COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut

COWLES FOUNDATION DISCUSSION PAPER NO. 171

CONVEXITY IN THE THEORY OF CHOICE UNDER RISK

Menahem E. Yaari

May 21, 1964

CONVEXITY IN THE THEORY OF CHOICE UNDER RISK

Menahem E. Yaari*

I. Introduction

Consider a family $\{E_1, \ldots, E_n\}$ of mutually exclusive and exhaustive events on some sample space S . A gamble $\langle x_1, \ldots, x_n \rangle$ is a contract which promises its owner x_1 dollars if E_1 occurs, x_2 dollars if E_2 occurs, and so on. We shall refer to gambles as being defined on the family $\{E_1, \ldots, E_n\}$, because each gamble is clearly a function from this family to the set of real numbers.

It is possible to regard gambles which are defined on $\{E_1, \ldots, E_n\}$ simply as commodity bundles in a world where there are n commodities, the first being "dollars if E_1 " the second being "dollars if E_2 " and so on. Having done this, one is tempted to proceed formally and use the conventional theory of choice among commodity bundles as a theory of choice among gambles. To this end, one assumes that a relation called the preference or indifference relation, and written $\gtrsim$, is defined on the set of all gambles and that

---

* William Brainard has narrowly escaped the co-authorship of this paper. I would like to thank William J. Fellner and my colleagues at the Cowles Foundation for their comments and for their help in obtaining some of the data. Part of the work leading to this paper was supported by the Office of Naval Research, under Contract Nonr-3055(00).

this relation satisfies the following five axioms:

(1) Transitivity: Given any three gambles, x, y and z, if $x \gtrsim y$ and $y \gtrsim z$ then $x \gtrsim z$.

(2) Completeness: For any two gambles, x and y, either $x \gtrsim y$ or $y \gtrsim x$.

(3) Continuity: For any gamble x, the sets $\{y | y \gtrsim x\}$ and $\{y | x \gtrsim y\}$ are closed.

(4) Monotonicity: For any two gambles, x and y, if $x \geqq y$ then $x \gtrsim y$. 1/

(5) Convexity: For any gamble x, the set $\{y | y \gtrsim x\}$ is convex.

These are the axioms which are usually written down without much ado, with the term "commodity bundle" replacing the word "gamble." There is one axiom, however, which is usually assumed to hold for gambles and not for conventional commodity bundles. Several versions of it exist (some weaker than others) and it usually goes by the name of the independence axiom or the dominance axiom. We shall use the version which corresponds to the latter of these terms.

To set the stage, suppose that the family $\left\{E_1, \ldots, E_n\right\}$ is partitioned into two sub-families $\left\{E_{11}, \ldots, E_{1k}\right\}$ and $\left\{E_{21}, \ldots, E_{2m}\right\}$ where $k + m = n$ and each of the E_{ij}'s is actually one of the original E_i's. Gambles which are defined on such sub-families are called conditional

---

1/ By $x \geqq y$ we mean that every component of x is greater than or equal to the corresponding component of y.

gambles, and it is assumed that the set of all conditional gambles defined on a given sub-family is ordered by a preference-or-indifference relation, for which the symbol $\gtrsim$ is also used. Now if x' is a conditional gamble defined on a given sub-family and x'' is a conditional gamble defined on its complement, then the pair $\langle x', x'' \rangle$ (often called a hyper gamble) can clearly be looked upon as a gamble which is defined on the original family $\left\{ E_1, \ldots, E_n \right\}$.

(6) Dominance: If x' and y' are conditional gambles defined on some sub-family, and if x'' and y'' are conditional gambles defined on its complement, and if, finally, $x' \gtrsim y'$ and $x'' \gtrsim y''$, then $\langle x', x'' \rangle \gtrsim \langle y', y'' \rangle$.

Not all of the above axioms are needed in order to develop a fruitful theory of choice under risk. Indeed, few (if any) authors subscribe to all six of them. The most common practice is to leave out the fifth axiom (convexity) and adopt all the others, possibly with some modifications. Recently, R. J. Aumann [3] investigated the theory which would result if the second axiom (completeness) were to be left out as well. At one point, a controversy developed concerning the soundness of the sixth axiom (dominance) with Allais [1], Manne [8] and Wold [13] leading the attack, and Markowitz [9], Savage [12] and Samuelson [11] leading the defense. The controversy subsided and, for the most part, the dominance axiom remained on the books, although attempts were made to weaken it as much as possible. (For the weakest version of all, see Herstein and Milnor [7].) Only in the study of general equilibrium in the presence of uncertainty (Arrow [2], Debreu [5]) is it common to leave out the dominance axiom. It is in the same area, however, that <u>convexity</u> must be assumed. In general, the absence of convexity spells trouble for the economic theorist. For this reason, a critical investigation of why convexity is usually left out may be in order.

The decline of convexity in the theory of choice under risk was brought about to a large extent by Friedman and Savage [6]. In this well known article, the authors set out to explain the co-existence of gambling and insurance, and they find themselves compelled to introduce a non-convexity in the decision maker's preferences in order to do so. Friedman and Savage chide their predecessors who were reluctant to abandon convexity, saying that this reluctance stemmed from a belief in diminishing marginal utility. They argue, in effect, that this belief in diminishing marginal utility is an anachronism, a remnant of Marshallian cardinal utility in the modern era of ordinal indifference curve analysis. This is a rather curious argument, for two reasons: (a) By adopting the dominance axiom Friedman and Savage find themselves in the very world which they have pronounced obsolete, i.e., in the world of cardinal utility. It is not Marshall's world, to be sure, but it is a world in which the sign of the second derivative of the utility function is invariant under the permissible transformations, so that diminishing marginal utility is a perfectly meaningful concept. (b) If one draws the indifference curves which correspond to the Friedman-Savage hypothesis (commodities being defined as above) one finds that in certain regions these indifference curves have the "wrong" curvature. This fact is, of course, a cause for concern in the ordinal theory just as much as it would be in the cardinal theory.

Friedman and Savage did not attempt to test convexity directly, using data from experiments or from other sources. In this essay we shall argue that such direct tests tend, by and large, to support the convexity axiom and, furthermore, that this convexity need not imply that a man cannot carry a lottery ticket and an insurance policy in the same pocket.

II. Acceptance Sets

The preference ordering $\gtrsim$ assumes a given wealth level. A gamble x may be superior to a gamble y at one wealth level, but inferior to it at another wealth level. Strictly speaking, then, one must add a label to the symbol $\gtrsim$ to indicate the level of wealth to which it refers. However, there is enough information in the preference ordering at one level of wealth to deduce the preference ordering at all other levels of wealth. Specifically, let $(x \gtrsim y)[\omega]$ be the statement " x is preferred or equivalent to y at the wealth level ω ." Furthermore, for each real number α , let $\bar{\alpha}$ be the n-tuple $\langle \alpha, \alpha, \ldots, \alpha \rangle$. The following assertion can be made:

$$(x \gtrsim y)\ [\omega + \alpha] \text{ if and only if } (x + \bar{\alpha} \gtrsim y + \bar{\alpha})[\omega] \ .$$

In other words, the preference ordering at any wealth level can be obtained from the preference ordering at any other wealth level by shifting the origin appropriately. Having said this, let us return to the simpler notation, writing $\gtrsim$ (without further labels) for preference or indifference, and agreeing to call the wealth level to which this relation corresponds the zero wealth level. In other words, let us agree to measure wealth in terms of deviations from the level which corresponds to $\gtrsim$.

The convexity axiom states that sets of the type $\{y | y \gtrsim x\}$ are convex for all x . Consider any gamble $x = \langle x_1, \ldots, x_n \rangle$. It is easy to show, given the completeness axiom, that there is always a "sure" gamble of the type $\langle \omega, \omega, \ldots, \omega \rangle$ which is indifferent to x . Let $\bar{\omega}$ be the vector $\langle \omega, \omega, \ldots, \omega \rangle$. Then, by looking at the sets $\{y | y \gtrsim \bar{\omega}\}$ for all values of ω , one is in fact looking at all the sets to which the convexity axiom applies. It is clear, however, that except for a shift of the origin, the set $\{y | y \gtrsim \bar{\omega}\}$ represents the set of all gambles which the decision maker accepts at the wealth level ω . [2] Since acceptance and rejection of gambles are observable actions, a way is open to check the various axioms directly.

It is convenient (although not entirely accurate) to refer to the set $\{y | y \gtrsim \bar{\omega}\}$ as the acceptance set for wealth ω . We shall use the symbol A_ω to denote this set. All the axioms of the foregoing section can be stated in terms of acceptance sets of this type:

(1') Transitivity: If $\omega \leqq \sigma$, then $A_\omega \supset A_\sigma$.

(2') Completeness: For any gamble x , there exists some ω such that $x \epsilon A_\omega$.

(3') Continuity: A_ω is closed for all ω

---

[2] The set $\{y | y + \bar{\omega} \gtrsim \bar{\omega}\}$ which is obtained from $\{y | y \gtrsim \bar{\omega}\}$ by shifting the origin, is precisely the set of gambles which the decision maker accepts at the wealth level ω .

(4') <u>Monotonicity</u>: If $x \in A_\omega$ and $y \geqq x$, then $y \in A_\omega$.

(5') <u>Convexity</u>: A_ω is convex for all ω.

(6') <u>Dominance</u>: Let A'_ω be the set of all conditional gambles (defined on some sub-family of events) which are <u>accepted</u> at the wealth level ω. Let A''_ω be the set of similar conditional gambles, only defined on the complementary sub-family. The axiom asserts that if $x' \in A'_\omega$ and $x'' \in A''_\omega$, then $\langle x', x'' \rangle \in A_\omega$.

As has already been indicated, our main concern will be the fifth of these axioms.

Consider the hypothesis of Friedman and Savage. Specifically, suppose that the decision maker's behavior in deciding whether to accept or to reject a gamble $x = \langle x_1, \ldots, x_n \rangle$ can be described as if the following steps were taken by him: First, the decision maker assigns non-negative real numbers $p_1, \ldots, p_n$, one to each of events in the family $\left\{ E_1, \ldots, E_n \right\}$ on which the gamble x is defined. These numbers satisfy $\Sigma p_i = 1$. He then consults a real function u, known as the <u>utility function</u> for wealth, and having the shape of the function in Figure 1:

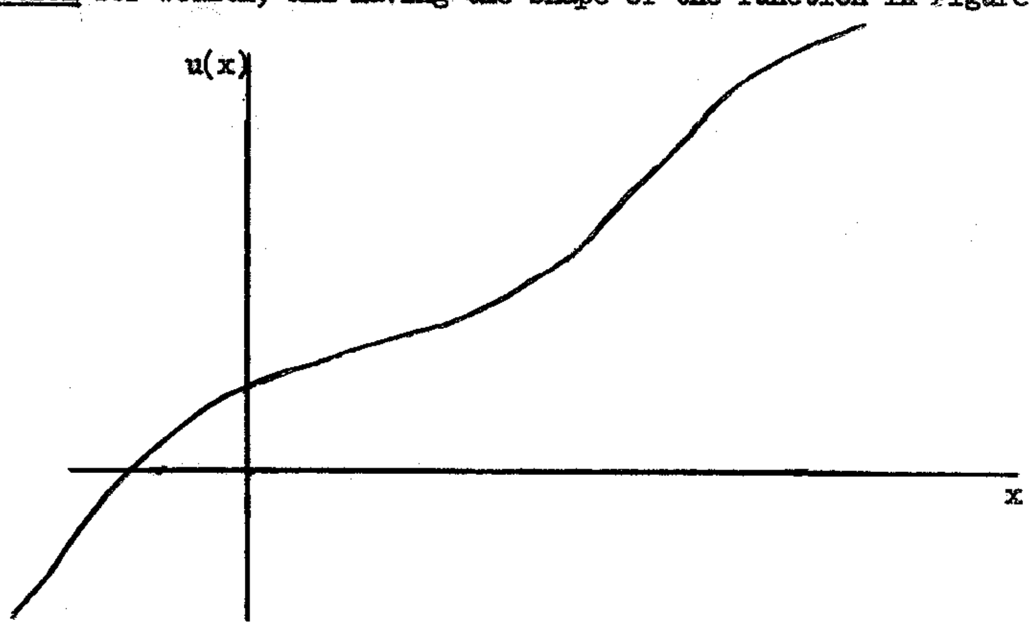

Figure 1

He computes a real number, which we might denote $U(x)$, and which is given by

$$U(x) = \sum_{i=1}^{n} p_i \ u\,(W + x_i)$$

where W is the decision maker's current wealth in dollars. Then, if $U(x) \geqq u(W)$ he accepts the gamble, and otherwise he rejects it.

Under this hypothesis, the set of all gambles which the decision maker accepts need not be convex. In particular, if the hypothesis leads to the acceptance of some gamble $< x_1, \ldots, x_n >$ for which the inequality

$$\sum_{i=1}^{n} p_i x_i < 0$$

holds, then the set of accepted gambles is non-convex. If $\gtrsim$ is the preference or indifference relation corresponding to the wealth level W , then this means that the set A_o is non-convex.

Since the numbers $p_1, \ldots, p_n$ are theoretical constructs (i.e. they are not observable) it is not possible to check convexity by constructing gambles for which $\Sigma\, p_i x_i < 0$ and then seeing if such gambles are accepted. Rather, one has to go back to the acceptance set itself and see if one can, indeed, locate any non-convexity.

III. Experimental Evidence

The experimental tests of convexity which are reported in this section pertain only to acceptance sets for gambles defined on families of events of the type $\{E, \sim E\}$, i.e., on families possessing but two members. Furthermore, there was no attempt to observe the subjects at different wealth levels, so acceptance sets are all of the type A_0 (i.e., A_ω for $\omega = 0$.) It is felt, however, that this restricted setting is adequate for the study of convexity, because it is not likely that non-convexities will arise in more complicated settings and yet fail to arise in the simple setting.

In this section, it will be convenient to use the symbol $A(E)$ to denote the set of all accepted gambles defined on a given dichotomy $\{E, \sim E\}$. By the monotonicity axiom, $A(E)$ is bounded from below by a monotonically decreasing function. Denoting this function f_E, one can write

$$f_E(x) = \min \{y \mid \langle x, y \rangle \in A(E)\}$$ [3]

for all real numbers x. (Clearly, if there is no y such that $\langle x, y \rangle \in A(E)$ then $f_E(x)$ is not defined.) We shall refer to f_E as the <u>offer function</u> (or <u>offer curve</u>) for the dichotomy $\{E, \sim E\}$. The statement that $A(E)$ is a convex set is clearly equivalent to the statement that f_E is a convex function.

Experimental construction of the offer curve f_E proceeds as

---

[3] By writing "min" rather than "inf" one is making use of the closedness of $A(E)$.

follows: An unambiguous random event E is described to a group of subjects. For example, the subjects may be told to let E be the event that so-and-so will receive the Republican Presidential nomination in 1964; or they may be told to let E be the event that a white ball will be drawn from an non containing 50 white balls and 50 red balls. The subjects are then asked to bet on the event E (and sometimes on the event $\sim E$.) More precisely, the subjects are told that they will <u>receive</u> (say) \$2.00 if E occurs and that they will have to <u>pay</u> x dollars if E fails to occur. They are then asked to quote how high x would have to be to make them refuse the bet. By varying the amount which the subject receives if E occurs, one obtains a series of points which may reasonably be regarded as points on (or near) the offer curve f_E .

The question "how high would x have to be for you to refuse to bet" is not a very good question for an experimental setting. A way must be found to get the subjects to quote the "highest" x which they would be willing to pay. This can be done in several ways, two of which are the following: First, the experiment may be conducted in the form of an <u>auction</u> in which the event E and the prize to be received if E occurs are announced. The subjects then submit closed bids (with no collusion) and the wager is awarded to the highest bidder. The second way consists of the following: Suppose that a subject had offered to pay y dollars if E fails to occur, in order to receive x dollars if E occurs. In other words, suppose that a subject had accepted the gamble $< x, - y >$. At a later stage of the game, offer him the gamble $< x, - 5y/4 >$

i.e., raise his bid by 25%. If he rejects this offer, then his original bid can be taken as "highest." A somewhat sloppy but convenient procedure is to let the subjects themselves raise their own bids by 25% and see if the resultant bid is "definitely too high." Most of the results which are reported below have been obtained using this latter procedure. Only in one case have results been obtained using the "auction" procedure.

All of these procedures require that the subject be presented with an entire sequence of gambles. But what the experiment is after is the behavior of the subject when presented with each gamble in isolation. One way to get around this problem is to tell the subjects in advance that at the end of the experiment one gamble (or very few gambles) will be selected at random, and only the gamble (or gambles) so selected will actually be contracted.

One point of the offer curve was obtained "free of charge" by assuming that the origin $< 0, 0 >$ lies on every offer curve, i.e., that $f_E(0) = 0$ for all E. This assumption can be broken in two: First, the assumption that gambles of the form $< x, y >$, where both x and y are negative, are never accepted. Second, the assumption that gambles $< x, y >$ for which both x and y are positive are always accepted. The first of these assumptions seems reasonable, and it actually represents a slight strengthening of the monotonicity axiom. The second assumption is in fact a statement of the requirement that the cost of entry (pecuniary or psychological) be zero for all gambles. In order to ensure this, subjects must in general be paid for participating in the experiment, or they must be told in advance that for the purposes of the experiment "no stakes are too small to bother with."

The actual results are divided into two parts, corresponding to two types of convexity tests. We shall refer to the first type as a test of <u>inter-quadrant</u> convexity, and to the second type as a test of <u>intra-quadrant</u> convexity.

<u>Inter-quadrant convexity</u>: Consider two gambles, $< x, y >$ and $< x', y' >$, both on the offer curve f_E. Suppose that $< x, y >$ is located in the southeast quadrant, while $< x', y' >$ is located in the northwest quadrant. Thus, clearly, $< x, y >$ represents a bet on the event E and $< x', y' >$ represents a bet on its complement, ~E. We shall say that the pair of gambles $< x, y >$ and $< x', y' >$ meets the test of interquadrant convexity if the line segment connecting $< x, y >$ and $< x', y' >$ does not pass below the origin (see Figure 2)

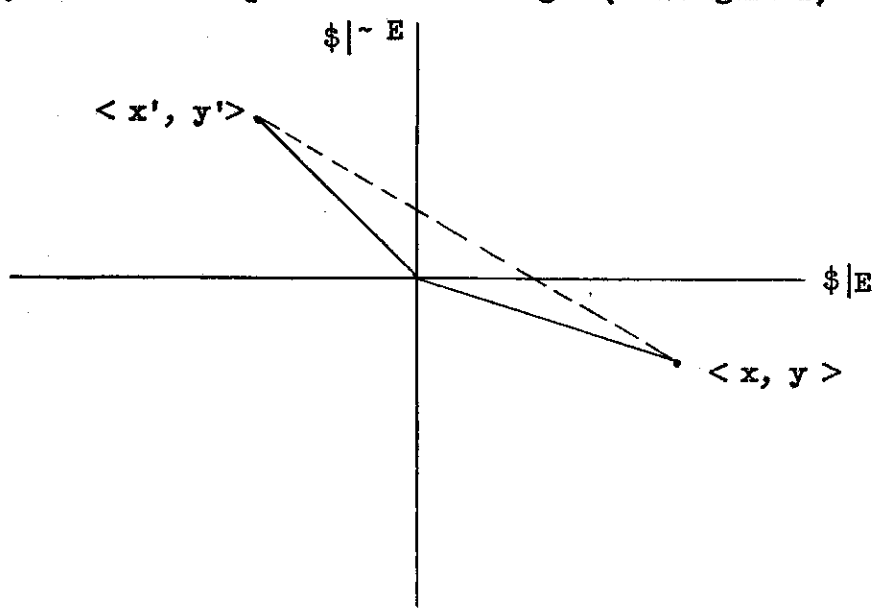

Figure 2

In symbols, inter-quadrant convexity holds if

$$\frac{xy' - x'y}{x - x'} \geqq 0$$

where $< x, y >$ and $< x', y' >$ are points in $A(E)$, satisfying $xx' \leqq 0$.

To check inter-quadrant convexity, 30 subjects (all of them Yale undergraduates) were asked to make offers for three pairs of bets. For each pair, an event E was defined (e.g. E = the event that the New York Yankees will enter and win the World Series in 1964) and the subjects were asked to offer separate bets on E and on $\sim E$, in the manner described above. The events chosen differed greatly in how likely they were generally considered to be. The sizes of the stakes also varied. The result was that out of 30 subjects, only one violated inter-quadrant convexity, and he did so only in one of three pairs of gambles. Specifically, he demanded to be paid $3.50 if E occurred in order to agree to pay $.50 if E failed to occur, but he demanded only $.07 if E failed to occur in order to agree to pay $.50 if E occurred. If he had demanded $.08 instead of $.07 in this latter gamble, he would have satisfied inter-quadrant convexity.

It was of specific interest to check inter-quadrant convexity for extremely likely or extremely unlikely events. To this end, a group of 15 subjects was presented with the following event: An urn contains 1000 marked balls, and one ball is drawn. Let E be the event that the ball marked "1" will be drawn. In bets on this event and it complement, inter-quadrant convexity was satisfied in all cases.

Intra-quadrant convexity: The second test concerns the convexity of the offer curve in each of the quadrants separately. Suppose that n points on the southeastern branch of the offer curve f_E are obtained. Let these points be denoted $< x_1, y_1 >$, $< x_2, y_2 >$ and so on, up to $< x_n, y_n >$. Assume that these points have been arranged in order of increasing abscissas, i.e., that $x_1 < x_2 < \ldots < x_n$. Finally, let $< x_o, y_o >$ be the origin. We shall say that intra-quadrant convexity holds if

$$\frac{y_i - y_{i-1}}{x_i - x_{i-1}} \leqq \frac{y_{i+1} - y_i}{x_{i+1} - x_i}$$

for $i = 1, 2, \ldots, n-1$.

This test was administered to the offer curves which resulted from responses of 50 subjects. Each subject made 15 responses, five for each of three events, from which three offer curves (each consisting of six points, including the origin) were constructed. Thus, there were 150 offer curves in all. Once again, the three events for which offer curves were constructed differed greatly in how likely they were generally considered to be.

The subjects' responses resulted in the following: In 42 cases, all three offer curves were convex. In 5 cases, two offer curves were convex and one was not. In 2 cases, one offer curves was convex and two were not. Finally, in one case all three offer curves were not convex.

Most of the non-convex offer curves were very nearly convex, but no attempt was made to work out a statistical test for the hypothesis of convexity. Most of the non-convex offer curves passed three (out of four) checks and failed in one.

Several experiments gave rise to an important observation, namely that there is no apparent connection between "betting at unfavorable odds" and non-convexity of the acceptance set. For example, an urn containing nine white balls and one black ball was presented to a group of 17 subjects. They were told that three balls were to be drawn from the urn, with replacement. Then they were asked to bet on the event that all three draws will yield a black ball. The amounts which they were to receive if they won the bet were fixed, and they were asked to submit bids of the amounts which they were willing to pay if they lost the bet. One subject produced the following series of bids. Let x be the amount which he would receive if he won the bet and let z be the amount which he offered to pay if he lost. Both x and z are in dollars:

x	z	x	z
3.00	.04	32.00	.24
5.00	.05	40.00	.28
7.00	.07	50.00	.32
10.00	.10	100.00	.50
17.00	.16	185.00	.60
24.00	.20	320.00	.66

For each of the bets, there was a chance that the gamble would actually be contracted and played. The gambles were imbedded randomly in a long series of other gambles, involving other events. We see that this subject exhibits "risk loving" behavior, i.e., he accepts gambles which, if we assume the probability of the event to be 10^{-3}, yield negative expected returns. Yet his offer curve is nice and convex. Other subjects were more cautious, and their bids were generally lower. However, seven out of seventeen made bids which, assuming a probability of 10^{-3}, yielded negative expected returns. Once again, there was no apparent connection between this behavior and violations of convexity.

To sum up, it seems that observation bears out the convexity postulate fairly well. There may be minor confort in this finding for the theorist, because to him a non-convexity is often a cause for concern.

IV. Subjective Probabilities

Let us proceed now to use acceptance sets in an operational definition of subjective probabilities. Several such definitions are possible, and the only distinction of the one which is about to be given is that it corresponds to the definition of subjective probability under the expected utility hypothesis.

Once again, consider gambles of the form $< x_1, \ldots, x_n >$ defined on a family $\{ E_1, \ldots, E_n \}$ of mutually exclusive and exhaustive events. Let A_o be the set of all gambles which are accepted if the cost of entry into the game is zero. The foregoing section was meant to lend credence to the hypothesis that A_o is a convex set. By "zero cost of entry" we mean that the origin $< 0, 0, \ldots, 0 >$ is on the boundary of A_o . Thus, there is a hyperplane which supports A_o at the origin. In other words, there exists a vector $p = < p_1, \ldots, p_n >$ such that

$$p \cdot x \geqq 0 \text{ for all } x \epsilon A_o .$$

By monotonicity, p must be a non-negative vector.[4/] Finally, since the length of p is arbitrary, it may be normalized to satisfy $\Sigma p_i = 1$. We may now refer to p_i as the subjective probability of E_i .

This definition identifies the subjective probability of E_i with the odds at which a person is willing to make small bets on E_i , which is entirely arbitrary. Indeed, it is easy to think of other definitions which do not seem less plausible. However, let us discuss the foregoing definition briefly.

The first thing to notice is that subjective probabilities, as defined above, need not be unique. In other words, there may be many different hyperplanes which support the set A_o at the origin. Indeed, in one of the experiments which led to this report, a subject consistently produced acceptance sets which looked like the set in Figure 3. There is a whole range of subjective probability distributions which would be applicable to this set under the foregoing definition.

[4/] To see this, suppose that $p_i < 0$ for some i . Let x be a vector in which the i-th component is positive and all other components are zero. By monotonicity, x is an accepted gamble. Yet, $p \cdot x < 0$, which is a contradiction.

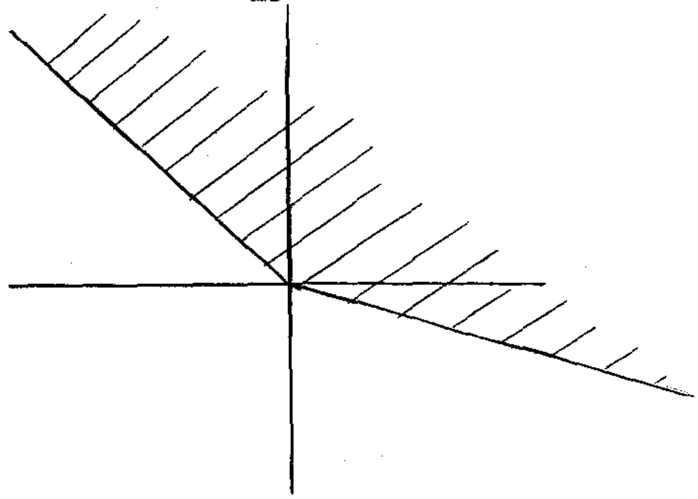

Figure 3

Secondly, the inner product $p \cdot x$ is, of course, the <u>expected return</u> of the gamble $< x_1, \ldots, x_n >$, given the subjective probability distribution $< p_1, \ldots, p_n >$. Thus what the statement

$$p \cdot x \geqq \text{ for all } x \epsilon A_o$$

says is that all accepted gambles are either fair or better than fair under this distribution. This is clearly not a behavioral statement. It does not say that if a gamble is not fair (in some "absolute" sense) then it will not be accepted. Rather, it says that probabilities have been defined in a manner which makes every accepted gamble fair or better than fair.

Consider the example of the "risk loving" subject in Section III. If we look at the offers which this subject made, we find that the subjective probability (according to the above definition) of the event "in each of three draws with replacement black will come up" is a number which is approximately equal to 1/75. At <u>this</u> probability the subject is indeed accepting gambles only if they yield a non-negative expected return.

In some cases (as in the case of this "risk loving" subject) the events on which gambles are defined are generated by a known probability mechanism. In such cases it is common to make certain assumptions on the behavior of the mechanism (e.g. that all balls in a given urn are equally likely to be drawn) and thus arrive at a set of "true" probabilities for the given events. If one compares the subjective probabilities which are obtained from the above definition with these "true" probabilities, one finds that some subjects tend to overstate low probabilities and to understate high probabilities. This is in line with the findings of Baratta and Preston [4] and with those of Mosteller and Nogee [10]. It should be emphasized, however, that this finding depends upon a specific (and arbitrary) definition of subjective probability, and it may not hold for other definitions.

References

[1] M. Allais, "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine." Econometrica 21; October 1953.

[2] K. J. Arrow, "Le Role des Valeurs Boursieves pour la Repartition la Mailleure des Risques." Econometrie, Centre National de la Recherche Scientifigue, Paris; 1953.

[3] R. J. Aumann, "Utility Theory without the Completeness Axiom." Econometrica 30; July 1962.

[4] P. Baratta and M. G. Preston, "An Experimental Study of the Auction Value of an Uncertain Outcome." American Journal of Psychology 61; April 1948.

[5] G. Debreu, Theory of Value. John Wiley & Sons; 1959.

[6] M. Friedman and L. J. Savage, "Utility Analysis of Choices Involving Risk." The Journal of Political Economy 56; 1948.

[7] I. N. Herstein and J. Milnor, "An Axiomatic Approach to Measurable Utility." Econometrica 21; April 1953.

[8] A. S. Manne, "The Strong Independence Assumption -- Gasoline Blends and Probability Mixtures." Econometrica 20; October 1952.

[9] H. M. Markowitz, Portfolio Selection. John Wiley & Sons; 1959.

[10] F. Mosteller and P. Nogee, "An Experimental Measurement of Utility." The Journal of Political Economy 59; October 1951.

[11] P. A. Samuelson, "Probability, Utility, and the Independence Axiom." Econometrica 20; October 1952.

[12] L. J. Savage, The Foundation of Statistics. John Wiley & Sons; 1954.

[13] H. Wold, "Ordinal Preferences or Cardinal Utility?" Econometrica 20; October 1952.

Footnotes

1. J. B. Rosen has presented a decomposition algorithm for nonlinear programming in [7].

2. For an interpretation of the present algorithm in the context of organizational decision making see Whinston [9] and for a general discussion of such applications see Whinston [8].

3. Without this assumption we would follow the argument as developed by Dantzig and Wolfe in [4] as it pertains to this point.

4. We assume that the constraint matrix has rank $\ell + 2$ and that the initial set of ρ_{ki} variables are associated with an independent set of columns. It should be observed that degeneracy can be handled exactly as in the case of linear programming.

5. These rules are based on [6].

6. This point will be discussed below in more detail.

7. For a discussion of these points see [6].

8. The slack variable y, which is required to be nonnegative, should be distinguished from an artificial variable, which must be zero during iterations of the algorithm. A slack variable is treated in the same fashion as any other primal variable.

9. See Dantzig [2] for a discussion of this topic.