

22

SUBJECTIVE PROGRAMMING

Robert J. Aumann

Econometric Research Program
Research Memorandum No. 22
10 February 1961

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

RETAIN THIS COPY
FOR YOUR FILES
USE JOB NO. BELOW
<input type="checkbox"/> TO CHECK YOUR BILLS
<input type="checkbox"/> TO REORDER
<input type="checkbox"/> FOR ALL REFERENCES
JOB NO. 3956-61
No. Copies 23 X 80
APR 14 1961
PRINCETON UNIVERSITY DUPLICATING BUREAU

SUMMARY

Under consideration are programming problems in which the objective function is given as a partial preference order rather than numerically. For example, many of the problems that are generally associated with the term "military worth" fall into this category. Conditions are given for such a problem to be transformable into a problem that can be handled numerically.

Introduction

Programming problems are defined by a constraint set and an objective function. The role of the constraint set is to define which courses of action—the technical word is activity vectors—are feasible, and the role of the objective function is to define the preference relation on the activity vectors. The aim is to find a feasible activity vector which is preferred or indifferent to all other feasible activity vectors. The objective function is assumed to be given as a numerical function defined on the space of all activity vectors, or at least on a neighborhood of the constraint set; preferred activity vectors correspond to higher values of the function. If the objective function is linear and the constraint set is a polyhedron, then the computational methods of linear programming are applicable. Computational techniques are also available for some kinds of problems in which the objective function is not linear, but in any case, all such techniques do use a numerically defined objective function.

There is, however, a large class of practical problems in which there is no à priori numerical objective function readily available. For example, many of the programming problems that arise in the military establishment are concerned with the maximization not of some definite and measurable quantity such as dollar profit, but rather of the ill-defined, vague "military worth." Mathematical techniques for dealing with such problems are useless unless military worth is adequately defined and some method is given for measuring it. Nor is this problem area limited to the military establishment; any organization which does not operate on a profit motive (including any government department) is in a similar

situation. Even ordinary profit-making business corporations may be faced with such a problem, when it is a question of deciding on activities whose effect on the profit-making mechanism, though vital, is too remote and complex to be effectively calculable. For example, many employee-assignment problems are of this kind; so are most problems involving activities whose main effect is on good-will.

The object in all these problems is still to find a "best possible" activity vector in the constraint set, but the term "best possible" is no longer defined by a numerical objective function. Instead, we must assume that there is some kind of preference structure defined on the set of all activity vectors, and solve our problem in terms of this structure. The preferences are presumably those of the individual responsible for making decisions in the problem involved. Hence the name "subjective programming"; the decision is to be made on the basis of subjective preferences rather than objectively measurable quantities. Perhaps "non-numerical programming" would have been a more accurate description of the mathematical context, but the name we have chosen is, we think, more suggestive of the applications.

What can we say about the preference structure? First, it will in general be a partial order, not a total order; it would be unreasonable and unrealistic to expect the decision-maker to have a well-defined preference (or indifference¹) between any pair of activity vectors.² Therefore

¹We stress the difference between indifference and incomparability: Indifference between two activity vectors involves a positive decision that it doesn't make any difference if 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 are not.

²Often the available preference information is restricted to a small subset even of the feasible activity vectors. For example, in the large-scale allocation problem discussed in [7], available preference information was restricted to a number of "token allocation plans" each of which involved only a single assignment.

we will not always be able to find a maximum feasible activity vector, i.e., one preferred or indifferent to all others; we will in general have to be content with finding a maximal activity vector, i.e., one to which no other feasible activity vector is strictly preferred. The second condition on our preference structure is transitivity.³ We have no wish to enter into a justification of this slightly controversial assumption, but state merely that without it the search for a maximal element is fruitless and the programming problem becomes meaningless. Possibly if the decision making remains the province of a single individual, the transitivity assumption is less objectionable than it might otherwise be.

Let us for the moment restrict our attention to the case of discrete (or integer) programming problems—problems in which the coordinates of the activity vectors take integer values only. Many real-life subjective programming problems have this form, and it is the more interesting case from the mathematical viewpoint. What we are seeking is a "bridge"—a method which would enable us to use the existing computational techniques, which were invented for numerical objective functions, on our non-numerical problems. The most obvious such bridge would be a numerical function u that "represents" the preference order, in the sense that if x and y are activity vectors such that $x \succ y$ (x preferred to y) then $u(x) > u(y)$.⁴ Maximization of such a function over the constraint set would lead to an element maximal in the preference order. Furthermore, such a function always exists; the given partial order can be extended to a total order, and any "monotone" function on the total order will do the

³This was previously implied by our use of the word "order."

⁴And if x and y are indifferent then $u(x) = u(y)$.

trick. The trouble is that this is too broad to be useful. The numerical function will have no regularity properties, will in general not be defined by a "formula," and therefore will be hard to compute with; it is no better than working with the original order. What we should ask ourselves is whether we can find functions of a specified kind to represent our preference order. In this paper we will investigate the question: Under what conditions on the preference order can a linear function be found to represent it? In many circumstances the existence of such a function would be of great importance—even if only from the practical, computational viewpoint. The (unsurprising but useful) answer we have found is as follows: Roughly speaking, if and only if $x \succ y$ implies $x + z \succ y + z$ for all activity vectors $x, y,$ and z (where $+$ stands for ordinary vector addition).

In the next section we give an exact description of our assumptions and an exact statement of our results. Section 2 is devoted to discussion and elucidation of the section preceding it. In Section 3 we describe a complex of problems, largely unexplored, which are suggested by this study; we consider that these problems are both fascinating from the mathematical viewpoint and crucial for the applications. Section 4 is devoted to generalizations and extensions of the basic results in various directions, including ordinary, non-discrete programming problems. Proofs are given in Section 5.

1. The Main Theorem

We denote by X the subset of euclidean n -space E^n consisting of all points with non-negative integer coordinates; X is the space of all activity vectors. We shall assume that on X there is imposed a

preference relation,⁵ denoted by \succsim and called "preference-or-indifference." If $x \succsim y$ and $y \succsim x$, then we shall write $x \sim y$, and say that x is indifferent to y . If $x \succsim y$ but $x \not\sim y$ (x not indifferent to y), then we shall write $x \succ y$ and say that x is preferred to y . Expressions of the form $x \succ y$ and $x \sim y$ will be called preference statements. The following assumptions are made about the preference relation ($+$ refers to vector addition, and all statements not otherwise quantified refer to all x, y , and z).

Transitivity If $x \succsim y$ and $y \succsim z$, then $x \succsim z$.

Reflexivity $x \sim x$

Additivity i) $x \succ y$ implies $x + z \succ y + z$

ii) $x \sim y$ implies $x + z \sim y + z$.

Finite Generation There is a finite set S of preference statements, such that every preference statement can be deduced from one of the statements in S by means of a finite chain of applications of the transitivity and additivity assumptions.⁶

Main Theorem There is a real-valued linear function u on X such that $x \succ y$ implies $u(x) > u(y)$ and $x \sim y$ implies $u(x) = u(y)$.

A linear function on X satisfying the conditions of the main

⁵Nothing is assumed about this relation that is not specifically stated in the sequel.

⁶This assumption may be paraphrased as follows: we shall say that a preference relation \succsim includes another one \succsim^* if $x \succ^* y$ implies $x \succ y$, and $x \sim^* y$ implies $x \sim y$. Further, for a given transitive, reflexive and additive preference relation \succsim , we shall say that a set S of preference statements generates the relation \succsim , if this relation is included in every transitive, reflexive and additive preference relation for which the statements in S are true. The assumption then says that the given relation is generated by a finite set of preference statements.

theorem is called a utility function. The main theorem has a partial converse (to which we were referring in the "only if" half of the rough statement in the introduction), which we will state in the following section.

2. Discussion

A) The space X , on which both the given preference order and the resulting utility are defined, consists of all the activity vectors,⁷ not just the feasible ones. This corresponds to the real-life situation; the preference relation and the constraint set are determined by completely different considerations, and there is no reason to restrict the former to the latter. The housewife, for example, may be able to express meaningful preferences as between various bundles of groceries, even though not all of them are feasible from the point of view of the week's budget. Furthermore, prices of groceries or the family's earnings might change without affecting preferences, so that this week's infeasible vector may be feasible next week. Thus the preference relation will ordinarily be defined on many pairs of vectors that are not in the constraint set, and this is perfectly reasonable.

B) We emphasize again that the preference order is partial, that is, it is in general not defined on all pairs of vectors. We also repeat that the lack of a preference between two vectors does not imply indifference between them.

C) Transitivity: The reader is asked to refer to the introduction (p.4) for a brief discussion of this assumption.

D) Additivity: This is the heart of the matter; we put off a full discussion to the end of this section. Here we note only that

⁷Vectors with non-negative integer coordinates.

additivity for the preference-or-indifference relation (i.e. $x \succsim y$ implies $x + z \succsim y + z$) would not be sufficient for our purposes; counter-examples to the main theorem for orders that satisfy only this weaker assumption are easily constructed.

E) Finite generation: This constitutes no real restriction, at least as far as the applications are concerned. Though the preference order must theoretically be defined on all of X (in order to enable us to prove our theorem), we are practically speaking only interested in its restriction to the constraint set, or to a "neighborhood" of the constraint set sufficiently large to contain all constraint sets that are liable to occur in practice. A utility for the preference order generated by this restricted order will accomplish everything that a utility for the original order would have accomplished; preference statements that are not contained in the restricted order have no practical relevance.⁸

F) The utility "represents" the preference order, but in a weaker sense than is usually understood under the word "utility." Though a preference $x \succ y$ implies the corresponding utility-inequality

⁸The following superficially appealing justification for the finite generation assumption is not valid: Obviously the decision maker can make only a finite number of "basic" preference statements; though an infinite number of preference statements can be deduced from the basic ones by the application of rules such as transitivity and additivity, which the decision maker accepts as reflecting his views, the whole preference order will always be generated by the finite set of basic decisions. The argument is invalid for the following reason: If the decision maker can adopt general rules—such as additivity—as reflecting his views, then he can also adopt other rules of the same nature; for instance he could adopt a lexicographic-type rule that a single unit of a certain commodity is preferred to any amount of another commodity. Such a rule leads to a preference order that is not finitely generated, and indeed the main theorem becomes false for such preference orders. The point is that it will still be possible to define a useful utility, by using the preference order generated by the restriction of the given one to a finite neighborhood of the constraint set; though this utility will not represent the original preference order outside this neighborhood, it will still yield the correct answer to all the programming problems that will arise in practice.

$u(x) > u(y)$, the converse is false; if $u(x) > u(y)$, x and y may be incomparable. Indeed, a moment's reflection will convince the reader that a utility that is to represent a partial order cannot be a two-way utility of the usual kind. This apparent defect does not, however, affect the usefulness of the utility in solving programming problems; maximization of the utility still leads to a maximal element of the constraint set. Usually there will be a number of different utilities; maximization of any of them will lead to a maximal member of the constraint set, and conversely, any maximal member of the constraint set can be obtained by maximizing an appropriate utility. Since the utility is to be used only for the solution of programming problems there is no need for it to be unique.

G) The converse of the main theorem: We are particularly concerned with the relation between the additivity assumption and the existence of a utility. Assume that a given preference relation has a utility; must the additivity assumption be satisfied? The answer is in general no⁹; but the additivity assumption must be satisfied whenever $x + z$ and $y + z$ are comparable. More precisely, let us define the preference order to be weakly additive if

- i) $x \succ y$ implies either that $x + z \succ y + z$, or that $x + z$ and $y + z$ are incomparable; and
- ii) $x \sim y$ implies either that $x + z \sim y + z$ or that $x + z$ and $y + z$ are incomparable.

Then every preference relation for which there is a utility must be

⁹Let X be the non-negative integers, and suppose the preference relation to contain (in addition to the statements $x \succ x$) only the single statement $1 \succ 0$ (from which it follows that $1 \succ 0$). This has a utility given by $u(x) = x$, but obviously the preference relation is not additive.

weakly additive. It must also be weakly transitive and reflexive, if these concepts are defined in the corresponding manner.

On the other hand, nothing of this kind can be said about finite generation. An alternative statement of the converse is that a preference relation for which there is a utility may be extended to a transitive and additive utility. The proof of the converse, unlike that of the main theorem, is trivial.

H) Computing the utility. Write

$u(1,0,\dots,0) = u_1, \quad u(0,1,0,\dots,0) = u_2, \dots, u(0,\dots,0,1) = u_n,$
and $x = (x_1,\dots,x_n), \quad y = (y_1,\dots,y_n)$. Obviously it is sufficient to determine the u_i . A preference statement of the form $x \sim y$ supplies us with the equation $\sum_i (x_i - y_i)u_i = 0$; if $x \succ y$ then $\sum_i (x_i - y_i)u_i > 0$. Let S be a set of preference statements that generates the preference relation—such a set exists because of the finite generation assumption. Each member of S is either an indifference or a strict preference; form the system S^* of equations and strict inequalities corresponding to the members of S . Any feasible solution (u_1,\dots,u_n) of this system gives us a utility, defined by

$$(1) \quad u(x) = \sum_i x_i u_i ;$$

conversely any utility u is of the form given in (1), where the u_i are a feasible solution of S^* .

From a practical viewpoint, the question of computing the utility is often more complicated. In this paper we wish to avoid a detailed discussion of the practical difficulties that arise¹⁰; we will merely

¹⁰The reader is referred to [7-10], where a fairly complete discussion is given of a complex real-life subjective programming problem, more or less in the spirit of this paper.

mention one of them, which has some theoretical interest and significance as well. The set S of basic decisions may be extremely large; it is often impractical explicitly to ascertain all of these decisions. The result of having a set of available decisions that is smaller than the true set is that the polyhedron of feasible solutions to the resulting system of equations and inequalities is larger than the true polyhedron of utilities. So we are not sure that a member of the larger polyhedron is actually a utility; in some sense, though, we may call it an approximate utility. The first question is, in what sense is it an approximation, and can we give any measure of how good the approximation is? We are not really interested in the utility as such; we are interested in its use as a tool for solving programming problems. Use of the approximate utility in a programming problem may be expected to lead to an answer that is not in general optimal; we may hope that it is "close" to optimal. The basic question here is how to define "closeness"; when we have done this we may be able to use the resulting measure of closeness on the space X of activity vectors to define an appropriate "closeness" measure on the space of utilities. One feels that the measure of closeness on X should be based on the polyhedron of true utilities, but it is not clear exactly how.

Once these basic questions have been answered, it is possible to ask whether there are any good techniques of approximation which have general validity. More precisely, suppose we have some control over the questions on which the decision maker will be asked to decide, the results of which will be used to find approximate utilities; how should we exercise this control in an efficient manner in order to make the

approximation good?

I) We now return to a conceptual discussion of the additivity assumption. This is a strong assumption, but as we saw under (G), it is implicit in the use of a linear objective function. Roughly, it may be said to hold whenever the effectiveness of an individual activity in an activity vector does not depend on the other activities being performed at the same time. For example, it will not hold in a situation which is governed by a law of diminishing returns, or in personnel assignment problems in which compatibility considerations play an important part.

Here are some examples of situations in which the additivity assumption does hold:

i) Problems in which the various activity units operate entirely independently of one another. For instance, consider an employment agency which has a number of candidates and a number of firms on its roster, where each firm has only one vacancy. To keep things subjective, assume that the agency gets a fixed fee for each assignment, and is therefore interested only in maximizing good will.

ii) Programming problems in which the interactions between the activity units are relatively minor and difficult to analyze. Allocation of clerical and other semi-skilled workers in a single organization might be an example. Another example of the same type is that discussed in [7-10].

iii) Certain situations governed by a law of diminishing returns, but where we are interested only in adding to current activity in relatively small amounts. Though the "utilities" here will not in general be linear (which we demand in our formal definition), they may be "marginally

linear." We give a more precise discussion at the end of Section 3.

J) There is another aspect of the additivity assumption that is open to some question. It is asserted that $x \succ y$ implies $x + z \succ y + z$ for all z , no matter how large. Now in reality, there may often be some practical limit on the set of activity vectors between which the decision maker is willing to express preferences. (The housewife of subsection (A) above, though willing to express preferences as between certain bundles of groceries, may call a halt when she is confronted with a pair of bundles each of which contains millions of cans of vegetables.) Thus the additivity assumption as it stands may often have to be regarded as an idealization of the true situation.

3. The Restricted Additivity Assumption

We wish to pursue further the question raised in the last subsection of the previous section. Let us say that the decision-maker wishes to restrict his preference statements to a certain set A of activity vectors. Within A the additivity assumption is assumed to hold; that is, it holds whenever $x, y, x + z$ and $y + z$ are in A . As before, the order within A is only partial, and the set A need by no means be restricted to the constraint set; on the contrary, in general it will be large enough to contain all the constraint sets that the decision maker thinks may arise in a given context, and may be a good deal larger. We wish to know whether the main theorem still holds in this situation, i.e., whether it is possible to define a utility on A .

Unfortunately, the answer is no. Roughly, we may say that it

is possible for the comparisons within A to carry within them the "seeds" of contradiction, but that this contradiction need not take place until we have gotten beyond A (and thus beyond the responsibility of the decision maker).

There are several directions in which we might proceed if we wish to save the situation. We might try in various reasonable ways to restrict the kind of set that A may be, for example by demanding that it be "convex" (i.e., the intersection of a convex set with the lattice points¹¹), or that it contain with a given point x all non-negative lattice points y which have coordinates no greater than those of x , or both. Neither of these conditions will do the trick, as is shown by the following example: Let A be the set of all lattice points x in euclidean 4-space E^4 satisfying $x_i \geq 0$ for $i = 1, \dots, 4$ and $\sum_{i=1}^4 x_i \leq 2$. Write $a = (1, 0, 0, 0)$, \dots , $d = (0, 0, 0, 1)$; the order on A is defined by $0 < a < b < c < d < 2a < a + b < 2b < a + c < a + d < b + c < 2c < b + d < c + d < 2d$; note that it is total. Let u be a utility; we obtain $u(b) + u(d) > 2u(c)$ and $u(b) + u(c) > u(a) + u(d)$. Hence $2u(b) + u(c) + u(d) > 2u(c) + u(a) + u(d)$; therefore $2u(b) > u(a) + u(c)$, contradicting $2b < a + c$. Thus even a total order on a perfectly "well-behaved" set does not always possess a utility.

We might hope to obtain positive results by restricting A to be of a certain geometric shape. Thus we might ask our question when A is a "box"—the set of all points of X which are coordinate-wise \leq a given point x . An equivalent question would be the corresponding one when A consists of the vertices of the unit cube. Unfortunately we do not even know the answer to this when the order is assumed to be

¹¹Points with integer coordinates.

total. It has been verified to be true for unit cubes of dimension up to 4 (when the order is total).

The problem of defining a utility on A is equivalent to that of extending the given order to all of the lattice points in E^n ; this follows from our main theorem. Thus the above problem may be formulated as follows: can a (partial or total) additive order on the vertices of the unit cube in E^n be extended to an additive order on all the lattice points of E^n ?

Possibly positive results can be obtained if some other kind of geometric shape (rather than a box) is assumed for A .

A completely different approach is as follows: Let A be the set on which the preference order is defined, and B the set on which a utility is needed (say the union of the constraint sets liable to occur in practice). In general A contains B . It seems intuitively clear, and indeed is not difficult to prove, that if A is sufficiently large with respect to B (but still finite), then a utility will be definable on B (any contradictions caused by the preference order on B must be realized "not too far" from B). Now the question arises, how large must A be, as a function of B , for a utility to be definable on B ? Alternatively, we might ask the following question: Let B be a fixed finite subset of X , and let A be a superset of B . How large must A be chosen so that every preference order on B that is extendable to A is already extendable to all of X ?

We are now in a position to explain what we meant by "marginal linearity" (cf. (iii) of subsection I of the previous section). We are concerned with a situation in which a large number of activities have

already been performed, and we are interested in performing a small number of additional activities. The additivity assumption is assumed to hold "in the small," i.e., as long as z stays in the "marginal" range (i.e., does not cause $x + z$ and $y + z$ to become too large compared with the activity vector already performed). If it is possible to extend the preference order in an additive manner beyond the marginal range, then it will be possible to define a utility within the marginal range, even though the extended preference order may be totally unrealistic outside the marginal range.

4. Extensions and Generalizations

These are possible in a great variety of directions, some of which we indicate below:

A) Non-discrete programming: Instead of assuming that X consists of the non-negative lattice points in E^n , we may assume that it contains the entire non-negative orthant. In this case it is natural to add the following linearity condition to the additivity assumption: $x \succcurlyeq y$ implies $\alpha x \succcurlyeq \alpha y$ for all positive reals α . The main theorem remains true.

B) Dropping the Finite Generation Assumption: This leads to an analogue of the main theorem in which ordinary utility must be replaced by multidimensional utility (cf. [1]). A multidimensional utility u of dimension m is the same as an ordinary utility, except that the range of u is lexicographically ordered euclidean m -space. The dimension m of the range of u can always be chosen so as to be no greater than the dimension n of X . The solution of programming problems using multidimensional utilities is perfectly straightforward; it consists of the

successive solution of m ordinary programming problems, each one using as objective function a different component of the multi-dimensional utility, and adding an additional linear constraint to the previous constraint set. Practically, there would be very little point in going through with such a procedure, for the reasons we explained in justifying the finite generation assumption (subsection E of Section 2). Anyway, all components of the multidimensional utility other than the first are completely insignificant when compared with the first component; the slightest error in the first component—even a computational round-off error—completely swamps all multidimensional considerations.

C) Generalizing the space X . The main theorem remains true when X is an arbitrary commutative semigroup. There is also an analogue of the theorems stated in (A) and (B); in (B), however, the utility may have infinitely many dimensions, and there may be no "first" component.

5. Proofs

There is a vast mathematical literature that is closely related to the ideas of this paper (and of utility theory in general) in that it deals with ordered structures (topological spaces, groups, semigroups, vector spaces, etc.). Much of this literature is concerned with classifying the structures, not with representing them by utility functions (i.e., order-preserving homomorphisms to the reals); but the two problems are closely related. It is possible to prove our theorems by making heavy use of some of this literature,¹² but even then the proof is

¹²Cf. [2-5].

comparatively long (though not particularly difficult). It is our belief that there should be available a comparatively short and elementary proof at least of our main theorem; this we proceed to give. The chief tool is the theorem of the supporting hyperplane. Our procedure will be first to prove the theorem on the existence of multidimensional utilities stated in (B) of the previous section, then apply the finite generation condition to obtain the desired result.

Let us define a weak m -dimensional utility to be a linear function u from X to E^m for which $x \succcurlyeq y$ implies $u(x) \geq u(y)$ in the lexicographic order on E^m . If u and v are two such utilities—possibly with different m —then u is stronger than v if u distinguishes between all elements that v does, but not vice-versa; more precisely, $v(x) > v(y)$ implies $u(x) > u(y)$, but there are x and y such that $x \succcurlyeq y$, $v(x) = v(y)$ and $u(x) > u(y)$. Our basic lemma says that for any weak multidimensional utility v which is not already a true multidimensional utility (in the sense that $v(x) > v(y)$ for all $x \succcurlyeq y$), there is a multidimensional utility that is stronger. Starting out with the trivial 0-dimensional utility $v = 0$, we thus build up a ladder of stronger and stronger weak multidimensional utilities; we show that the process must finally terminate in a true multidimensional utility, and indeed in not more than n steps.

Lemma 1 If $x \succcurlyeq y$ and $z \succcurlyeq w$, then $x + z \succcurlyeq y + w$.

Lemma 2 Let B be a subset of E^n containing rational points only, i.e., points all of whose coordinates are rational. Then any rational point in E^n that is expressible as a linear combination, with non-negative coefficients of points of B , is also expressible as a linear combination, with non-negative rational coefficients, of points of B .

The proofs of lemmas 1 and 2 are left to the reader.

Lemma 3 Let v be a function on X . Assume

(4) there are $x, y \in X$ for which $v(x) = v(y)$ and $x \succ y$.

Then there is a linear function w from X to the reals such that

(5) if $x, y \in X$, $v(x) = v(y)$, and $x \succeq y$, then $w(x) \geq w(y)$; and

(6) there are $x, y \in X$ for which $v(x) = v(y)$, $x \succ y$, and $w(x) > w(y)$.

Remark This is the basic lemma referred to above. If v is a weak m -dimensional utility, then (v, w) is a weak $(m+1)$ -dimensional utility stronger than v . Assumption (4) says that v is not already a true multidimensional utility.

Proof Let W be the set of all members of E^n of the form $x - y$, where $x, y \in X$, $v(x) = v(y)$, and $x \succeq y$. Let W_1 be the convex cone generated by W , and let W_2 be the linear subspace of E^n spanned by W . Let m be the dimension of W_2 . From (4) it follows that W contains points other than the origin 0 . Hence $m \geq 1$. Let us assume for the moment that W_1 does not fill W_2 . Then 0 must be in the frontier of W_1 when considered as a subset of W_2 , and therefore by the theorem of the supporting hyperplane (see, for example [6], Theorem 8, p.20), W_1 has an $(m-1)$ -dimensional supporting hyperplane H in W_2 that contains 0 . Let H' be the $(n-1)$ -dimensional hyperplane in E^n that is spanned by H and the orthogonal complement of W_2 in E^n . We define w to be the distance from H' (appropriately directed); the verification of (5) and (6) is straightforward.

It remains to prove that W_1 does not fill W_2 . Suppose it does. Let $z = x - y$ be a non-zero member of W , where $x, y \in X$ and $x \succ y$. Then $-z \in W_2$, and hence $-z \in W_1$. Applying lemma 2, we

deduce that there are members z_1, \dots, z_k of W , and non-negative rationals $\alpha_1, \dots, \alpha_k$, such that $-z = \sum_{i=1}^k \alpha_i z_i$. Multiplying by a positive common denominator p of the α_i , we obtain non-negative integers q_i such that $pz + \sum_{i=1}^k q_i z_i = 0$. But $z_i = x_i - y_i$, where $x_i, y_i \in X$ and $x_i \succsim y_i$; hence it follows that

$$(7) \quad px + \sum_{i=1}^k q_i x_i = py + \sum_{i=1}^k q_i y_i .$$

On the other hand, if we recall that $x \succ y$ and $x_i \succsim y_i$ and apply lemma 1 a number of times, we obtain

$$px + \sum_{i=1}^k q_i x_i \succ py + \sum_{i=1}^k q_i y_i ,$$

which contradicts (7). This completes the proof of lemma 3.

The true dimension of a weak m -dimensional utility u is the dimension of the linear subspace of E^m spanned by $u(X)$.

Lemma 8. If u is stronger than v , then the true dimension of u is larger than the true dimension of v .

Proof Extend u and v to all of E^n ; they are then linear functions and can be represented by matrices, which we also denote by u and v respectively. Since u is stronger than v , $u(z) = 0$ implies $v(z) = 0$, but there are z such that $v(z) = 0$ and $u(z) \neq 0$ (set $z = x - y$ in the definition of "stronger"). It follows that the nullity of v is greater than the nullity of u , and hence the rank of u exceeds that of v . Since the rank is precisely the true dimension, the proof of the lemma is complete.

Lemma 9 The true dimension of a weak multidimensional utility on X is not larger than the dimension n of X .

Proof A linear mapping cannot raise the dimension of a vector space.

Theorem 10 A preference order on X which is transitive, reflexive

and additive (but not necessarily finitely generated) has an m -dimensional utility, where $m \leq n$.

Proof Starting out with the weak 0-dimensional utility $v = 0$, successively apply lemma 3 to obtain stronger and stronger weak multidimensional utilities; by lemma 8 each of these must have true dimension greater than the preceding one. Hence by lemma 9 the process terminates in at most n steps; the end result must be a true multidimensional utility, and since each step raises the nominal dimension by only 1, we cannot have nominal dimension greater than n at the end.

Proof of the Main Theorem Let $v = (v_1, \dots, v_m)$ be a multidimensional utility, and let ϵ be a positive real number; write $u_\epsilon(x) = \sum_{i=1}^m v_i(x) \epsilon^i$. If ϵ is small then u_ϵ "behaves like" v , because each term in the expression for u_ϵ is "much more significant" than the succeeding term. More precisely, if $x \sim y$ then $u_\epsilon(x) = u_\epsilon(y)$ for all ϵ ; if $x \succ y$ then $u_\epsilon(x) > u_\epsilon(y)$ for sufficiently small ϵ , i.e., when $0 < \epsilon < \epsilon_0 = \epsilon_0(x, y)$. Now choose a positive ϵ which is less than $\epsilon_0(x, y)$ for all pairs (x, y) such that the preference statement $x \succ y$ is in S (the set of preference statements that generates the preference relation). Then u_ϵ is a utility.

REFERENCES

- [1] M. Hausner, "Multidimensional utilities," Chapter XII of Decision Processes (edited by Thrall, Coombs, and Davis), John Wiley, 1954, pp. 167-180.
- [2] _____ and J. G. Wendel, "Ordered vector spaces," Proc. Amer. Math. Soc. 3 (1952), pp. 977-982.
- [3] H. Hahn, "Über die nichtarchimedischen Grössensysteme," Wiener Berichte 116 (1907), pp. 601-655.
- [4] C. J. Everett, "Note on a result of L. Fuchs on ordered groups," Amer. Jour. Math. 72 (1950), p. 216.
- [5] O. Nakada, "Partially ordered abelian semigroups I. On the extension of the strong partial order defined on abelian semigroups." Jour. Fac. Sci. Hokkaido Univ. 11 (1951), pp. 181-189.
- [6] H. G. Eggleston, Convexity, Cambridge Univ. Press, 1958.
- [7] R. J. Aumann and J. B. Kruskal, "Assigning quantitative values to qualitative factors in the Naval electronics problem," Nav. Res. Log. Quart. 6 (1959), pp. 1-16.
- [8] _____, "The coefficients in an allocation problem," ibid. 5 (1958), pp. 111-123.
- [9] J. W. Smith, "A plan to allocate and procure electronic sets by the use of linear programming techniques and analytical methods of assigning values to qualitative factors," ibid. 3 (1956), pp. 151-162.
- [10] G. Suzuki, "Procurement and allocation of Naval electronic equipments," ibid. 4 (1957), pp. 1-7.