypothetical and is a far cry from the behavioral foundations of Savage's original model;
- (ii) even if one uses a "canonical" set of states of the world, there is no "canonical" way to attach a prior to such a set. Thus it is far from clear that the (mostly hypothetical) preference relation over acts in \bar{A} will satisfy axioms such as Savage's. Nor is it obvious how should one construct a prior on this set for normative applications of the theory;
- (iii) finally, as mentioned in the Introduction, the humongous model which results from such a construction can be hardly claimed to be

psychologically plausible or practically implementable.

6.6 As mentioned above, EUT and CBDT may be viewed as purely behavioral theories; that is, they may describe a DM's choices without claiming that the DM is aware of the decision process or the primitives of the model (such as states of the world or cases, probabilities or similarities etc.). Yet when viewed as normative theories, or when judged for psychological plausibility as descriptive ones, one has to assume that the DM actually thinks in terms of these models.

From this viewpoint there is a crucial difference between the two: CBDT, as opposed to EUT, does not require the DM to think in hypothetical or counterfactual terms. In EUT, whether explicitly or implicitly, the DM considers states of the world and reasons in propositions of the form "If the state of the world were ω and I choose a then r would result." In CBDT, by contrast, the (implicit) conditional statement takes the form "If I choose a , r results." That is, the conditional statement depends only on the DM's choices and not on the counterfactual choice of "Nature."²

Similarly, there is a difference between EUT and CBDT in terms of the informational requirements they entail: to "implement" EUT, one needs to "know" the utility function u , i.e., its values for all consequences which may result from some acts. For CBDT, on the other hand, it suffices to know the u -values of those outcomes which were actually experienced.

²Indeed, the axiomatic derivations of both EUT and CBDT require "hypothetical choices." In CBDT, however, these are only assumed on the modeler's part, not the DM's. That is to say, the DM will be asked hypothetical questions only when the modeler tries to elicit the similarity and utility functions from "observed" choices. If these functions are assumed as primitives, or if enough historical data exists, these hypothetical questions need not be asked. In EUT, on the other hand, the very definition of an "act" à la Savage is a list of hypothetical conditional statements, and thus the DM has to reason in counterfactuals.

7. Concluding Remarks

7.1 A few words on the normative interpretation of CBDT are in order. It seems undeniable that CBDT does not have the same normative appeal that EUT has. A CB decision maker (CBDM) does not "consider" all choices, does not "attempt" to "maximize utility" and so forth.

Yet we would like to suggest a defense of CBDT as a "second-best normative theory." In cases where EUT's recommendations are blatantly impractical, its normative appeal is also tainted. And then there is also room, normatively speaking, for other, less idealized theories.

7.2 We will not dwell here on potential normative implications of CBDT. Let us only briefly mention that, if we accept CBDT as a second-best theory then perhaps the second-best may be improved. For instance, one may try to change one's similarity function so that it be symmetric, ignore primacy effects and so forth. It may even be argued that it is more useful to train professionals (doctors, managers, etc.) to make efficient and probably less biased CB decisions rather than to teach them very rational but impractical EUT.

However, we decline to make any such claim here (or elsewhere), and it is given merely as an illustration of possible implications.

7.3 It may seem that our CBDM's are extremely conservative and boring creatures: whenever an act achieves their aspiration level, they would stick to it. A CBDM, it would seem, is an animal which always eats the same food at the same place, chooses the same form of entertainment (if at all) and so forth.

Although this is true at some level of description, it does not have to be literally true: as most entities in theoretical models, "acts" are

language dependent. Thus an act—which is, say, chosen over and over again—need not be "Have lunch at X"; it may also be "Have lunch at a place I did not visit this week." Repetition in this level of description will obviously be translated to an extremely diverse lunch pattern.

7.4 In a more general model, one may try to capture manipulations of the similarity function. In phenomena as diverse as advertising and legal argumentation, people try to influence other peoples' perceived similarity of cases. Along similar lines, the similarity and/or the utility functions for existing memory may change due to new information. Learning new facts may, for instance, make the DM aware of some similar or distinguishing features of past cases, or change the evaluation of past results. These research directions are, however, beyond the scope of this paper.

7.5 It goes without saying that CBDT, especially in its descriptive interpretation, may greatly benefit from additional psychological insights into the structure of memory and the evolution of aspiration levels. For instance, one may hypothesize that the "satisficing" nature of decision making is revealed not only in a dynamic context, but also within each decision: rather than computing the U-value of all possible acts, it may be more realistic to suggest that the DM stops at the first act which obtains a positive U-value. There are, however, several ways in which "first" could be defined. For example, the DM may ask him/herself, "When did I choose this act?", and only after the evaluation of a given act will the next one be considered. Alternatively, the DM may focus on the problem and ask "When was I in similar situation?", and as the cases are retrieved from memory one by one, the function U is updated for all acts—until one act exceeds the aspiration level. Since these two crude models already induce different

decision rules, it is obvious that CBDT may be refined and improved by incorporating empirical findings regarding the recollection process.

Similarly, the process by which the aspiration level $u(r_0)$ is adjusted (as mentioned in subsection 3.1 above) calls for further study.

7.6 The model we present here should be taken merely as a "first approximation." Especially in view of some descriptive failures of the monumental expected utility theory, case based decision theory could hardly be expected to fare any better. Thus it makes sense that for some applications CBDT will have to be generalized to allow for semi-orders (see Luce (1956)), non-additive measures on the space of cases (see Schmeidler (1989)) and so forth.

Our main goal in this paper was to explore the possibility of a formal, axiomatically-based decision theory which uses a different, less "rational" but at times more realistic paradigm than EUT. We believe that case-based decision theory may be such an alternative.

Appendix 1: Proof of Propositions1. Proposition 1

Fix p, H and denote $\succeq = \succeq_{p, H}$. W.l.o.g. assume $H \neq \emptyset$. First note the following.

Observation: If \succeq satisfies A1 and A4, then:

- (i) For all $x, y \in X$ with $x * y = 0$,

$$x \succeq y \Leftrightarrow -y \succeq -x.$$
- (ii) For all $x, y, z, w \in X$ with $x * y = 0$, $(x + z) * (y + w) = 0$ and $z \sim w$,

$$x \succeq y \Leftrightarrow (x + z) \succeq (y + w).$$

Proof: (i) Assume $x \succeq y$. Consider $z = w = -(x + y)$ and use A4 (where $z \succeq w$ follows from A1).

- (ii) Under the provisions of the claim, $z \succeq w$ and A4 implies

$$x \succeq y \Rightarrow (x + z) \succeq (y + w).$$

As for the converse, define $z' = -z$, $w' = -w$. By (i), $-z \succeq -w$ and A4 can be used again to conclude $x \succeq y$. ■

We now turn to the proof of Proposition 1. Define $\succeq' \subseteq X \times X$ by

$$x \succeq' y \Leftrightarrow x - y \succeq 0$$

for all $x, y \in \mathbb{R}^n - X$.

We state without proof the following facts:

Fact 1: For $x, y \in X$ with $x * y = 0$,

$$x \succeq y \Leftrightarrow x \succeq' y.$$

Fact 2: \succeq' is complete, i.e., for all $x, y \in X$ $x \succeq' y$ or $y \succeq' x$.

Fact 3: \succeq' is transitive.

Fact 4: \succeq' is monotone, i.e., for all $x, y \in X$, $x \geq y$ implies $x \succeq' y$.

Fact 5: \succeq' is continuous, i.e., for all $x \in X$ the sets

$$(y | y \succeq' x), (y | x \succeq' y)$$

are closed in \mathbf{R}^n . (In view of the above, this is equivalent to the sets

$$(y | y \succ' x), (y | x \succ' y)$$

being open.)

Fact 6: \succeq' satisfies the following separability condition: for all $x, y, z, w \in X$, if $x \succeq' y$ and $z \succeq' w$ then $(x + z) \succeq' (y + w)$.

Fact 7: If $x \succeq' y$, then

$$x \succeq' (1/2)(x + y) \succeq' y.$$

Furthermore, if $x \succ' y$, then

$$x \succ' (1/2)(x + y) \succ' y.$$

Fact 8: For every $y \in X$, the sets

$$(x | x \succ' y), (x | x \succeq' y)$$

$$(x | y \succ' x), (x | y \succeq' x)$$

are convex.

Fact 9: Define

$$A = \{x | x \succeq' 0\}$$

$$B = \{x | 0 \succ' x\}$$

(where 0 denotes the zero vector in \mathbf{R}^n).

Then A is closed and convex, B is open and convex, $A \cap B = \emptyset$ and

$$A \cup B = \mathbf{R}^n.$$

Fact 10: If $B = \emptyset$, the function $s(\bullet) = 0$ satisfies the representation condition. If, however, $B \neq \emptyset$, there exist a linear functional $S: \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$S(x) \geq 0 \quad \text{for all } x \in A$$

$$S(x) < 0 \quad \text{for all } x \in B.$$

The function $s: H \rightarrow \mathbb{R}$ defined by it is nonnegative (by Fact 4) and satisfies the desired representation for \succeq' on \mathbb{R}^n .

Note that Facts 1 and 10 complete the proof. ■

Notice that additional standard results would also hold here: (1) Proposition 1 can be strengthened to be an "iff" statement; (2) under some appropriate non-triviality conditions, the similarity function is unique up to a multiplicative positive constant and the utility function (here implicitly assumed to be the identity) is unique up to a positive linear transformation.

2. Proposition 2

It follows from Proposition 1 that the function $s_{p,H}: H \rightarrow \mathbb{R}_+$ defines, for each $q_1, q_2 \in H$ with $s_{p,H}(q_2) > 0$ a unique relation

$$\delta_{p,H}(q_1, q_2) = \frac{s_{p,H}(q_1)}{s_{p,H}(q_2)}$$

in the sense that any other similarity function would give rise to the same ratio $\delta_{p,H}$ both in terms of its domain and in terms of its values.

It is easy to see that A5 implies that this ratio is independent of H . That is, as long as $H_1, H_2 \subseteq P$ satisfy $q_1, q_2 \in H_1, H_2$ and $p \notin H_1, H_2$ we obtain

$$\delta_{p,H_1}(q_1, q_2) = \delta_{p,H_2}(q_1, q_2) = \delta_p(q_1, q_2).$$

Moreover, notice that whenever p , q_1 and q_2 are three distinct problems, there are subsets $H \subseteq P$ such that $q_1, q_2 \in H$, $p \in H^c$.

We now turn to define the similarity function s . Fix $p \in P$. If for all $H \subseteq P \setminus \{p\}$ and all $q \in H$, $s_{p,H}(q) = 0$, set $s(p, \bullet) = 0$. Otherwise, choose $q \in H$ such that $s_{p,H}(q) > 0$ and set

$$s(p, q) = 1.$$

For any other $q' \neq p$ define

$$s(p, q') = \delta_p(q', q).$$

(Notice that under A5 $\delta_p(q', q)$ is well-defined for all q' .)

We need to show that for every $H \subseteq P \setminus \{p\}$ there exists a constant $\alpha > 0$ such that

$$s(p, q') = \alpha s_{p,H}(q') \text{ for } q' \in H.$$

Applying A5 with $q_1 = q_2 = q'$ we conclude that both sides of the equality vanish together. (This concludes the proof for the case $|H| = 1$.)

First consider H such that $q \in H$. It follows from our construction (and from A5) that $\alpha = [s_{p,H}(q)]^{-1}$ will do.

Next assume that $q \notin H$. Considering $q', q'' \in H$, we need to show that

$$\frac{s(p, q')}{s(p, q'')} = \frac{s_{p,H}(q')}{s_{p,H}(q'')}$$

whenever these are well-defined.

By definition of δ ,

$$\frac{s_{p,H}(q')}{s_{p,H}(q'')} = \delta_p(q', q'').$$

On the other hand,

$$\frac{s(p, q')}{s(p, q'')} = \frac{\delta_p(q', q)}{\delta_p(q'', q)}.$$

Consider $H_0 = H \cup \{q\}$. Using the definition of δ we may write

$$\delta_p(q', q) = \frac{s_{p,H_0}(q')}{s_{p,H_0}(q)}; \quad \delta_p(q'', q) = \frac{s_{p,H_0}(q'')}{s_{p,H_0}(q)}.$$

Hence,

$$\frac{\delta_p(q', q)}{\delta_p(q'', q)} = \frac{s_{p,H_0}(q')}{s_{p,H_0}(q'')} = \delta_p(q', q'').$$

Finally, since P is finite we may normalize the function s such that it takes values in $[0,1]$. This concludes the proof of Proposition 2. ■

We note that it is straightforward to check that A5 is also a necessary condition for the conclusion of Proposition 2.

3. Proposition 3

First consider the optimal (i.e., U -maximizing) strategies. Consider a buyer who has reservation value r_i and aspiration level a_i (the proof for a seller is symmetric.) At stage 0, an act p , $r_i - h_i \leq p \leq r_i$ is evaluated by

$$U(p) = (r_i - p - h_i)\epsilon$$

whence the maximizer is $p = r_i - h_i$.

Assuming no trade was made, at stage 1 this act is evaluated by

$$U(p) = (r_i - p - h_i)\epsilon - h_i = -h_i$$

where for $p = r_i - h_i + 1$ the value is

$$\begin{aligned} U(r_i - h_i + 1) &= [r_i - (r_i - h_i + 1) - h_i]\epsilon = \\ &= -\epsilon > -h_i \end{aligned}$$

and it becomes the next maximizer. Proceeding in this fashion, one proves that all acts will be tried once, and after they all failed once the problem is identical (up to a shift in h_i units) to the original one.

To complete the proof of the proposition it suffices to show the "furthermore" part, i.e., that after stage $4N^2$, $r_i < r_j$ for all $i \in B_t$ and $j \in S_t$. Suppose that this were not the case, and $r_i \geq r_j$ for some $i \in B_t$, $j \in S_t$. i 's bids cycle along $(r_i - h_i, r_i - h_i + 1, \dots, r_i)$ while j 's offers cycle along $(r_j + h_j, r_j + h_j - 1, \dots, r_j)$. At stages of the form $k(h_i + 1) - 1$ (for some $k > 0$), i is bidding r_i . At stages of the form $m(h_j + 1) - 1$ (for some $m > 0$), j is offering r_j . Thus at stage $(h_i + 1)(h_j + 1) - 1$, if not sooner, i and j can be matched. Since they were not matched in the first $4N^2 \geq (h_i + 1)(h_j + 1)$ stages, we arrive at a contradiction. ■

4. Proposition 4

Let us first prove (ii). Let $p_{t-1} = N + K$ and assume, w.l.o.g. that $K > 0$. (The case $K = 0$ is trivial, while $K < 0$ is symmetric.)

The buyers $\{1, \dots, N + K - 1\}$ have zero aspiration levels, and they will be bidding their reservation prices at all stages. Similarly, the sellers

$\{3N + K + 1, \dots, 4N - 1\}$ have zero aspiration levels, and they will be offering their reservation values. (As opposed to the zero-aspiration buyers, however, all these sellers will not trade.)

$(N - K)$ buyers, i.e. $\{N + K, N + K + 1, \dots, 2N - 1\}$ bid $p_{t-1} = N + K$, while $N + K$ sellers, namely $\{2N + 1, \dots, 3N + K\}$, offer $N + K$. Thus any mechanism μ will match the $(N - K)$ buyers with $(N - K)$ sellers. Let L be the set of matched sellers.

Note that after stage 0 the buyers cease to be an interesting partner: only zero-aspiration buyers are left, and they keep bidding their reservation values. The sellers, on the other hand, all decrease their offers by 1 (apart for $3N + K$, who is already at his/her reservation price.) Thus at stage 1 one more transaction will take place, this time at price $N + K - 1$.

The process continues in this fashion until all sellers have offered their reservation prices at least once. (A seller who has gone into a second cycle will not trade any more since all bids are below his/her reservation price.) The number of the trade stages depends on the choice of sellers by the mechanism μ : suppose, for instance that

$$L = \{3N + 1, \dots, 3N + K\}$$

i.e., that the sellers who get to trade at stage 0 were all with reservation prices above the equilibrium price N . Thus, the "cheap" sellers are still in the market and the process may continue even below the equilibrium price. However, the maximal number of stages is $2K + 1$: after stage $2K$ the total number of units traded is

$$(N - K) + 2K = N + K = p_{t-1}$$

i.e., it is equal to the total supply at price p_{t-1} . (Recall that the sellers $(3N + K + 1, \dots, 4N - 1)$ will never offer the good at a price lower than $3N + K + 1$.)

On the other hand, the minimal number of stages will be obtained by a mechanism which "uses up" the "cheap" sellers as soon as possible, and, in particular, sets

$$L \subseteq \{2N + 1, \dots, 3N\}$$

However, since there are N sellers who are willing to trade at price N , at least K stages of trade will occur. This concludes the proof of (ii).

It is now quite simple to verify that (i) holds: assume, again, that $p_t = N + K$ and (w.l.o.g.) $K \geq 0$. Suppose that μ determines $J + 1$ stages of trade, where $K \leq J \leq 2K$. Then

$$p_{t+1} = [(N - K)p_t + \sum_{j=1}^J (p_t - j)] / [(N - K) + J]$$

Obviously, $p_{t+1} \leq p_t$. To verify that $p_{t+1} \geq N$, notice that, since $J \leq 2K$,

$$(1/J) \sum_{j=1}^J (p_t - j) \geq N$$

Finally, we turn to prove (iii). Again consider $p_t = N + K$ with $K \geq 0$. Notice that the upper bound on p_{t+1} is obtained if trade occurs in precisely $J = K + 1$ stages. Then it is readily verified that the average price is

$$N + K - (K^2 + K)/2N$$

Thus p_{t+1} which is the (rounded-off) price will be strictly smaller than

p_t as long as

$$(K^2 + K)/2N > 1$$

or

$$K > \frac{\sqrt{4N+1} - 1}{2}$$

Hence the price will be fixed only if

$$K \leq \frac{\sqrt{4N+1} - 1}{2} < \sqrt{N+1}$$

It follows from part (ii) that the maximal number of trade stages is $2K + 1$, hence it is bounded by

$$2\sqrt{N+1} + 1$$

Finally, let us address the question of convergence. Assume that $p_0 = N + K$ ($K \geq 0$). From the foregoing analysis it is clear that

$$p_{t+1} \leq p_t - 1$$

at least as long as

$$|p_t - 1| > \sqrt{N+1}$$

Hence, in view of (i), at most N rounds are needed for p_t to converge. ■

Remark: It is easy to see that one can find better bounds. For instance, defining α_t by

$$p_t = N + \alpha_t N \quad 0 \leq \alpha_t \leq 1$$

the upper bound on p_{t+1} yields

$$\alpha_{t+1} \leq \alpha_t(1 - \alpha_t/2)$$

Thus, if $p_0 = 2N$, p_1 is bounded by $(N + N/2)$, p_2 is bounded by $(N + 3N/8)$ and so forth. Correspondingly, one may prove that the number of rounds (before the price becomes fixed) is bounded by any of the following:

N

$N/2 + 1$

$3N/8 + 2$

$39N/128 + 3$

and so forth

Appendix 2: An Alternative Model

In this appendix we outline the axiomatic derivation of case-based decision theory with act-dependent similarity function. We also show that V-maximization can be axiomatically derived.

We assume that the set P , A , R , C and M are defined and interpreted as in Section 2. Given a problem $p \in P$ and memory $M \subseteq C$ we define

$$E = E(M) = \{(q, a) \mid \exists r \in R, (q, a, r) \in M\}$$

$$H = H(M) = \{q \in P \mid \exists a \in A, (q, a) \in E\}$$

and

$$B = B(M) = \{a \in A \mid \exists q \in P, (q, a) \in E\}.$$

For each $a \in B$ denote $H_a = \{q \in H \mid (q, a) \in E\}$ and let F_a be the set of hypothetical acts

$$F_a = \{x: H_a \rightarrow \mathbf{R}\}.$$

(Again, we identify the set of outcomes R with the real line and implicitly assume it is measured in "utils.")

We will assume $|B| \geq 2$ and define

$$F = \bigcup_{a \in B} F_a.$$

For every p, E (with $|B| \geq 2$) we will assume that $\succ_{p, E} \subseteq F \times F$ is a binary relation satisfying the following axioms. For simplicity of notation, the subscripts " p, E " will be dropped whenever possible.

A1' Order: \succeq is reflexive and transitive, and for every $a, b \in B$, $a \neq b$, $x \in F_a$ and $y \in F_b$, $x \succeq y$ or $y \succeq x$.

A2' Continuity and Archimedeanity: For every $a, b \in B$, $a \neq b$ and every $x \in F_a$, the sets

$$(y \in F_b | y \succ x), (y \in F_b | x \succ y)$$

are non-empty and open (in F_b , endowed with the standard topology).

A3' Monotonicity: For every $a, b \in B$, $a \neq b$, $x, z \in F_a$, $y \in F_b$,

if $x \geq z$ then

$$z \succeq y \text{ implies } x \succeq y$$

and

$$y \succeq x \text{ implies } y \succeq z.$$

A4' Separability: For every $a, b \in B$, $a \neq b$, $x, z \in F_a$, $y, w \in F_b$, if $z \succ w$ then $(x \succeq y) \Leftrightarrow (x + z) \succeq (y + w)$.

Proposition A2.1: \succeq satisfies A1'–A4' iff there exists a function $s: H \rightarrow \mathbf{R}_+$ with $\sum_{q \in H_a} s(q) > 0 \forall a \in B$ such that for all $a \neq b$, $x \in F_a$, $y \in F_b$

$$x \succeq y \Leftrightarrow \sum_{q \in H_a} s(q)x(q) \geq \sum_{q \in H_b} s(q)y(q).$$

Notice that A1' requires that \succeq be transitive, which implies that acts belonging to the same space F_a be comparable. However, if $|B| \geq 3$ one may start out by assuming that transitivity holds only if all pairs compared belong to different spaces, and then consider the transitive closure of the

original relation.

Proof (Outline): Fix $a \in B$ and consider the restriction of \succeq to F_a . It is easy to see that on F_a \succeq is complete (hence a weak order), continuous and monotone (in the weak sense, i.e., $x \geq z$ implies $x \succeq z$).

Finally, if $x, y, z \in F_a$, we get

$$x \succeq y \text{ iff } (x + z) \succeq (y + z).$$

We therefore conclude that for every $a \in B$ there is $s_a: H_a \rightarrow \mathbf{R}_+$ such that for all $x, y \in F_a$

$$x \succeq y \iff \sum_{q \in H_a} s_a(q) x(q) \geq \sum_{q \in H_a} s_a(q) y(q).$$

Furthermore, by A2, \succeq is non-trivial on each F_a , whence for some $q \in H_a$, $s_a(q) > 0$.

Next, let 1_a denote the element of F_a consisting of 1's only ($1_a(q) = 1$). Fix $a \in B$ and for each $b \in B$ let δ_b satisfy

$$1_a \sim \delta_b 1_b.$$

Define $s: H \rightarrow \mathbf{R}_+$ by

$$s(q) = [\delta_b \sum_{q' \in H_b} s_b(q')]^{-1} s_b(q)$$

for all $q \in H_b$.

Finally, it is straightforward to see that this similarity function satisfies the desired representation condition and that the axioms are also

necessary. ■

To obtain the representation by the functional V in (**), consider the following axiom.

A6 (Experience Invariance): For all $a, b \in B$, $l_a \sim l_b$.

Without judging its reasonability, let us note that A6 means that the "quality" of the experience is all that matters, rather than its "quantity."

Finally, one may further demand that the following be satisfied.

A7 (Constant Similarity): For every $a, b \in B$ $a \neq b$, $q, q' \in H_a$, $x \in F_b$,

$$l_q \succeq x \text{ iff } l_{q'} \succeq x$$

where $l_q, l_{q'}$ stand for the corresponding unit vectors in F_a .

Obviously, A1'-A4', A6 and A7 will give rise to frequentist expected utility.