UTILITY THEORY WITHOUT THE COMPLETENESS AXIOM

Robert J. Aumann

Econometric Research Program
Research Memorandum No. 26
April 3, 1961

Princeton University
Econometric Research Program
92-A Nassau Street
Princeton, N. J.

One of the basic postulates of the von Neumann-Morgenstern utility theory is the completeness[1] postulate. It is assumed that any two outcomes are comparable, that there is preference or indifference[2] between any pair of outcomes. This is perhaps the most questionable of all the assumptions of utility theory. Von Neumann and Morgenstern themselves say that "it is very dubious, whether the idealization of reality which treats this postulate as a valid one, is appropriate or even convenient." [5,p.630]. Here we present a variation of the von Neumann-Morgenstern theory which makes no use of the completeness postulate.

The fact that up to the present no utility theory has been able to dispense with the completeness axiom[3] is no accident, but is due to the definition of utility that is usually accepted at present. In fact, no theory using this definition is possible without the completeness axiom. Under this definition, the utility is a function u from the outcome space to the real numbers which faithfully represents the preference order; to put it in symbols, we must have $u(x) > u(y)$ if and only if $x \succ y$

---

[1]Sometimes called "connectedness." This postulate is assumed also in the non-numerical indifference-curve approach to utility.

[2]Not to be confused with incomparability: Indifference between two possible outcomes involves a positive decision that it is immaterial whether the one or the other is chosen, whereas incomparability means that the decision maker refuses to decide between them. Indifferent activity vectors are comparable in the preference order, incomparable ones of course not.

[3]Von Neumann and Morgenstern mention very briefly [5,p.29] that if the completeness axiom is dropped, "a mathematical theory . . . is still possible. It leads to what may be described as a many-dimensional vector concept of utility. This is a more complicated and less satisfactory set-up." Details were never published. Professor Morgenstern has informed me that what they had in mind was not Hausner's multidimensional utility [2], but some kind of mapping into a partially ordered euclidean space (Hausner's mapping is into a completely ordered euclidean space).

(x preferred to y). Since the condition is necessary and sufficient and the real numbers are completely ordered, the original outcome space must also be completely ordered.

A number of ways out suggest themselves. One is to let the range of u be some canonical partially ordered space rather than the real numbers, in the spirit of footnote 3; we will not pursue this idea further. Another possibility is to relax the demand that the representation be faithful. More precisely, we shall demand that $x \succ y$ imply $u(x) > u(y)$, but not conversely.

Although such a utility will not have the uniqueness properties of the von Neumann-Morgenstern utility, it will have many of the other useful properties. For example, we can solve maximization problems with it: Maximization of such a "one-way" utility over a given constraint set will always lead to a maximal[4] element of the constraint set; conversely, for every maximal element x there is a utility whose maximization leads to x. Following up an idea of Shapley [6], we can even set up a theory of games in which the outcomes to an individual player are only partially ordered. Just as ordinary utilities are used to give a numerical treatment of games in which the outcomes are totally ordered (for each individual player), we will be able to use the "one-way" utilities to give a numerical treatment of such partially ordered games. In fact we will be able to "solve" these games in a manner analogous to the completely ordered case—obtaining an analogue of saddle points for zero-sum games, and of Nash equilibrium points [4] for general games.

---

[4] I.e., an element to which no other element in the constraint set is preferred. We cannot expect to get a maximum element—i.e., an element preferred or indifferent to all others in the constraint set—because such an element may not exist.

We remark that the present theory is a genuine generalization of the von Neumann-Morgenstern theory, in the sense that in case the outcome space does happen to be completely ordered, our utilities are the same as the von Neumann-Morgenstern utilities.

There is a significant parallel between the notion of utility proposed here and the notion of representation as used in algebra. A representation of a group G is a homomorphism u from G to some fixed canonical group H, such as a group of matrices or the group of rationals modulo 1. The range group H generally has more "structure" than the domain group G; there are more relations in H than in G (for example H may be commutative and G not). Thus $x = yz$ always implies $u(x) = u(y)u(z)$, but the converse is usually false. Although they are homomorphisms rather than isomorphisms, representations are useful because they enable us to use our knowledge about the structure of the range group H as a tool in studying the domain group G. Moreover, we can often say a lot about G by considering the set of all representations of G in H, rather than a single one. For example, this set often determines the structure of G completely.

The parallel with our situation is clear. Here G is the space of outcomes, H the real line; we are familiar with the structure of H, but not with that of G. An "order-isomorphism" (what we called a faithful representation above) is obviously out of the question if G is not completely ordered. We seek an "order-homomorphism" u, a mapping which gives us for each relation in G a corresponding relation in H, but not vice versa. Even a single mapping of this kind gives us important information on the original preference order; and as might be expected, the set of all such utilities gives us even more information.

In fact, we shall see that the preference order on the outcome space is to a large extent determined by the set of all utilities.

We have made an attempt to concentrate the less technical part of the paper in the first five sections; the remaining sections become progressively more technical.

1. The Outcome Space: Formal Assumptions

The space on which the utility will be defined is called a mixture space: intuitively, this is the set of all probability combinations of a set of "sure outcomes" or "pure prospects." Formally, it is a space X with a convex structure; that is, if $\{\gamma_1, \ldots, \gamma_k\}$ is a set of probabilities (i.e. $\gamma_i \geq 0$, $\Sigma\gamma_i = 1$), and if $x^1, \ldots, x^k \in X$, then there is defined in X the convex combination $\Sigma_{i=1}^k \gamma_i x^i$. One operates with these combinations in all ways like with ordinary vector-space sums, keeping in mind only that the coefficients must always be non-negative and sum to unity.[5] A set $\{x^1, \ldots, x^k\}$ of members of X is said to be independent if no two distinct combinations of the x^i are equal, and a maximal independent subset of X is said to span X. We shall assume in the sequel that X has a finite spanning subset, or in other words that it is finite dimensional. In particular, this condition will always be satisfied when only a finite number of "sure outcomes" are possible.

We assume that on our mixture space X there is defined a transitive and reflexive relation called preference-or-indifference and denoted by $\gtrsim$. If $x \gtrsim y$ and $y \gtrsim x$ we will say that x is indifferent

[5] For a set of formal axioms for a mixture space, see [2,p.169]. The treatment here is similar to that of [3].

to y and write $x \sim y$; if $x \gtrsim y$ but not $x \sim y$, we will say that x is preferred to y and write $x \succ y$. We assume that the following conditions hold:

(1.1) if $0 < \gamma < 1$, then $x \gtrsim y$ if and only if $\gamma x + (1-\gamma)z \gtrsim \gamma y + (1-\gamma)z$

(1.2) if $\gamma x + (1-\gamma)y \succ z$ for all $\gamma > 0$, then not $z \succ y$.

Assumption (1.2) is the "archimidean" or "continuity" assumption.

The relation $\gtrsim$ will be called a partial order; the space together with the partial order $\gtrsim$ will be called a partially ordered mixture space. The symbol X (and occasionally Y) will denote a partially ordered mixture space, but may sometimes also be used to denote the underlying (unordered) mixture space; no confusion will result.

2. The Utility

A utility on a partially ordered mixture space X is a function from X to the reals for which

(2.1) $u(\gamma x + (1-\gamma)y) = \gamma u(x) + (1-\gamma)u(y)$

(2.2) $x \succ y$ implies $u(x) > u(y)$

(2.3) $x \sim y$ implies $u(x) = u(y)$.

Condition (2.1) is the familiar "expected utility hypothesis," whereas (2.2) and (2.3) state that u represents the preference order.

Our basic result is:

Theorem A. There is at least one utility on X .

3. Two Examples

An example of a mixture space is Euclidean n-space R^n , considered as a vector space over the real numbers. Two of the partial orders most frequently encountered in the literature are the weak and the strong partial orders on R^n, which we denote by $\textcircled{>}$ and $\boxed{>}$ respectively. Using subscripts to denote coordinates, we write $x \textcircled{>} y$ if $x_i > y_i$ for all i ; we write $x \boxed{>} y$ if $x_i \geq y_i$ for all i , but $x \neq y$.[6] Both orders satisfy all assumptions of the previous section; they are also both pure, i.e., indifference holds only in the case of equality. If we normalize the utilities by setting $u(0) = 0$, then the utilities for $\boxed{>}$ are of the form $u(x) = \Sigma_{i=1}^{n} u_i x_i$, where $(u_1, \ldots, u_n) \textcircled{>} 0$; the utilities for $\textcircled{>}$ are of the same form, except that now we need only have $(u_1, \ldots, u_n) \boxed{>} 0$.

4. Maximization Problems, Games, and Shapley's Theorem

The convex hull[7] of a finite set of points in X is called a convex polyhedron.

Theorem B. Let E be a convex polyhedron in X , and let $x \in E$. Then x is maximal in E under the partial order $\succsim$, if and only if there is a utility u on X such that x maximizes u over E .

Theorem B enables us to deal with two-person zero-sum games played over X . These games are similar to ordinary matrix games in all respects, except that the payoffs are in X rather than being real numbers. As usual, the two players each have a finite set of pure strategies, denoted by $(p_1, \ldots, p_k)$, $(q_1, \ldots, q_\ell)$; there is a payoff function which associates with each pair of pure strategies p_i and q_j a member a_{ij} of X .

---

[6] The terminology may sound reversed to the reader, but it has some justification. One partial order is stronger than another if it has more relations; we consider a total order stronger than a partial one.

[7] The convex hull of a set D is the set of all convex combinations of members of D .

If the players use mixed strategies $c = (\gamma_1, \ldots, \gamma_k)$ and $d = (\delta_1, \ldots, \delta_\ell)$ respectively, then the outcome is the point $\Sigma_{i,j}\gamma_i\delta_j a_{ij}$ (abbreviated cAd) in X. The preference order $\succsim$ is associated with the first player; the second player has the opposite order, i.e., he prefers x to y or is indifferent between them, if and only if $y \succsim x$. Corresponding to a saddle-point in ordinary matrix games, we here have equilibrium points; these are pairs of mixed strategies (c^o, d^o) which are "good against each other" in the sense that c^oAd^o is maximal in the set F of all points in X of the form cAd^o, and minimal in the set G of all points in X of the form c^oAd. These equilibrium points have just about all the nice properties of saddle points in ordinary matrix games. For example, the interchangeability property holds: each player has a set of "good" strategies, such that the equilibrium points are precisely the pairs of good strategies. Furthermore, a player not only achieves "best possible" for himself by playing a good strategy, but he also protects himself against loss; if the other player changes his strategy, the result will either be another equilibrium point, or a point that (from our player's point of view) is actually preferred to an equilibrium point. However, there is nothing in this kind of game that corresponds to the unique value of ordinary matrix games.

Do equilibrium points always exist? If so, how can they be calculated? These questions are answered by the following theorem:

Theorem C. (c^o, d^o) *is an equilibrium point in the matrix game* A *if and only if there is a pair* (u, v) *of utilities on* X *such that* (c^o, d^o) *is an equilibrium point (in the sense of Nash [4]) in the bimatrix game* $(u(A), -v(A))$.[8]

---

[8] I.e., the two-person non-zero sum game whose strategy spaces are the same as in the original game, but in which the payoff to (s_i, t_j) is $u(a_{ij})$ to player 1 and $-v(a_{ij})$ to player 2.

This theorem generalizes a result of Shapley [6]. Shapley considered the case in which the underlying mixture space is R^n and the order is either the weak or the strong order. He defined "weak" and "strong" equilibrium points accordingly, and by exhibiting the utilities explicitly, proved what amounts to Theorem C for each of these two special cases separately. Our proof (see section 8) is essentially the same as one of Shapley's two proofs. We quote Theorem C here chiefly as an application of our utilities; it serves to unify Shapley's two results, includes a far larger class of preference orders, and, we believe, exhibits his results in their proper context.

Theorem C can be extended to n-person games. Each player i has a finite set P^i of pure strategies, and an outcome space X^i satisfying the assumptions of section 1. With each n-tuple $(p^1,\ldots,p^n)$ of pure strategies, there is associated an n-tuple $(x^1,\ldots,x^n)$ of payoffs, where $x^i \in X^i$. We now define an n-tuple $(c^1,\ldots,c^n)$ of mixed strategies to be an equilibrium point if each c^i is "good" against the combination of the n-1 others. The result is that $(c^1,\ldots,c^n)$ is an equilibrium point if and only if it is a Nash equilibrium point for some n-tuple of utilities $u^1,\ldots,u^n$ on $X^1,\ldots,X^n$.

5. Discussion

A. The Assumptions All of our assumptions hold in the von Neumann-Morgenstern model, except for the finite dimensionality assumption; with this exception, therefore, our theory is a true generalization of the von Neumann-Morgenstern theory.[9] As we remarked in section 1, finite dimensionality holds in all cases in which there are only a finite number of pure prospects; this includes most cases of practical interest.[10] The assumption of finite dimensionality cannot be dropped. To see this, let X be the set of all infinite sequences of real numbers, and impose the strong order. Then X satisfies all our assumptions except finite dimensionality, but has no utility.[11]

Assumption (1.1) is taken from Hausner's set of axioms [2].

Assumption (1.2) is an extremely weak version of the "archimidean" or "continuity" principle; it is weaker than any variant I have seen. It serves only to exclude the case in which the direction of strict preference between a point z and a closed line segment [xy] goes in one direction for one of the end points y and in precisely the opposite direction for the entire remainder of the segment. Two cases which are <u>not</u> excluded are illustrated by the weak and the strong orders respectively (Figure 1).

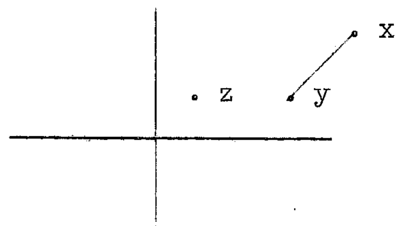

Figure 1

[9]We can strengthen our theory in a trivial fashion so as actually to include the von Neumann-Morgenstern theory, by extending it to all spaces which become finite dimensional when the indifference relation is divided out.

[10]But not all, because it is occasionally useful to build a model in which a certain variable can take a continuum of values, though in actuality there are only finitely many possibilities; price is an example. In most of these cases, however, our theory still applies, because by dividing out the indifference relation we can often get a finite-dimensional space (though usually not one that is the convex hull of a finite number of points).

[11]The dimensionality of X is the cardinality of the continuum. It would be interesting to know whether or not there is a counter-example with denumerable dimensionality.

Here y and z are on the same horizontal line. For both orders all points in the half-open[12] segment [xy) are preferred to z ; the difference between the two orders is expressed in the relation between z and the end-point y . For the weak order y and z are incomparable, and for the strong order y is preferred to z ; in neither case, though, is z actually preferred to y . It can also happen that all points in [xy) are preferred to z , while y and z are indifferent (not pictured).

In practice the effect of (1.2) is to exclude the lexicographic order[13] and other orders inspired by the lexicographic order. We remark that if we drop (1.2) we can still build a utility theory, but the values of the utility functions will be points in a lexicographically ordered euclidean space rather than real numbers; this generalization of the present theory is analogous to Hausner's generalization [2] of the von Neumann-Morgenstern theory. It will still be possible to solve maximization problems and games under exactly the same conditions as before (compare [7]).[14]

---

[12] I.e., the segment [xy] without the point y .

[13] The lexicographic order on R^2 is the pure order for which $x > y$ if and only if: either $x_1 > y_1$, or $x_1 = y_1$ and $x_2 > y_2$. The definition may be generalized to R^n .

[14] I personally believe the archimidean principle to be very compelling, notwithstanding some of the counter-intuitive examples that have been offered in the literature. For example, it is sometimes argued that a trivial prize such as two pins may not be worth any probability of death, no matter how small. But many people drive their cars every day for, say, $50, although they know that this involves a positive probability of death; and by using postulates of utility theory other than the archimidean principle, one can convince oneself that $50 is "comparable" to two pins (by going up a pin at a time, say). The counter-intuitive flavor of the example may be traceable to aspects of the preference axioms other than the archimidean principle; for example, the idealization that asserts the ability to differentiate between probability combinations that are very close to each other may be involved. In spite of all this, there may certainly be situations in which the lexicographic order or something similar constitutes the most convenient model, so it is desirable to have a theory that covers it. (I am indebted to J. Brand for this argument.)

B. Linear Transformations As in the von Neumann-Morgenstern theory, if u is a utility, then so is $\alpha u + \beta$, where $\alpha > 0$ and β is an arbitrary real number. Two utilities connected in this way are called equivalent.

C. Uniqueness The utilities are not unique up to linear transformation; a given preference order may (and usually will) have many different inequivalent utilities.

6. The Structure of Partially Ordered Mixture Spaces

In this section we will give a constructive characterization of spaces X satisfying the assumptions of section 1. Let us first consider the case in which the mixture space involved is R^n. Assumptions (1.1) and (1.2) may then be restated as follows:

(6.1) $x \gtrsim y$ implies $x + z \gtrsim y + z$;

$\alpha > 0$ implies $\alpha x \gtrsim \alpha y$;

(6.2) $\alpha x \succ z$ for all $\alpha > 0$ implies not $z \succ 0$.

A utility in this context is merely a real function on X which represents the order in the sense of (2.2) and (2.3), and which is linear in the ordinary (vector-space) sense; that is, there is a vector $(u_1, \ldots, u_n)$ and a scalar c, such that $u(x) = c + \Sigma_{i=1}^n u_i x_i$. Different c's yield "equivalent" utilities; we will usually normalize[15] by setting $c = 0$. We will denote the vector $(u_1, \ldots, u_n)$ by u, and call it a utility as

[15] This normalization sets $u(0) = 0$; there is also a multiplicative parameter that could be normalized, but there seems to be no unique natural way in which to do this. Note that the "natural" way in which we have fixed the additive parameter depended on the existence of an origin; this is a feature of R^n when considered as a vector-space, but it is not inherent in the mixture-space structure of R^n.

well. Thus $u(x)$ is the same as the inner product ux ; no confusion will result.

For a geometric characterization, we turn to the set S $(= S_X)$ of points in R^n that are $\gtrsim 0$. It is not difficult to see that the order is completely determined by S . Of more significance than S in the analysis, however, is the set T $(= T_X)$ of points in R^n that are $\succ 0$; this may be defined in terms of S by $T = S \setminus (-S)$, where $\setminus$ denotes set-theoretic subtraction. From (1.1) it follows that

(6.3) S is a convex cone,[16]

and from (1.2) that

(6.4) $\overline{T} \cap (-T) = \emptyset$,

where the bar denotes closure; conversely, if these conditions are satisfied, then S defines a partial order. Note that T is also a convex cone, but does not contain the origin. A utility is geometrically characterized by an open support of T , i.e., an open half-space containing T , and whose bounding hyperplane contains the origin; the inner normal to the bounding hyperplane provides the utility.

For the examples of section 3, T is the open positive orthant for the weak order, and for the strong order it is the closed positive orthant minus the origin. Orders on R^n "between" the weak and the strong order are obtained by choosing T to be between these two; for example, for $n = 2$ we could stipulate $T = \{x: x_1 > 0, x_2 \geq 0\}$. Other possibilities for T are open half spaces, open half-spaces of linear subspaces of R^n , circular cones. Excluded are closed half-spaces, or half-spaces that are partly open and partly closed (such as the open

---

[16]A cone is a subset C of R^n such that $x \in C$ and $\alpha > 0$ imply $\alpha x \in C$.

half-plane $x_1 > 0$ to which has been adjoined the positive x_2-axis, which would yield the lexicographic order on R^2). Candidates for S can sometimes be obtained from candidates for T by judiciously adding to T points from $\bar{T}$; details are omitted, but we mention that a closed half-space is a possibility for S, but a partly open half space like the one described above is excluded.

Up to now we have assumed that the underlying mixture space of X is R^n. Hausner [2] has proved that any mixture space may be imbedded in a real vector-space, and from our finite-dimensionality assumption it follows that the vector space will be an R^n. It is not difficult to extend the partial order as well. Thus any partially ordered mixture space X can be described as a convex subset of a partially ordered space Y whose underlying space is R^n, such that the order on X is the restriction to X of the order on Y. Furthermore, the utilities on X will be precisely the restrictions to X of the utilities on Y.

7. Duality

In this section we wish to answer the question: To what extent does the set of utilities on X determine the order on X ?

To this end, we introduce the duality notion. The dual of a cone C in R^n is defined to be the cone C^* consisting of all $u \in R^n$ such that $ux > 0$ for all $x \in C$. For example, the open positive orthant and the closed positive orthant without the origin are mutually dual, as are R^n and $\emptyset$, an open half-space and the ray normal to its bounding hyperplane, and concentric open and closed right circular cones (the latter without the origin) whose half-angles add to 90°. The cone $\{x \in R^2: x_1 > 0, x_2 \geq 0\}$ is self-dual.

In all the above examples $C^{**} = C$. It is of interest to ask under what general conditions this holds. If we calculate C^{**}, we find that it is precisely the intersection of the open supports of C. Thus we have

Theorem D. *A necessary and sufficient condition that* $C^{**} = C$ *is that* C *be the intersection of its open supports.*

The importance of Theorem D lies in the fact that the condition given is of wide applicability. Let us call a cone satisfying the condition *regular*. A regular cone must be convex, and unless it is all of R^n, it may not contain the origin; but aside of these restrictions, almost any cone "liable to come up in practice" is regular. Of course the examples we brought above all involve regular cones. More generally: Any open cone is regular. If C is a convex cone obtained from a closed cone by removing the origin, then C is regular. The set of all x satisfying a given set of homogeneous linear inequalities, *which may contain both weak and strong inequalities*, is regular if it contains at least one strong inequality (which may, for example, serve only to remove the origin). If C is an open circular cone, then any cone between C and $\bar{C}$ that does not contain 0 is regular. On the other hand, if we add the positive half of one of the axes to the open positive octant in R^3, the result is a cone which is not regular, though it is convex and does not contain the origin.

If X is a partially ordered copy of R^n, then the set of all (normalized) utilities on X is precisely T_X^* (where T_X is the set of all points preferred to 0). Hence if we know that T_X is regular, we can recover the order from the set of all utilities. Thus the set of all utilities on a given X "almost" determines the order, and determines it completely if the set of orders under consideration is suitably restricted.

Our definition of duality is somewhat different from the ordinary definition, in which C^* is defined to be the set of all u such that $ux \leq 0$ for all $x \in C$. Under that definition the necessary and sufficient condition that $C^{**} = C$ is that C be the intersection of its closed supports, or equivalently that it be convex and closed.

8. Proofs

Proof of Theorem A. We assume that the underlying mixture space of X is R^n; this involves no loss of generality because any finite-dimensional mixture space can be imbedded in such a mixture space. The proof is by induction on n. If $n = 1$ the order must either be total or all elements are incomparable; in either case the theorem is trivial. Suppose the theorem has been proved for all dimensions up to but not including n. If there is an element of X other than 0 that is indifferent to 0, then we may "divide out" the indifference relation, i.e., consider equivalence classes under indifference; this yields a space of lower dimension, to which the induction hypothesis applies. We may therefore assume without loss of generality that the order on X is pure, so that $S = T \cup \{0\}$ (cf. section 6). Suppose first that the closure of $T \cup (-T)$ is not all of X, and let w be a point not in that closure. Let Y be a subspace of X such that every $x \in X$ is uniquely of the form $\beta w + y$, where β is real and $y \in Y$; for example, take Y to be the orthogonal complement of the line L_w spanned by w. Define an order on Y by $y \gtrsim 0$ if and only if there is a β such that $y + \beta w \gtrsim 0$ in X; geometrically, S_Y is the projection of S_X on Y in the direction of L_w. Hence S_Y is a convex cone, and to prove that the order on Y satisfies our assumptions, it remains only to establish that $\overline{T}_Y$ and $-T_Y$ do not meet. Indeed, suppose y is in their intersection. Noting that T_Y is the projection of T_X on Y in the direction of L_w, we deduce that there is a β and sequences $\{\beta_1, \beta_2, \ldots\}$ and $\{y_1, y_2, \ldots\}$ such that $y_i \to y$, $y_i + \beta_i w \succ 0$ in X, and $0 \succ y + \beta w$ in X. Let β_∞ be a limit point (possibly infinite) of the β_i; without loss of generality we may assume that it is actually the limit. If $\beta_\infty = \beta$, then $y + \beta w \in \overline{T}_X \cap (-T_X)$, contrary to (6.4). If $\beta_\infty \neq \beta$ but is finite, then

$$w = \frac{(y_i + \beta_i w) - (y + \beta w)}{\beta_\infty - \beta} + \frac{y - y_i}{\beta_\infty - \beta} + \frac{\beta_\infty - \beta_i}{\beta_\infty - \beta} w ,$$

and hence w is the sum of a term that is either $\succ 0$ or $\precsim 0$ (according as $\beta_\infty > \beta$ or $\beta_\infty < \beta$) and terms that tend to 0, contrary to the assumption that w is not in the closure of $T_X \cup (-T_X)$. If $\beta_\infty = \pm \infty$, then

$$w = \frac{y_i + \beta_i w}{\beta_i} + \frac{y - y_i}{\beta_i} - \frac{y}{\beta_i} ,$$

and again w is the sum of a term that is $\succ 0$ or $\precsim 0$ and terms that tend to 0, yielding a contradiction. This proves that the order on Y satisfies our assumptions, and since Y is of lower dimension than X, we can apply the induction hypothesis to construct a utility on Y. This utility can now be extended to X by setting $u(y + \beta w) = u(y)$.

Finally, suppose that the closure of $T \cup (-T)$ exhausts X. If we could show that T is open, then since $0 \notin T$, it would follow that there must be a hyperplane through 0 that does not intersect T (cf. for example [1,p.19, Theorem 7]); the normal to this hyperplane in the direction of the half-space occupied by T would then provide a utility. It remains therefore to show that T is open. Contrariwise, suppose $x \in T$ is on the boundary of T. Let H be a support hyperplane for T through x. Any neighborhood of x will contain points on both sides of H, and therefore in particular it will contain a point that is not in $\overline{T}$; this point must therefore be in the closure of $-T$. Therefore x itself will be in the closure of the closure of $-T$, which is the same as $-\overline{T}$. Therefore $-x \in \overline{T}$. But since $x \in T$, it follows that $-x \in -T$; so $-x \in \overline{T} \cap (-T)$, contradicting $\overline{T} \cap (-T) = \phi$.

Proof of Theorem B. The "if" statement follows from the definition of utility (2.2). To prove the "only if" half, we assume again that the underlying space is R^n and that the order is pure. The remainder of the proof follows Shapley's proof precisely; it is included only for the sake of completeness. Let D be the set of points dominated by members of E , i.e.

$$D = \{y: \exists z \in E \text{ such that } z \gtrsim y\} .$$

D is a polyhedral set; let D^r be the unique r-dimensional face of D whose (relative) interior contains x . Let H be a supporting hyperplane for D that meets D precisely in D^r . Then if u is the normal to H , we have $u(x-y) \geq 0$ for all $y \in D$, with equality only if $y \in D^r$. We claim that $u(x-y) > 0$ whenever $x \succ y$. If not, there would be a $y \in D^r$ such that $x \succ y$. But, since x is in the relative interior of D^r , there would also be a $z \in D^r$ (of the form $z = x + \epsilon(x-y)$, $\epsilon > 0$) such that $z \succ x$. This contradicts the maximality of x . Thus our claim is substantiated, and it follows that u is a utility.

Theorem C follows at once from Theorem B.

REFERENCES

[1] H. G. Eggleston, <u>Convexity</u>, Cambridge University Press, 1958.

[2] M. Hausner, "Multidimensional utilities," in [8], pp. 167-180.

[3] R. D. Luce and H. Raiffa, <u>Games and decisions</u>, John Wiley, 1957.

[4] J. F. Nash, "Non-cooperative games," <u>Ann. of Math.</u> 54 (1951), pp. 286-295.

[5] J. von Neumann and O. Morgenstern, <u>Theory of games and economic behavior</u>, Princeton University Press, 1944, Third edition, 1953.

[6] L. S. Shapley, "Equilibrium points in games with vector payoffs," <u>Nav. Res. Log. Quart.</u> 6 (1959), pp. 57-61.

[7] R. M. Thrall, "Applications of multidimensional utility theory," in [8], pp. 181-186.

[8] R. M. Thrall, C. H. Coombs, and R. L. Davis, Editors, <u>Decision processes</u>, John Wiley, 1954.