13 Subjective Programming

Programming problems are defined by a constraint set and an objective function. The role of the constraint set is to limit the courses of action—the technical word is *activities*—to those that are feasible, and the role of the objective function is to give the preference relation on the activities.[1] The aim is to find a feasible activity that is preferred or indifferent to all other feasible activities. The objective function is assumed to be given as a numerical function on the space of activities; preferred activities correspond to higher values of the function. If the objective function is linear and the constraint set is a polyhedron, 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 a priori numerical objective function readily available. For example, many of the programming problems that arise in the military establishment are not concerned with the maximization of some definite and measurable quantity, such as dollar profit, but rather with the maximization of vague concepts like "military worth." Mathematical techniques for dealing with these problems are useless unless military worth is adequately defined and some method is given for measuring it. Nor are they limited to the military establishment; any organization that does not operate on a profit motive (including any government department) is in a similar situation. Even ordinary profit-making business corporations may face such problems when they must decide among 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 those that involve activities whose main effect is on good will. Also, in a business operation there are often multiple motives that may conflict—for example, advancing the interests of a division or of the whole corporation of which it is a part. In such cases it may well be impossible to measure the objective simply in dollar terms.

1. The word "alternative" is sometimes used in the sense that we use "activity." Other terms that have been used more or less synonymously in various contexts are "outcome" or "prize" (in von Neumann–Morgenstern utility theory), "commodity bundle" (in economics), and "activity vector" (in linear programming).

The object in all these problems is still to find a "best possible" activity 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 activities and solve our problem in terms of this structure. The preferences are presumably those of the individual responsible for making decisions. Hence the name "subjective programming;" the decision is to be made on the basis of subjective preferences rather than objectively measurable quantities. Perhaps "nonnumerical 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.

One general method of attack on subjective programming problems is to follow the natural bent of the mathematician: rather than solve the problem directly, reduce it to one that has already been solved or at least that we know more about. Specifically, the idea is to transform the given "subjective" nonnumerical problem into one of "objective" numerical maximization. The basic tool in this transformation is the notion of *utility*—a numerical function that in some appropriate sense "represents" the given preference structure.

The aim of this chapter is to investigate the implications of this attack, the difficulties that result, and to review and synthesize the work that has been done. There are many open questions in this area. Some are mathematically well defined and must be answered either "yes" or "no;" in others that are more conceptual in nature the chief difficulty at this stage is one of formulation. One of our major purposes here is to put these important problems into a coherent framework, with the hope that they will attract the attention they deserve.

Our approach is novel in two basic respects. First, the preference orders with which we work are in general *not* complete; that is, it is not assumed that all activities are comparable in the preference order. We feel that not only is the completeness assumption not satisfied in practice, but, more important, it is not even defensible as a norm of rational behavior. Heretofore it has been assumed almost universally in discussions of preference; but probably this has been more because of conceptual and mathematical convenience or inertia (previous investigators having assumed it) than out of conviction.

Second, we consider the word "utility" to be a generic name for a number of related but distinct concepts, all of which are tools for solving subjective programming problems. However, different situations call for different utilities. Most subjective programming problems have some inherent mathematical structure other than the order, and to solve them it is convenient to use a utility that not only represents the order but also

"fits" the structure. Different problems have different structures, and this leads to different utilities. For example, the utilities of von Neumann and Morgenstern (1944) are particularly suited to problems involving risk,[2] because they have the "expected utility property," according to which the utility of a risk equals the expected utility of its possible outcomes. This property is a great mathematical convenience in dealing with risky situations. But in problems not involving risk the expected utility property is useless, and it is often more convenient to use a kind of utility other than that of von Neumann and Morgenstern.

As an illustration of this idea, we devote a considerable portion of the sequel to an investigation of a particular class of nonrisky subjective programming problems in which the activities have an additive structure, that is, they can be meaningfully added to one another. It turns out to be convenient in this situation to have available an additive (i.e., linear) utility function; such a utility, when it exists, makes ordinary linear programming techniques applicable. An example is an employee-assignment problem in which the value of having a particular employee in a particular job is not given numerically at the outset; all that is given is a partial preference order on the assignment plans. Conditions for the existence of additive utilities in this situation are then presented. Even in this special area, more questions will be raised than answered.

The basic definitions concerning orders, maximality, and so on, are presented in the following section, with a brief discussion of transitivity and completeness. Subsequent sections are devoted to the problems with additive structures that we mentioned in the preceding paragraph: their treatment is motivated by giving two important examples; conditions for the existence of additive utilities are stated and proved; the assumptions and conclusions from both an intuitive and a mathematical viewpoint are discussed and some of the open problems are pointed out; and an important variant of the basic additivity assumption, about which little is known to date, is outlined. The remainder of the chapter is devoted to a survey of the utility concept, and its application to subjective programming, in the context of various mathematical structures. We discuss the area from a general intuitive viewpoint and give our views concerning the uses and misuses of utility. We go into technical details concerning the mathematical structure of programming problems by means of six examples that describe various contexts in which they may occur. These examples are followed up in four sections, in which utilities are defined and illustrated in each of the six contexts, and their existence, uniqueness,

---

2. Strategic games, though they may not explicitly involve risk, usually do involve it implicitly, because of the need for using mixed strategies to obtain optimal results.

and application to subjective programming is discussed. Finally, we review briefly the outstanding problems that have been mentioned in preceding sections and add two others that have not been mentioned.

The chapter demands only a modicum of mathematical sophistication. Nevertheless, the problems in this area are really fascinating. We have tried to keep the nonmathematical material as independent as possible of the mathematical.

Preference Orders

Let X be a set of points. A *partial preference order* or simply *order* on X is a relation $\gtrsim$ on X, called preference or indifference, which satisfies the conditions of

TRANSITIVITY   *$x \gtrsim y$ and $y \gtrsim z$ imply $x \gtrsim z$*

and

REFLEXIVITY   *$x \gtrsim x$*

for all x, y, z in X. The relations $\succ$ (*preference*) and $\sim$ (*indifference*) are defined in terms of $\gtrsim$: $x \succ y$ if $x \gtrsim y$ but not $y \gtrsim x$; $x \succ y$ if $x \gtrsim y$ and $y \gtrsim x$. If neither $x \gtrsim y$ nor $y \gtrsim x$, then x and y are called *incomparable*; otherwise they are *comparable*. The order is *complete* if it satisfies the condition of

COMPLETENESS   *For all x and y in X, either $x \gtrsim y$ or $y \gtrsim x$ (or both).*

Many of the orders we consider here are *not* complete. An order is called *pure* if the indifference relation is essentially absent, that is, if $x \sim y$ implies $x = y$.

Let A be a subset of X. A point x of A is said to be a *maximum* of A if it is preferred or indifferent to all other members of A, that is, if for all y in A, $x \gtrsim y$. It is said to be *maximal* in A if no member of A is preferred to it, that is, if there is no y in A such that $y \succ x$. Thus a maximal member of A can be advertised only as being "no worse than anything else," whereas a maximum element can be advertised as being "at least as good as anything else." If the order is complete, then "maximum" and "maximal" are equivalent. More generally, if A has a maximum, then all maximal elements are indifferent to it and are, a fortiori, themselves maximum. Often, however, there will be a number of incomparable maximal elements in A and no maximum element. Consider, for example, the pure order on the Euclidean plane E^2 defined by $x \succ y$ if and

Figure 1

only if $x^1 > y^1$ and $x^2 > y^2$. Let A be the triangle defined by $x^1 \geqslant 0$, $x^2 \geqslant 0$, $x^1 + x^2 \leqslant 1$ (see Figure 1). Then all points in A in the line $x^1 + x^2 = 1$ (heavy line) are maximal, but they are all incomparable and A has no maximum.

Of course, not every A need have even a maximal element; for example, if X is the real line ordered in the usual way and A is an open interval, then A has no maximal element.

A real-valued function u on X is said to *represent* the order if for all x, y in X,

$$x \succ y \quad \text{implies} \quad u(x) > u(y)$$

$$x \sim y \quad \text{implies} \quad u(x) = u(y). \tag{1}$$

The representation is called *faithful* if

$$x \gtrsim y \quad \text{if and only if} \quad u(x) \geqslant u(y). \tag{2}$$

It is easily seen that every representation of a complete order is faithful, and, conversely, if an order has a faithful representation, it is complete. *Thus a representation of an incomplete order cannot be faithful*; if $u(x) > u(y)$, x and y may be incomparable. We will say that u *represents the order on A* if (1) holds for all x, y in A.

The most general formulation for a subjective programming problem is the following: given a set X with an order, and a subset A of X, find the maximal elements in A. One way to attack such a problem is to look for numerical functions u that represent the order and try to maximize these functions over A. Clearly, the result will always be a maximal element of A. Little further of interest can be said in this extremely general context. In the sequel it is shown how this principle can be applied in specific situations to the solution of subjective programming problems.

We close this section with an intuitive discussion of the transitivity and completeness assumptions. The transitivity assumption is slightly controversial, especially if the preferences in question are those of a group. But without it a given constraint set may well have no maximal element, even under the best of "regularity" conditions on it (for instance, even if

it is finite), so that the programming problem becomes meaningless. What is worse, without transitivity[3] representation by a numerical function is usually impossible, so that we cannot apply our basic tool. We shall return to this subject on page 254; for the present let us merely say that our considerations apply to those situations in which transitivity does obtain, and this is a large and significant group of problems.

The completeness assumption constitutes a restriction that is probably much more significant than that of transitivity. In spite of the fact that it has been almost universally assumed in discussions of preference, it almost never holds in real-life situations. There is no reason to expect a decision maker to have well-defined preferences between any two possibilities, some of which may be very complicated or hypothetical. The situation illustrated by Figure 1 is a perfect example—the decision maker prefers one point to another if it is larger in both coordinates, but he is unable or unwilling to decide, or is uninterested in deciding, on a substitution rate between the two coordinates. A more detailed discussion of the completeness assumption in a more special context is given by Aumann and Kruskal (1958, pp. 446–447). Here we shall usually not assume completeness. This means that the solutions to our programming problems will be maximal rather than maximum in the constraint set.

Problems with an Additive Structure: Motivation

For the next four sections the space X of activities is the set of all points in Euclidean n-space E^n whose coordinates are nonnegative integers, that is, the set of lattice points in the nonnegative orthant of E^n. *Addition* on X is defined as ordinary vector (coordinatewise) addition.

This is the typical setting of "integer programming" problems, which often, though basically subjective, are "objectivized" by a more or less arbitrary choice of a numerical objective function. Consider the employee-assignment problem mentioned on page 223: There are k jobs and k candidates, and it is desired to assign the candidates to the jobs "optimally." An "activity" (point in X) is a particular assignment plan. Each assignment plan can be thought of as a vector with k^2 coordinates x_{ij}, in which i and j run from 1 to k and x_{ij} is 1 or 0 according as the ith candidate is or is not assigned to the jth job; X can therefore be taken to be the set of lattice points in the nonnegative orthant of Euclidean k^2-space. The constraint set A is then the subset of X satisfying

---

3. More precisely, if the relation $\succ$ has cycles.

$$\sum_{i=1}^{k} x_{ij} = 1, \quad j = 1, \ldots, k$$

$$\sum_{j=1}^{k} x_{ij} = 1, \quad i = 1, \ldots, k.$$

The treatment of assignment problems in the literature often starts out by saying "let us suppose, for simplicity, that each candidate is either totally suited or totally unsuited to a given job." In the first case the assignment of the candidate to the job is given the "value" 1, in the second, the "value" 0; all values for a particular assignment plan are added to yield the "value of the plan," and this total is then maximized over the set A of feasible plans. No justification is given for the addition of the values, and the assumption of total suitability or unsuitability is so restrictive that it makes the whole discussion almost useless. Even in the best of circumstances, that is, even if a given candidate really is either "totally suited" or "totally unsuited" to a given job, it is not clear exactly what has been achieved when the total value has been maximized; has the whole procedure any meaning or is it simply a magic rite to appease the god of operations research?

Sometimes it is assumed that assigning candidate i to job j has value c_{ij}, and it is not assumed that c_{ij} is necessarily 0 or 1. When c_{ij} can be directly measured, say in dollars, then this procedure has some validity. But more often this is not the case, even in business organizations; and then it is not clear how the c_{ij} should be determined nor even what they mean.

A similar problem has been treated at some length in a series of papers and letters to the editor in *Naval Research Logistics Quarterly* (Smith, 1956; Suzuki, 1957; Aumann & Kruskal, 1958, 1959; McShane & Solomon, 1958; Davis, 1959; Kruskal, 1959; Aumann, 1960). A number of ships, or rather types of ships, and a number of types of electronic equipment (such as radar, sonar, radio transmitters, all of various kinds) are given. There is usually more than one ship of each type, and more than one of each type of equipment. The electronic equipment, of which only a limited amount is available, must be assigned to the ships. In some variants of the problem a budget (i.e., sum of money) to buy the equipment is given, rather than the equipment itself; this makes it necessary to decide what equipment to buy in addition to allocating it once it is bought.

The writers of the papers cited assumed that there was some military authority with a preference order on the set of all possible allocation plans and that this preference order could be represented by a numerical

objective function obtained by giving a fixed value c_{ij} to the assignment of an equipment of type i to a ship of type j and then summing all the values for a particular allocation plan. Having made this assumption, they then concentrated attention on practical methods for finding the c_{ij}.

In the following sections we investigate and discuss to what extent and under what conditions the methods described in the foregoing examples can be justified.

Problems with an Additive Structure: Mathematical Treatment

Let X be as in the preceding section, that is, the set of lattice points in the nonnegative orthant of E^n. Assume that there is a partial preference order $\gtrsim$ on E^n satisfying the condition of

$$\text{ADDITIVITY} \qquad \begin{aligned} x \succ y \quad &\text{implies } x + z \succ y + z, \\ x \sim y \quad &\text{implies } x + z \sim y + z, \end{aligned} \qquad (3)$$

for all $z \in x$. A *utility* on X is a function u on X to the reals that represents the preference order and is additive; that is, it satisfies

$$u(x + y) = u(x) + u(y). \qquad (4)$$

An additive function is necessarily linear; for if $x = (x^1, \ldots, x^n)$ and $c^1 = u(1, 0, \ldots, 0), \ldots, c^n = u(0, \ldots, 0, 1)$, then it follows easily from (4) that $u(x) = c^1 x^1 + \cdots + c^n x^n = c \cdot x$. We wish to investigate conditions under which utilities on X exist and what can be done with them once they are found.

The first point to notice is that a utility need not always exist. For a counterexample, let X be the set of lattice points x in the nonnegative orthant of Euclidean 2-space E^2, and let the order be the lexicographic order; that is, $x \succ y$ if and only if either $x^1 > y^1$ or $x^1 = y^1$ and $x^2 > y^2$. It is easily seen that there is then no utility.

To obtain the existence of a utility, we add a further condition, that of finite generation. Intuitively, the order (not the space!) is said to be finitely generated if there is a finite set of preferences and indifferences from which all others can be deduced by means of repeated applications of the transitivity and additivity (3) assumptions. More precisely, we shall say that a preference order $\gtrsim$ is *included* in another one $\gtrsim^*$ if $x \succ y$ implies $x \succ^* y$ and $x \sim y$ implies $x \sim^* y$. A set S of preference statements and indifference statements (statements of the form $x \succ y$ or $x \sim y$) is said to generate a given additive preference order $\gtrsim$ if that order is the minimum (with respect to inclusion) additive order in which the statements of S are true; if S is consistent, that is, if there is any additive order

in which the statements of S are true, then there is also a minimum such order. An additive order generated by a finite set S of preference statements and indifference statements is said to be *finitely generated*.

THEOREM 1   Every finitely generated additive order on X has a utility.

The proof is based on the following:

THEOREM OF THE ALTERNATIVE   Let $z_1, \ldots, z_m$ be arbitrary vectors in E^n, and let $1 \leqslant k \leqslant m$. Then **either** there is a vector c such that

$$c \cdot z_i > 0, \quad i = 1, \ldots, k,$$

and

$$c \cdot z_i = 0, \quad i = k + 1, \ldots, m,$$

or there are nonnegative numbers $r_1, \ldots, r_k$ not all of which vanish, and numbers (possibly negative) $r_{k+1}, \ldots, r_m$, such that

$$\sum_{i=1}^{m} r_i z_i = 0.$$

Note that the last equation is a vector equation. This version of the theorem of the alternative is a consequence, for example, of Tucker's Corollary 2A(ii) (1956, p. 10; Tucker's A_2 must be set equal to zero).

We also make use of the lemma that follows at once from transitivity and additivity.

LEMMA 1 $x \succ y$ and $z \succsim w$ imply $x + z \succ y + w$.

To prove Theorem 1, let S be a set of preference statements generating the order. Let the statements in S be $x_1 \succ y_1, \ldots, x_k \succ y_k, x_{k+1} \sim y_{k+1}, \ldots, x_m \sim y_m$. If $k = 0$, there is nothing to prove, for then there is no strict preference in our order, and 0 is a utility. Suppose therefore that $k \geqslant 1$, set $z_i = x_i - y_i$ for $1 = i, \ldots, m$, and apply the theorem of the alternative. If the first alternative holds, we claim that $c \cdot x$ is a utility. Indeed, let us define an order $\succsim^*$ by $x \succsim^* y$ if and only if $c \cdot (x - y) \geqslant 0$. Since the first alternative holds, the statements of S are true in $\succsim^*$; since $\succsim$ is the minimum order for which the statements in S hold, $\succsim^*$ includes $\succsim$. Hence $x \succ y$ implies $c \cdot x > c \cdot y$ and $x \sim y$ implies $c \cdot x = c \cdot y$, and so $c \cdot x$ is a utility.

On the other hand, suppose the second alternative holds. We may suppose that all the r_i are nonnegative; for, if some r_j is negative, we may reverse the roles of x_j and y_j to obtain r_j positive. Moreover, we may suppose that the r_i are rational; for the z_i are all rational (in fact integral),

and the proof of the theorem of the alternative holds without change in the rational field. (Alternatively, it is easily seen that a linear system with rational coefficients that has any solution also has a rational solution.) It follows that we can take the r_i to be nonnegative integers, for we can multiply through by a common denominator. From $x_1 \succ y_1, \ldots,$ $x_k \succ y_k, x_{k+1} \sim y_{k+1}, \ldots, x_m \sim y_m$, Lemma 1, and the fact that not all the $r_1, \ldots, r_k$ vanish, it follows that

$$\sum_{i=1}^{m} r_i x_i \succ \sum_{i=1}^{m} r_i y_i;$$

but this contradicts the second alternative, since $\succ$ is by definition irreflexive. This completes the proof of Theorem 1.

As we have already pointed out, if a function u represents the order, then maximization of u over a constraint set A leads to a maximal element of A. If u is a utility, then it is linear as well as representing the order; this means that the problem has been reduced to a numerical "integer programming" problem, for which an algorithm[4] is available (Gomory, 1958).

The finite generation condition on the preference order clearly holds in many real-life situations, and it is to these situations that Theorem 1 is most obviously applicable. For the purposes of solving subjective programming problems, Theorem 1 can often be applied even when the order is not finitely generated. In fact, we are interested only in the behavior of the order on the constraint set A. Maximization of a function that represents the order on A yields a maximal member of A, even though the representation may fail outside A. Now if A is finite—as it usually is—we can look at the set S of all preferences and indifferences between members of A. This set is certainly finite, and it is consistent because all the statements in S hold in the given order $\succsim$ on X. Hence S generates an order $\succsim'$ on X; $\succsim'$ coincides with the original order $\succsim$ on A, but it is finitely generated, even though $\succsim$ may not be. Theorem 1 assures us that $\succsim'$ has a utility; for the purpose of solving the given subjective programming problem, this utility is just as good as a utility for the original order $\succsim$ would have been. We have demonstrated

THEOREM 2   Let $\succsim$ be an additive preference order on X, and let A be a finite subset of X. Then there is a linear function u that represents $\succsim$ on A. A fortiori, the maximization of u over A leads to a maximal element of A.

---

4. Certain restrictions must be placed on the constraint set A in order for the Gomory algorithm to be applicable.

We shall call u a utility on A or an A-utility. Practical methods for finding A-utilities are identical with those for finding utilities (see p. 243).

It should be noted that the utilities, and a fortiori the A-utilities, need not be unique. A simple counterexample is the pure order generated on the lattice points in the nonnegative orthant of E^2 by $(1,0) \succ 0$, $(0,1) \succ 0$. All functions of the form $c^1 x^1 + c^2 x^2$, where $c^1 > 0$ and $c^2 > 0$, are utilities.

Problems with an Additive Structure: Discussion

The space X, on which the given preference order and the resulting utility are defined, consists of *all* the activity vectors, not just those in the constraint set. 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 one to the other. The housewife, for example, may be able to express meaningful preferences 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 is ordinarily defined on many pairs of vectors that are not in the constraint set, and this is perfectly reasonable.

The Additivity Assumption

It is this assumption that enables us to find linear functions that represent the order. There is no doubt that it is a strong assumption, that it severely limits the range of problems to which the theory developed in the preceding section may be applied. In particular, unlike transitivity, it cannot be considered a generally applicable "rationality" assumption. Roughly, it holds whenever the usefulness of an activity does not depend on other activities being performed at the same time. Otherwise, it will usually fail: for example, four tires may be preferred to a Cadillac without tires (at least the tubes can be used for swimming), but many would prefer a Cadillac with tires to eight tires; so we have reversed preferences simply by adding four tires. Other examples in which the additivity assumption does not hold are those governed by a law of diminishing returns, or personnel assignment problems in which compatibility considerations play an important part.

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

1. *Problems in which the various activity units operate entirely independently of one another*. For instance, consider an employment agency with

a number of candidates and a number of firms on its roster; 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.

2. *Programming problems in which the interactions between the activity units are relatively minor and difficult to analyze.* Allocation of clerical and other semiskilled workers in a single organization might be an example. Another example of the same kind is the naval electronics problem (see p. 237).

3. *Certain situations governed by a law of diminishing returns, but in which we are interested only in adding to current activity in relatively small amounts.* Intuitively, 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 description at the end of the next section.

Another aspect of the additivity assumption 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 size of the activities involved in preference statements. For example, the housewife, though willing to express preferences 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 have to be regarded as an idealization of the true situation.

The Converse of Theorem 1

We are particularly concerned with the relation between the additivity assumption and the existence of a utility in the sense described on page 238. Assume that a given preference relation has a utility; must the additivity assumption be satisfied? The answer is in general no[5]; 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

1. $x \succ y$ implies either that $x + z \succ y + z$ or that $x + z$ and $y + z$ are incomparable; and

2. $x \sim y$ implies either that $x + z \sim y + z$ or that $x + z$ and $y + z$ are incomparable.

5. Let X be the nonnegative integers and suppose that the preference relation contains (in addition to the statements $x \gtrsim x$) only the single statement $1 \gtrsim 0$ (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.

Every preference relation for which there is a utility must be 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 one that is transitive and additive. The proof of the converse is trivial.

Significance of the Converse

The converse says that we cannot expect a linear objective function for a programming problem unless additivity holds. On page 241 we discussed the intuitive meaning of additivity and pointed out that it is a strong assumption. We now see that though it is strong it necessarily holds in any situation in which linear programming is even potentially applicable. In particular, this should serve as a warning to those who blithely add together the "worths" of components in order to get the "worth" of a whole system; such a procedure is valid only in certain circumstances (see p. 241).

Finding All the Maximal Elements of the Constraint Set

As we have pointed out, it is obvious that maximization of a utility over the constraint set yields a maximal element. The question is whether every maximal element can be obtained in this way. In this context the answer is no. Consider, for example, the pure order on the nonnegative integers generated by $3 \succ 2 \succ 0$. This order is additive and we have $0 \prec 2 \prec 3 \prec \cdots$, but 1 and 2 are incomparable. An essentially unique utility is $u(x) = x$. If we let $A = \{1, 2\}$, then 1 is maximal in A, but no utility maximizes it.

Conditions under which the question can be answered positively are discussed on page 253.

Computing the Utilities

At the end of the last section we mentioned that in general there is no unique utility; thus what we are looking for is the set of all utilities. Let $u(x) = c \cdot x$ be a utility; we have already remarked that all utilities have this form. Then if $x \succ y$, it follows that $c \cdot (x - y) > 0$; if $x \sim y$, it follows that $c \cdot (x - y) = 0$. If the order is finitely generated, then there is a finite set S of preference statements and indifference statements that generate the order. Form the system S^* of strict inequalities and equations corresponding to the statements in S; any feasible solution of this system is a utility, and conversely any utility is given by a feasible solution of the system.

If the order is not finitely generated, the same idea may be applied to obtain an A-utility. Of course, it is not necessary to use all comparisons within A; a set of comparisons that generates the order on A would suffice.

From a practical viewpoint the question of computing the utility is often more complicated. In this chapter we wish to avoid a detailed discussion of the practical difficulties[6]; we merely mention one which has some theoretical interest and significance. The set S of basic (generating) decisions may be extremely large; it is often impractical to ascertain all of these decisions explicitly. The result of having a set of available decisions 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 questions are 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, but in its use as a tool for solving programming problems. Application of the approximate utility in a programming problem may be expected to lead to an answer that is not generally 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. We feel 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 of 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?

Generalizations are discussed in later sections (see pp. 245–53).

The Restricted Additivity Assumption

We wish to pursue further the question raised on pp. 241–42 in the preceding 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

---

6. The reader is referred to the papers cited on page 237 in connection with the naval electronics problem for a discussion of some of the problems.

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 Theorem 1 still holds in this situation, that is, 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 gone beyond A (and thus beyond the responsibility of the decision maker).

There are several directions in which we might proceed 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 points), 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. None 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 $x_i \geqslant 0$ for $i = 1, \ldots, 4$, and $\sum_{i=1}^{4} x_i \leqslant 2$. Write $a = (1,0,0,0), \ldots, d = (0,0,0,1)$; the order on A is defined by $0 \prec a \prec b \prec c \prec d \prec 2a \prec a + b \prec 2b \prec a + c \prec a + d \prec b + c \prec 2c \prec b + d \prec c + d \prec 2d$; 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 $2b \prec a + c$. Thus even a complete order on a perfectly "well behaved" set does not always possess a utility.

Another counterexample has been given by Kraft, Pratt, and Seidenberg (1959). They exhibited an additive order on the unit cube in E^5 (i.e., the set of all lattice points in E^5 whose coordinates are 0 or 1) in connection with an attempt to define numerical subjective probabilities from qualitative comparisons of probabilities of the kind "this is more probable than that."

Possibly positive results can be obtained if some other kind of geometric shape 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 likely to occur in practice). In general, A contains B. It seems intuitively clear, and indeed is not difficult to prove, that if A is large enough with respect to B (but still finite) 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?

The problem of defining a utility on B is equivalent to that of extending the given order to all of the lattice points in E^n; this follows from Theorem 1. Thus the foregoing problem may be formulated as follows: let B be a fixed 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" (see p. 241). 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," that is, 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, 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. Since we are concerned only with the marginal range, the unrealism of the preference order outside it need not disturb us.

General Preference Structures and Utilities

We have already investigated a situation in which preferences are defined on a space with an additive structure. These preferences were assumed to "fit" the additive structure—that was the content of the additivity assumption (3). It was found that to solve subjective programming problems in this context, it was possible and convenient to make use of functions that represented the preference order and that also "fitted" the additive structure (4). Such functions were called "utilities."

That situation is typical. Subjective programming problems are usually set in the context of a space with, in addition to the preference order, an underlying mathematical structure of a different kind. The preferences are assumed to "fit" this structure in some sense; to solve subjective programming problems, it then turns out to be possible and convenient to use functions that represent the preference order and also fit the structure. We propose to call such functions "utilities," in whatever context they occur; thus the precise meaning of the word varies from context to context, depending on the underlying structure.

Actually this is not a new proposal but merely a confirmation of existing usage. Much of mathematical economics is set in the context of a space of commodity bundles; preference orders on this space satisfy a

certain continuity condition, and a utility is defined to be a function that represents the order *and* is continuous. The "structure" in this case is a topological one; roughly, this means that there is defined in the space a notion of nearness on which the notion of continuity is based. In game theory and statistics, on the other hand, the space consists of lotteries, that is, risky outcomes. Here the structure is an algebraic one, which allows us to combine lotteries to obtain new ones. Preference orders are required to "fit" this structure, as are utilities; the fact that the utilities "fit" is expressed by the expected utility hypothesis. Our usage of the word "utility" in the preceding sections, which corresponds to the occasional use of the word "utility" for "objective function" in numerical (objective) programming problems, is again different.

Much confusion and pointless polemic could be avoided if it were realized that the word "utility" is used with different (but related) meanings in different, contexts and that this is perfectly proper. Utilities are tools, and different tools are needed to solve different problems. Nobody would think of saying "the only *real* axe is the climber's ice-axe; a woodsman's axe is not really an axe at all." Still less would anybody use an ice-axe to chop down a tree, when both kinds are available; nor does the question "how should an axe be built?" have meaning when the job to be done with it is not specified. It is just as meaningless to try to "measure" utility or military worth without taking into account how these measures are to be used. These facts should be obvious once they are pointed out, and they have been pointed out repeatedly. Unfortunately, they are still not generally understood, and much of what has been and is being written about utility and related notions is rendered meaningless by the failure to understand them.

As an example of what we have been saying, consider a problem with an additive structure, say an assignment problem (see pp. 235–43); suppose the additivity assumption (3) holds. Now imagine that instead of defining an additive utility in the sense described on page 238 and solving the problem with such a utility, we were to try to apply von Neumann–Morgenstern utilities. This could, in fact, be done. Each feasible activity would have to be considered, as well as the set of all lotteries whose outcomes are feasible activities. The original preference order on the activities would have to be extended to the set of all lotteries over the activities; the extended preference order would yield a von Neumann–Morgenstern utility. Since, in particular, this utility represents the original order, its maximization would actually solve the original problem, that is, it would yield a feasible assignment that is maximal in the original preference order. So the attempt to use N–M utilities here would seem to have been crowned with success.

In fact, though, such an attempt would be both highly impractical and completely absurd. It would be necessary to ask many questions of the decision maker in order to establish preferences between the lotteries, that is, risky outcomes; all of them would be irrelevant to the original problem, which contained no element of risk. Moreover, the resulting utility would presumably be nonlinear, which would make the numerical solution much harder, if not completely impractical. For example, although the decision maker has additive preferences, he might be unwilling to take risks; thus he might be indifferent between a sure prize of 3 activity units on the one hand and a 1/2–1/2 lottery of 0 activity units and 10 activity units on the other hand. This yields $u(10) = 10$, $u(3) = 5$, $u(0) = 0$, a nonlinear function. The nonlinearity has nothing to do with the original problem; it is a result of the decision maker's attitude to risk. This is irrelevant, since neither the original problem nor its solution contains risky alternatives. On the other hand, use of a utility that fits the relevant structure of the problem, that is, an additive (linear) utility, would permit the application of standard linear programming techniques.

Of course, this works both ways; where an N–M utility is called for, an additive utility should not be used. Consider, for example, a two-person game whose outcomes happen to be imbedded in a space with an additive structure; suppose that the preference orders of both players on this space obey the additivity assumption. It might be tempting to use additive utilities (which are usually easier to calculate than N–M utilities) as entries in the payoff matrix of such a game. But "optimal" mixed strategies calculated on this basis would be completely meaningless and, in fact, not optimal. A special case of this mistake is setting the N–M utility of money equal to money. Similar remarks apply to many problems in statistical decision making.

The remainder of this chapter is devoted to a review, from the mathematical viewpoint, of various spaces on which there is defined both an underlying mathematical structure and an order that in some sense "fits" or is "consistent with" this structure. Such spaces are called "preference spaces;" this is therefore a loose generic term with an intuitive meaning only. We are particularly interested in utilities on such spaces and in subjective programming problems on them. In the following section the more important examples of such spaces are presented in some detail; these examples have been chosen because of their applications to various problems of human judgments and optimality. In subsequent sections the concepts and methods in which we are interested are applied to each of these examples in turn, with the object of bringing out the parallels that exist between the various kinds of preference spaces.

Examples of Preference Spaces

The symbol X is used to denote a preference space.

1. In this example X is assumed to be a *topological space*[7] on which there is defined a preference order. The consistency requirement imposed is that the preference order be "continuous," that is, if $x_n \to x$, $y_n \to y$, and $x_n \gtrsim y_n$ for[8] all n, then $x \gtrsim y$; an equivalent formulation is that the set $\{(x, y) : x \gtrsim y\}$ is closed in the product space $X \times X$. When the order is complete, this in turn is equivalent to saying that the sets $\{x : x \gtrsim y\}$ and $\{x : y \gtrsim x\}$ are closed[9] in X. However, this is not so for a partial order. This situation has been extensively studied in economics; X is then the nonnegative orthant of the Euclidean space E^n of n dimensions, and the members of X represent commodity bundles. The sets $\{x : x \sim y\}$ are the "indifference curves." The topology is the ordinary Euclidean topology, that is, it is induced by the Euclidean metric. For a detailed treatment see Debreu (1959).

2. X is an *Abelian semigroup*, a space on which there is defined a commutative and associative addition. On X there is also defined a preference order, which is assumed to be "consistent" with the semigroup structure in the sense that $x \succ y$ implies $x + z \succ y + z$, and $x \sim y$ implies $x + z \sim y + z$ (this is what we called "additivity" (3)). An example of an Abelian semigroup is the set of lattice points in the nonnegative orthant of E^n; this example and its application have already been discussed in some detail. We shall see that much of this discussion remains valid in the more general context of an arbitrary Abelian semigroup. It also applies to Abelian *groups*, which are special cases of semigroups. An example of a group is the set of *all* lattice points in a Euclidean space (not just those in the nonnegative orthant).

7. Those readers not familiar with this term may substitute the notion of *metric space*, that is, space on which the notion of "distance between two points" is defined.

8. If X is not separable, then the convergence must be taken in the sense of Moore–Smith (Kelley, 1955); that is, $\{x_n\}$ and $\{y_n\}$ are nets rather than sequences. The applications usually involve subspaces of Euclidean spaces, which are all separable.

9. Suppose that $\{x: x \gtrsim y\}$ and $\{x: y \gtrsim x\}$ are closed. Then by completeness, $\{x: x \succ y\}$ and $\{x: y \succ x\}$ are open. Let $A = \{(x, y): x \succ y\}$ and suppose $(x, y) \in A$, that is, $x \succ y$. Either there is a z such that $x \succ z \succ y$, or for all $z \in X$ either $z \gtrsim x$ or $y \gtrsim z$ (again by completeness). In the first case $\{w: w \succ z\} \times \{v: z \succ v\}$ is an open neighborhood of (x, y) entirely contained in A. In the second case $\{w: w \gtrsim x\} = \{w: w \succ y\}$, and so the right side is open; similarly $\{v: y \gtrsim v\}$ is open. Then $\{w: w \gtrsim x\} \times \{v: y \gtrsim v\}$ is an open neighborhood of (x, y) entirely contained in A. In any case, A is open, hence its complement $\{(x, y): y \gtrsim x\}$ is closed. The proof of the converse is trivial.

3. X is a *mixture space*, that is, a space in which convex combinations are defined. More precisely, if $x_1, \ldots, x_n \in X$ and $\alpha_1, \ldots, \alpha_n$ are nonnegative real numbers summing to 1, then the *convex combination* $\sum_{i=1}^{n} \alpha_i x_i$ is defined as a member of X. Convex combinations are assumed to obey all the rules of ordinary vector space sums, but the coefficients are always nonnegative and sum to 1. The idea of convex combination has topological as well as algebraic aspects: the topology of the space of coefficients α is a significant structural element. On X there is defined a preference order, on which two consistency requirements are imposed, corresponding to the algebraic and topological aspects of the structure. The algebraic requirement is that for $0 < \alpha \leqslant 1$ and for all $z, x \gtrsim y$ if and only if $\alpha x + (1 - \alpha)z \gtrsim \alpha y + (1 - \alpha)z$. The topological requirement is that $\{\alpha : \alpha x + (1 - \alpha)y \gtrsim z\}$ is closed (Herstein & Milnor, 1953). This set-up is most familiar in connection with the utility theory of von Neumann and Morgenstern (1944). There the members of X are probability combinations of "prizes" or "pure outcomes," and the order is complete. Aumann (1962) has investigated a variant of this theory in which the order is partial; in that investigation the topological requirement was somewhat weaker than the one quoted here. Another example of a mixture space is the set of all fixed amounts of, say, liquids; convex combinations are taken by mixing in the appropriate proportions, and the space is ordered by the relation "hotter than."

4. X is a *vector space* over the real field. This is a variant of example 2, corresponding to subjective noninteger programming problems. A preference order is imposed on X, and, in addition to the consistency requirement imposed in example 2, we require here that $x \gtrsim y$ imply $\alpha x \gtrsim \alpha y$ for all positive scalars α. This is a purely algebraic requirement, making no use of the topology of the real line.

Since a vector space over the reals is, in particular, a mixture space, we may ask what the connection is between the requirements of examples 3 and 4. The answer is that the requirements of 3 are stronger; for example, they exclude the lexicographic order on E^2, whereas this order is admitted under 4 (under the lexicographic order on E^2, $x \gtrsim y$ if and only if either $x^1 > y^1$ or $x^1 = y^1$ and $x^2 \geqslant y^2$). Hausner (1954) has investigated completely ordered vector spaces. He has also investigated the more general situation of completely ordered mixture spaces in which no topological requirement is imposed.

This example could be generalized slightly by substituting a "semimodule" for a vector space. A semimodule (over the positive real numbers, say) is a set that is provided with a commutative and associative addition and a multiplication by positive real scalars, satisfying the appropriate conditions. The consistency requirements are as above. This

corresponds to (nondiscrete) programming problems for nonnegative variables, but little is gained from this generalization.[10]

5. A degenerate example is the one in which X is simply a point set. The consistency requirement for the preference order is then vacuous, and we simply have a set with an order.

6. If we let X be the constraint set of pp. 244–45, we will have a space such that $x + y$ is defined for some, but not all, pairs of points x, y in X; when it is defined, the addition is associative and commutative. For preference orders on such X the consistency requirement in example 2 is imposed, but it is required to hold only when all the terms appearing in the requirement are meaningful, that is, members of X.

Ordered mathematical structures have been the objects of mathematical interest in contexts that are not related to the notion of preference, and there is a large literature on the subject.[11] Workers in the field of optimality and human judgments would do well not to ignore this work.

We close this section with the remark that the notion of "consistency" or "meshing" between two kinds of structures imposed on a space is quite common in mathematics. Thus in a ring—a system in which both multiplication and addition are defined—the consistency requirement is the distributive law. In a topological group it is required that the group operation be continuous. Preference spaces are mathematical structures on which two kinds of structures are defined—one a preference order—and the two structures are assumed to "mesh."

Utilities

A utility u on a preference space X is a function that represents the order on X and fits the underlying structure of the space. Thus, like "preference space," it is a loose generic term whose precise meaning varies from context to context. When X is a topological space, u is continuous. When X is a semigroup, u is a homomorphism, that is, $u(x + y) = u(x) + u(y)$. When X is a mixture space, u has the expected utility property, that is,

10. Added in proof: A different generalization is obtained by replacing the real field by an arbitrary ordered field (compare, for example, Charnes & Cooper, 1961, pp. 280 ff.). Utilities, too, may sometimes be generalized to have values in an ordered field, rather than real values.

11. See, for example, papers in *Mathematical Reviews* listed under "Order, lattices," under "Groups and generalizations—ordered groups and semigroups," and under "Functional analysis—partially ordered vector spaces," or under similar subject headings (the headings change slightly from year to year).

$$u(\alpha x + (1 - \alpha)y) = \alpha u(x) + (1 - \alpha)u(y).$$

When X is a vector space, u is linear, that is, u is a homomorphism and $u(\alpha x) = \alpha u(x)$. When X has no underlying structure, no requirement is placed on u (other than that it represents the order). When X is a system in which $x + y$ is sometimes defined, we have $u(x + y) = u(x) + u(y)$ whenever $x, y, x + y \in X$.

A utility represents the preference order, but not in general faithfully; that is, it gives partial, not complete, information about the original preference order. Given a utility on a space X, we cannot always recon-struct the preference order from it. This is entirely typical of mappings on mathematical structures: a continuous mapping does not determine a topological space, a homomorphism does not determine a group, and so on. The typical situation is that the image space—the real line (or part of it) in our case—has more structure, more relations than the original space. If two points in a topological space are "close" to each other, then so are their images under a continuous function; but the images can be close without the original points being close. If $x = y + z$ in an Abelian group G and u is a homomorphism, then $u(x) = u(y) + u(z)$; but the lat-ter equation may hold even when the former does not. Similarly, though $x \succ y$ implies $u(x) > u(y)$ for a utility u, we may have $u(x) > u(y)$ with-out $x \succ y$.

When the order is complete, then every utility *does* determine the pref-erence order. This is the classical situation, as treated, for example, by von Neumann and Morgenstern (1944), Debreu (1954, 1959), Hausner (1954), Herstein and Milnor (1953), and so on. All of these authors demanded faithful representation; but they were concerned only with complete orders, and, for these, representation and faithful representation are equivalent. Until recently the term "utility" was reserved for com-plete orders; its first use in connection with partial orders was made by Aumann (1962).

Existence of Utilities

Once the notion of utility has been defined, it is natural to ask whether a given type of preference space has a utility. When X is a completely ordered separable[12] topological space, Debreu (1954) has shown that there is a utility. When X is a mixture space with complete order, the

---

12. Having a countable basis; Debreu calls this "perfectly separable." All subsets of Eucli-dean spaces have this property.

existence of a utility was first demonstrated by von Neumann and Morgenstern (1944). For partial orders on a finite dimensional[13] mixture space the existence of a utility was established by Aumann (1962); on the other hand, when the space is not finitely generated, there may be no utility.[14] We have seen (p. 238) that a preference-ordered semigroup need not have a utility; essentially the same example (the lexicographic order) yields a counterexample for the existence of a utility in a preference-ordered vector space. However, we know that a utility can be defined when the preference order is finitely generated. The same holds in general: *let X be an Abelian semigroup, or a vector space, on which there is imposed a finitely generated preference order consistent with the structure of X. Then there is a utility on X.* Here "finitely generated" has the obvious meaning corresponding to that given on page 238 (the precise meaning is different for semigroups and vector spaces). The proof for semigroups proceeds by reducing the situation to that given on page 238; for vector spaces, it proceeds by analogy but turns out to be even simpler.

It would be interesting to know whether Debreu's result for topological preference spaces with complete orders (quoted above) could be extended to the general case of partial orders.[15] Development of utility theory for infinite dimensional partially ordered topological linear spaces, and for the corresponding mixture spaces, would also be desirable. In this connection see Aumann (1962, p. 461).

Uniqueness of Utilities

For completely ordered preference spaces, utilities, when they exist, are in some sense unique. The usual statement is that the utilities are unique "up to" some class of functions. More precisely, there is usually a class G of strictly increasing functions from the real line into itself such that

(5) if u and v are utilities, then there is a g in G such that $v = gu$ (where gu is the composition of g and u);

(6) if u is a utility and g is in G, then gu is also a utility.

13. Roughly, one that is generated by a finite number of "pure" prospects (see Aumann, 1962).

14. Added in proof: For additional work on spaces that are not finitely generated, see Kannai (1963).

15. Added in proof: The following example due to B. Peleg (unpublished) shows that it cannot: X is the unit circle in the plane (boundary only). The order is given by: $x \gtrsim x$ and $x \gtrsim (1,0)$ for all x in X; $(\alpha, \beta) \gtrsim (-\alpha, \beta)$ for all (α, β) in X such that $\alpha > 0$ and $\beta > 0$. Possibly, if the definition of continuity is strengthened, a positive result can be obtained.

The class G of functions, like the utilities themselves, depends on the kind of preference space under consideration. Thus for topological spaces[16] it consists of increasing continuous functions; for Abelian semi-groups, of increasing additive functions (which are a fortiori linear); for mixture spaces and vector spaces, of increasing linear functions; and for unstructured spaces, of arbitrary increasing functions.

When the preference space is only partially ordered, there is no uniqueness theorem. In the extreme case, when the order is empty (i.e., $x \succsim y$ never holds), any function that is consistent with the underlying structure is a utility. See also the remark following Theorem 2 (p. 240).

Programming Problems

Let X be a preference space and A a subset of X. It is desired to find an element of A that is maximal in the order on X. Since every utility represents the preference order, it is obviously sufficient to maximize a utility over A to find a maximal element of A. For this purpose it is not even necessary that the utility be consistent with the underlying structure. The requirement of consistency is imposed on utilities only because it is more convenient to work with functions that represent the order and fit the structure of the space than just with representing functions.

When the order is complete, every utility constitutes a faithful representation, so that every maximum element of A can be obtained by maximizing any utility over A. When the order is not complete, it may happen that some maximal elements of A are achieved as the maxima of utilities, whereas others are not. An example is given on page 243. A less pathological example is the following: let X be the Euclidean plane considered either as a vector space or as a mixture space; let the order be given by $x \succsim y$ if and only if $x^1 \geqslant y^1$ and $x^2 \geqslant y^2$ (whence $x \succ y$ if and only if $x^1 \geqslant y^1$, $x^2 \geqslant y^2$, and $x \neq y$); and let A be the unit disk. The utilities are the functions of the form $c^1 x^1 + c^2 x^2$, $c^1, c^2 > 0$ (plus possibly a constant in case X is considered as a mixture space). The point $(1, 0)$ is maximal in A, but it is not a maximum of any utility.

We would like to find conditions on A and/or the preference order which exclude this possibility, that is, under which

(7) an element x of A is maximal in A if and only if there is a utility u on X that is maximized over A at x.

---

16. In this case (5) holds only when X is connected.

For mixture spaces it has been shown (Aumann, 1962, Theorem B) that (7) holds whenever A is a convex polyhedron.[17] It can also be demonstrated for vector spaces when A is a convex polyhedron and the preference order is finitely generated. The topological case remains unexplored. It would be interesting to see results in this direction.

Index of Open Questions

This is simply a listing of the open questions mentioned in the foregoing discussion, to which we add two that we have not covered.

Computing the utilities for problems with an additive structure (see p. 243). Existence of utilities under the restricted additivity assumption, either by limiting the shape of the constraint set or by surrounding it with a sufficiently large neighborhood (the second approach seems more hopeful—see p. 244). Topological preference spaces: existence of utilities for partially ordered spaces (see p. 252), finding all maximal elements of the constraint set. Development of utility theory for partially ordered topological linear spaces and the corresponding mixture spaces (see p. 251 and footnote 14).

We close with a statement of two additional questions. On page 235 we discussed briefly the transitivity condition. Can this condition be dispensed with? If the preference order contains "significant inconsistencies"—if, for example, all members of the space X are members of one large cycle—then clearly the programming problem has no meaning. It could happen, though, that only "small inconsistencies" are in the order and that it would still be meaningful to look for an "approximately optimal" member of the constraint set. How can this notion be formalized, and once it is formalized how can these "approximately optimal" members be found? One suggestion for the second question is to define an "ε-representation" as a function u so that $x \gtrsim y$ implies $u(x) \geq u(y) - \varepsilon$; precautions must be taken to avoid trivial ε-representations. One can then work with "ε-utilities," that is, ε-representations that "fit" the structure of the space.

On page 251 we remarked that in case the order is not complete it is usually not determined by a given utility. Our second question is, to what extent, and under what conditions, does the set of *all* utilities determine the order? For mixture spaces, this was discussed in Aumann (1962,

---

17. The demonstration given in Aumann (1962) is incorrect and the theorem as stated there is false. However, the theorem is correct if the consistency requirement in Example 4 stated on page 249 is used rather than the weaker one in Aumann (1962 [2.2]). A correction will be published (Aumann, 1964).

section 7). What can be said for topological preference spaces? The question is similar in principle to questions occurring in algebra (group representations or characters) and in functional analysis (reflexivity of Banach spaces).

References

Aumann, R. J., & Kruskal, J. B. (1958). The coefficients in an allocation problem. *Naval Res. Logistics Quart.*, **5**, 111–123 [Chapter 11].

Aumann, R. J., & Kruskal, J. B. (1959). Assigning quantitative values to qualitative factors in the naval electronics problem. *Ibid.*, **6**, 1–16 [Chapter 12].

Aumann, R. J. (1960). Letter to the Editor. *Ibid.*, **7**, 85–88.

Aumann, R. J. (1962). Utility theory without the completeness axiom. *Econometrica*, **30**, 445–462 [Chapter 14a].

Aumann, R. J. (1964). Utility theory without the completeness axiom: a correction. *Econometrica*, **32**, 210–212 [Chapter 14b].

Charnes, A., & Cooper, W. W. (1961). *Management models and industrial applications of linear programming.* New York: Wiley.

Davis, R. G. (1959). Letter to the Editor, *Naval Res. Logistics Quart.*, **6**, 183–185.

Debreu, G. (1954). Representation of a preference ordering by a numerical function. In R. M. Thrall, C. H. Coombs, & R. L. Davis (Eds.), *Decision processes.* New York: Wiley. Pp. 159–165.

Debreu, G. (1959). *Theory of Value*, New York, Wiley.

Gomory, R. (1958). Outline of an algorithm for integer solutions to linear programs. *Bull. Amer. Math. Soc.*, **64**, 275–278.

Hausner, M. (1954). Multidimensional Utilities. In R. M. Thrall, C. H. Coombs, & R. L. Davis (Eds.), *Decision processes.* New York: Wiley. Pp. 167–180.

Herstein, I. N., & Milnor, J. (1953). An axiomatic approach to measurable utility. *Econometrica*, **21**, 291–297.

Kannai, Y. (1963). Existence of a utility in infinite dimensional partially ordered spaces. *Israel Jour. Math.* **1**, 229–234.

Kelley, J. (1955). *General topology.* Princeton, N.J.: Van Nostrand.

Kraft, C. H., Pratt, J. W., & Seidenberg, A. (1959). Intuitive probability on finite sets. *Ann. Math. Statist.*, **30**, 408–419.

Kruskal, J. B. (1959). Letter to the Editor. *Naval Res. Logistics Quart.*, **6**, 261.

McShane, R. E., & Solomon, H. (1958). Letter to the Editor. *Ibid.*, **5**, 363–367.

Smith, J. W. (1956). A plan to allocate and procure electronic sets by the use of linear programming techniques and analytical methods of assigning values to qualitative factors. *Naval Res. Logistics Quart.* **3**, 151–162.

Suzuki, G. (1957). Procurement and allocation of naval electronic equipment. *Naval. Res. Logistics Quart.*, **4**, 1–8.

Tucker, A. W. (1956). Dual systems of homogeneous linear relations. In H. W. Kuhn & A. W. Tucker (Eds.), *Linear inequalities and related systems. Ann. Math. Study* 38. Princeton, N.J.: Princeton Univer. Press. Pp. 3–18.

Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior.* Princeton, N.J.: Princeton Univer. Press.