



"Making book against oneself," the Independence Axiom, and Nonlinear Utility Theory

Citation

Green, Jerry. 1987. "Making book against oneself," the independence axiom, and nonlinear utility theory. *Quarterly Journal of Economics* 102(4): 785-796.

Published Version

<http://dx.doi.org/10.2307/1884281>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3203640>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

“MAKING BOOK AGAINST ONESELF,” THE INDEPENDENCE AXIOM, AND NONLINEAR UTILITY THEORY*

JERRY GREEN

An individual with known preferences over lotteries can be led to accept random wealth distributions different from his initial endowment by a sequential process in which some uncertainty is resolved and he is offered a new lottery in place of the remaining uncertainty. This paper examines the restrictions that can be placed on an individual's preferences by axioms that stipulate that such a process not be able to generate a new wealth distribution that is *prima facie* inferior to the original. The relationship of these axioms to the independence axiom of von Neumann and Morgenstern and to the quasi convexity of preferences in the wealth distribution are explored.

I. INTRODUCTION

Expected utility theory is the cornerstone of most of the modern economic analysis of behavior in the presence of risk. It is based on three assumptions about preference over distributions of payoffs. First, preferences are described by a complete preorder over the space of distributions. Second, this ordering is continuous in an appropriate sense. Third, the ordering satisfies an independence condition: if the individual is indifferent between distributions F and G , then he is indifferent between $\alpha F + (1 - \alpha)H$ and $\alpha G + (1 - \alpha)H$ for all probabilities α and all distributions H .

The independence condition is necessary for the representability of preferences by a linear functional over the space of all probability distributions. This linearity has been extremely useful as an analytical tool. The behavior of an individual with a linear preference functional over distributions has many attractive properties. His choices will be dynamically consistent when beliefs are updated using Bayes's Rule. Moreover, a knowledge of his preferences for small gambles, that is, for low probability deviations from his current situation, is sufficient to determine his preference globally.

Ever since it was first proposed as an axiom, the independence hypothesis has been under attack. Criticism has been based on theoretical and empirical grounds. Examples of *reasonable* behav-

*This research was supported by a grant from the National Science Foundation. The author wishes to thank Eddie Dekel and Andreu Mas-Colell for many discussions on this topic.

iors that are incompatible with independence have been given. The consequences of independence have been viewed by some as counterintuitive, if not actually self-contradictory. Furthermore, in experimental tests of ever improving design, the independence hypothesis has been repeatedly refuted. An important theoretical task, therefore, is the development of alternative hypotheses that place some testable constraints on behavior, less restrictive than those entailed by independence.

Recently, Machina [1982] has shown that the principal results of expected utility theory survive the relaxation of the independence axiom. Using only the first two axioms, he demonstrates that many of the observed violations of the independence axiom can be accommodated, while at the same time retaining the features of behavior most commonly assumed in economic analysis. For example, the monotonicity of preferences with respect to the relation of stochastic dominance of distributions has a straightforward interpretation in the absence of the independence axiom. Likewise, the idea of risk aversion and increasing absolute risk aversion, which are usually formulated in terms of the von Neumann-Morgenstern utility function, can equally well be recast in the more general framework. They can stand on their own in the absence of the independence hypothesis.

The goal of this paper is to develop qualitative restrictions on the nature of the nonlinear preference functional from new behavioral hypotheses. The origin of these hypotheses lies in an early but incorrect defense of expected utility theory, widely quoted but apparently never actually put forth in written form (see Machina [1983], p. 92 or Samuelson [1983], pp. 517-18). The statement is that, in the absence of the independence condition, an individual could be induced to "make book against himself."¹

A formal description of the hypotheses concerning the absence of any possibility that an individual can be induced to make book against himself will be given in the next section. Here, by way of introduction, we give a less than precise account. We hypothesize the existence of an outside agent who knows the preferences of the

1. There is a literature on the related hypothesis of "No Dutch Book," in the theory of subject probability and Bayesian updating rules. It says that, if a risk-neutral individual does not behave according to an additive probability distribution which he updates as a Bayesian, it would be possible to lead him into gambles with a negative expected payoff, or even into gambles which have nonpositive payoffs with probability one. See Freedman and Purves [1969] or deFinetti [1974], Chapter 3. Related work, in the nature of restricting preferences rather than updating rules, is the paper by Yaari [1985].

individual and who can do two things. He can formulate the risks facing the individual as a compound lottery, to be resolved in stages in a way that the individual can observe. And the agent can offer to exchange whatever risk the individual is facing, at an intermediate stage of the compound lottery, for something he prefers.

Thus, the outside agent leads the individual from the random variable upon which he initially has a claim to some other random variable. In the process the outsider must absorb the difference between them. If this difference were a positive random variable, outsiders would flourish. And agents with preferences such as this would lose money—or “make book against themselves.” Even if this difference were not positive, the distribution effectively received by the agent, taking into account the effects of the manipulation, might be stochastically dominated by that of the original random variable to which he had a claim. The outsider would receive a nonnegative amount of money, and if he could repeat this procedure with many agents, he would profit handsomely. Although the agent would have been led to accept the manipulation in accordance with his own preferences, he would retrospectively regret it. If preferences were learned, then preferences such as these would be discouraged, and presumably would eventually disappear.

The axioms that I propose take two forms. The first says that these manipulations should never lead the agent to an alternative that is (first-order) stochastically dominated by his initial prospect. The second form of this axiom says that the outsider should not be able to profit from this manipulation, either by obtaining a positive random payoff, or one with a positive mean.

The conclusions of this paper can be summarized as follows. If the individual has nonlinear preferences and if the realizations of the initial random variable can be observed by the outsider, there will always exist some initial random variable such that the outsider can extract a nonnegative amount of money from the agent with probability one and a strictly positive amount of money with some probability. In this sense the individual would have “made book against himself.” And it is in this sense that the defense of expected utility theory on these grounds is valid. The observability of the realization of initial random wealth is an important, unstated, aspect of the argument. If manipulative activity cannot be based on the realization of initial wealth, the outsider never can make money in this sense.

Stronger results are proved in the case of nonstochastic initial wealth. Linearity can be weakened to quasi convexity. Whenever

the preferences fail to be quasi-convex, there will exist some initial situation w such that the individual can be led from w to a random variable y whose distribution is stochastically dominated by w . Conversely, for any quasi-convex preference relation and any initial w , no manipulation exists that leads the agent to a stochastically dominated alternative. Note that a quasi-convex but nonlinear preference relation can be manipulated from *some* initial *random* wealth but not from *any nonrandom* wealth. Thus, such preferences cannot be dismissed on the grounds sketched above, unless there is reason to think that the particular initial random wealths from which detrimental manipulations can be made are actually present. In summary, this paper presents a strong argument for quasi convexity of preference and only a qualified bolstering of the argument for linearity.

The paper proceeds as follows. In Section II I present the preliminaries that will permit a discussion of lotteries whose outcomes are realized through a sequential procedure. Section III covers the concept of the manipulation of lotteries—the method by which individuals can effectively be led to accept payoff distributions that they would have rejected had they been offered as direct replacements for their initial wealth. Section IV presents the equivalence of independence for the nonmanipulability condition mentioned above, when correlated lotteries are allowed, and, for the case of nonstochastic initial wealth, the equivalence of the nonmanipulability condition with the quasi convexity of the preference functional. A concluding set of comments is given in Section V.

II. LOTTERIES

A *tree* $(N, <)$ is a finite set N and a partial ordering $<$ over the points of N , having two properties: $<$ has a least element, n_0 , called the *root* of the tree, and, for each $n \in N$, $<$ completely orders $\{n' \in N | n' < n\}$. The points of N are called *nodes*. If $n_1 < n_2$, then n_1 is said to be a *predecessor* of n_2 , and n_2 is said to be a *successor* of n_1 .

If $n_1 < n_2$ and there is no n_3 such that $n_1 < n_3 < n_2$, then n_1 and n_2 are said to be the *immediate predecessor* and *immediate successor* of each other. A node with no successors is said to be a *terminal node*. The set of all terminal nodes in a tree is denoted T . A *payoff function* $z: T \rightarrow R$ is any real valued function on the terminal nodes of the tree. Its values indicate the amount of money received by the player if that terminal node is reached.

Each pair (n_1, n_2) of two nodes consisting of an immediate predecessor and successor of each other is called a *step* in the tree. The set of all steps in a tree is denoted S . A *transition rule* $p: S \rightarrow [0,1]$ is a function specifying the conditional probability of taking each of the possible steps (n_1, n_2) having reached n_1 . Thus, for each $n_1 \notin T$, p satisfies $\sum p(n_1, n_2) = 1$, where the sum is taken over all immediate successors n_2 of n_1 . A transition rule induces a probability distribution over terminal nodes in the obvious way.

An *extensive lottery* $L = (N, <, z, p)$ consists of a tree, $(N, <)$, a payoff function z , and a transition rule p . To every extensive lottery L we can associate its payoff distribution F_L , induced by the distribution over terminal nodes and the payoffs associated with these nodes.

A *subtree* $(N', <')$ of $(N, <)$ is a subset N' of N consisting of a node n and all its successors, together with the restriction $<'$ of $<$ to N' .

A *sublottery* $L(n) = (N', <', z', p')$ of $(N, <, z, p)$ is the subtree $(N', <')$ with root n , the restriction z' of the payoff function z to the terminal nodes in $(N', <')$, and the restriction p' of the transition rule p to the steps in $(N', <')$. Special attention will be devoted to the sublotteries $L(n)$, where n is an immediate successor of n_0 , the root of $(N, <)$. These sublotteries will be termed *immediate sublotteries*.

III. PREFERENCES AND MANIPULATIONS

In this section we describe the methods through which an agent can be manipulated, dynamically, by an outsider who knows his preferences, his initial situation, and who has the ability to construct certain alternatives that can be offered to the agent. First, we shall discuss the manipulation of an agent whose situation is described by the claim to a lottery L . Then we shall treat the case in which the agent's initial wealth is described by a random variable x . The concept of manipulation will be different according to whether or not the outsider can or cannot use the realization of x to condition his offers to the agent.

Let $D(I)$ be the set of all probability distributions F with support on a compact interval I . The points of I represent possible payoffs in lotteries. Distributions concentrated on single points $w \in I$ are denoted G_w .

Let $V(\cdot)$ be a numerical representation of the agent's preference relation on $D(I)$. This representation exists under quite

general assumptions. We assume that $V(\cdot)$ is continuous and strictly increasing with respect to the partial ordering of first-order stochastic dominance. For every $F \in D(I)$, there is a point $c(F) \in I$ such that $V(G_{c(F)}) = V(F)$. The point $c(F)$ is termed the *certainty equivalent* of F .

Agents' preferences over lotteries are given by their preferences over the induced distributions. Letting $W(L) = V(F_L)$, $W(\cdot)$ serves as a numerical representation of this preference.

Similarly, if the agent's random wealth is described by a random variable x , the induced distribution of x will be denoted F_x and the numerical representation of preferences over random variables will be defined by $U(x) = V(F_x)$.

Consider a lottery L , and let the set of all *immediate successors* of its root, n_0 , be denoted S_L . Note that

$$F_L = \sum_{S_L} p(n_0, n) F_{L(n)}.$$

Suppose that for each $n \in S_L$, L_n is a lottery such that $W(L_n) \geq W(L(n))$. We could resolve the initial uncertainty in the lottery L , obtaining n with probability $p(n_0, n)$, and then offer to substitute L_n for the sublottery, $L(n)$, that would have remained at that point. Such replacements would all be accepted by an individual whose preferences are described by W . This individual would effectively have obtained a distribution,

$$F' = \sum_{S_L} p(n_0, n) F_{L_n},$$

instead of F_L .

When $V(\cdot)$ is nonlinear, it can be the case that $V(F') < V(F_L)$, even though $V(F_{L_n}) \geq V(F_{L(n)})$ for all $n \in S_L$. This fact motivates the basic definitions of this paper, which allow us to make precise the method by which such reversals of preference can be induced.

A *simple* (or one-stage) *manipulation of a lottery* L is a function $M(L; \cdot)$ defined on the immediate successors of the root of L , and taking values in the space of all lotteries. The lottery $M(L; n_1)$ replaces the sublottery $L(n_1)$ whenever n_1 is reached. Thus, the simple manipulation $M(L; \cdot)$ changes the lottery L into a lottery in which the payoff distribution is $\sum_{S_L} p(n_0, n_1) F_{M(L; n_1)}$.

A *two-stage manipulation of* L consists of a one-stage manipulation $M_1(L; \cdot)$ and, for each $n_1 \in S_L$, a one-stage manipulation of $M_1(L; n_1)$, denoted $M_2(M_1(L; n_1), \cdot)$.

We can proceed by compounding manipulations in this fashion for as many stages as desired. A (*compound*) *manipulation C of L* consists of a collection of 1-stage manipulations:

$$\begin{aligned}
 & M_1(L; \cdot), \\
 & M_2(M_1(L; n_1); \cdot) \text{ for each } n_1 \in S_L, \\
 & M_3(M_2(M_1(L; n_1); n_2); \cdot) \text{ for each } n_1 \in S_L, \text{ and } n_2 \in S_{M_1(L; n_1)},
 \end{aligned}$$

etc.

A simple manipulation of L , $M(L; \cdot)$ is *acceptable* to an individual with utility $W(\cdot)$ if for each $n \in S_L$,

$$W(M(L; n)) \geq W(L(n)).$$

A compound manipulation C is *acceptable* if each of the simple manipulations of which it is composed is acceptable.

Manipulations of lotteries alter the distribution of payoffs that the individual is actually receiving, when viewed from the root of the original lottery. We denote the distribution induced by a manipulation C by F_C .

An initial lottery F can be manipulated into any lottery F_C where C is an acceptable manipulation of F . The set of all such F_C includes, of course, all F' for which $V(G) \geq V(F)$. Whenever V is nonlinear, it will be much larger, and will not even define a transitive ordering over distributions. Our hypothesis, however, will be that the set of all such F_C should not include any distribution that F stochastically dominates.

We now consider an individual whose initial situation is described by a real-valued random variable x . A *manipulation of x* will be defined to consist of a substitution of a lottery L instead of x and a further manipulation of the lottery L . Let us denote a manipulation of x by the pair (L, C) , where C is a compound manipulation of L . A manipulation (L, C) is *acceptable* if $W(L) \geq U(x)$ and if C is an acceptable manipulation of L .

The outcome of a manipulation of (L, C) is a random payoff y whose distribution will be written as $F_{(L, C)}$. This distribution determines the effective outcome attained by the agent. That is to say, the agent's welfare depends only upon $F_{(L, C)}$. On the other hand, the effect of having undertaken this manipulation on the net wealth of the outside agent who has made these manipulative offers and exchanges may depend on *joint distribution* of the random payoff finally received by the agent, y , and on the agent's original payoff, x , which he has given up in exchange for L . Whether or not it is in the

interest of the outsider to undertake the manipulation (L, C) of x will depend on both his preferences and on this joint distribution.

The lottery L can be correlated with x in the sense that the outcomes of x are associated with terminal nodes t of L according to a function $x(t)$ which induces the random variable x . The interpretation of this is that the outside manipulator offers to pay $z(t)$ at terminal node t , but retains the right to the outcome $x(t)$ that the agent forgoes. Thus, the outside manipulator receives the random variable $x(t) - z(t)$ if the initial manipulation of x is accepted. Superimposing the compound manipulation C on L , the outside manipulator's payoff is altered further. Because of his ability to observe the outcome of x , the manipulator can receive any random variable $x - y$, where y is the outcome of the compound manipulation and x is the initial wealth, with the correlation between x and y controlled by the manipulator. If the outcome of x were not observable, the manipulator could only achieve those y for which x and z were independent.

IV. RESULTS

This paper attempts to limit the family of preferences that the agent might display by imposing, as a behavioral hypothesis, the conditions that he cannot be adversely manipulated or that the outside manipulator cannot benefit by doing so. The agent will surely be hurt by a manipulation if he can be led to accept a random variable whose distribution is stochastically dominated by that of his initial wealth. For any particular functional, some manipulations may actually make him worse off, even though they lead to distributions that are not dominated. But such transformations will not be beneficial to the manipulative outsider. Whether or not a given transformation is beneficial to the outsider depends upon whether we require that he be made better off with probability one, or merely that he reach a stochastically dominating distribution. For if the initial x stochastically dominates the result y , the outsider's payoff $x - y$ will have positive variance as well as positive mean. Therefore, to insure that the outsider profits at the expense of the agent, we might require that $x - y$ be a nonnegative random variable with positive mean. The weaker concept of stochastic dominance is relevant if the outsider is risk neutral, or if he can independently perform the indicated manipulation on many agents.

These remarks motivate the following results.²

2. A theorem very much like Theorem 2 below is stated in Yaari [1985].

THEOREM 1. If V satisfies the independence axiom, then for all x and all manipulations (L,C) , the distribution $F_{(L,C)}$ does not stochastically dominate F_x . A fortiori, $x - y$ cannot be a nonnegative random variable with positive mean.

Proof. Under the expected utility hypothesis the individual will be dynamically consistent. Therefore, $U(y) \geq U(x)$. The result is thus an obvious consequence of the monotonicity of V with respect to first-order stochastic dominance. ||

THEOREM 2. If V fails to satisfy the independence axiom, and if the outsider can observe the outcome of x , then there exists an initial x and a manipulation (L,C) that induces a random variable y such that $x - y$ is nonnegative and has a positive mean. A fortiori, F_x is stochastically dominated by $F_{(L,C)}$.

Proof. The proof is a direct consequence of the independence axiom and the fact that the steps in L can be perfectly correlated with the outcomes of x . When independence fails, we can find distributions F, G , and H and a probability α such that $V(F) = V(G)$ and $V(\alpha F + (1 - \alpha)H) < V(\alpha G + (1 - \alpha)H)$. Let x be any random variable on an atomless measure space Ω having the distribution $\alpha F + (1 - \alpha)H$. Without loss of generality, using a continuity argument, we can suppose that x takes only finitely many values. Let Ω_F and Ω_H form a partition of Ω having probability α and $1 - \alpha$ and such that the conditional distributions of x given Ω_F and Ω_H are F and H , respectively. The existence of such a partition follows from the fact that Ω is atomless.

Define the lottery L as follows. There are two immediate successors of the root of L , denoted n_F and n_H . The steps to n_F and n_H are taken in the events Ω_F and Ω_H , respectively. Following n_F and n_H , there are as many immediate successors as there are possible values of x . The transition rule of L specifies that these steps are taken in precisely the events on which x , conditional on Ω_F or Ω_H , takes the corresponding values. These nodes are all terminal, and the payoff function specifies the corresponding value of x at each node. Clearly L is just a two-stage way of realizing the same random variable x , and thus $V(L) = U(x)$.

Now manipulate L as follows: at n_F , substitute a lottery L_G instead of the subplottery $L(n_F)$. The lottery L_G consists of a root and a set of immediate successors of this root, each identified with a possible value of the distribution G , which, without loss of generality, is presumed to have finite support. The payoff and transition rule in L_G are such that G is the distribution of L_G . At n_H , no change is made. Let us call the resulting lottery L^1 .

Viewed from the root of L , this one-stage manipulation is acceptable because $V(\alpha F + (1 - \alpha)H) = W(L) < W(L^1) = V(\alpha G + (1 - \alpha)H)$. Indeed, a positive amount can be subtracted at every realization of L^1 without affecting this acceptability.

Finally, we manipulate L^1 by substituting, at n_F , a lottery precisely equal to the subplottery $L(n_F)$, which was taken away at the previous stage. Because $V(F) = V(G)$, this is acceptable. And if a positive amount were subtracted from the payoffs in L_G , then it is strictly preferred. ||

This two-stage manipulation of L has led to a lottery that is strictly dominated, for every $\omega \in \Omega$, by x . It is the ability of the outsider to construct lotteries whose outcomes are determined by events in Ω that is responsible for this pointwise dominance of the resulting random variables.

We now come to the main results of this paper. Here, the outside manipulator's inability to observe x is not important because the initial wealth of the consumer is assumed to be nonstochastic. In these circumstances, the linearity of $V(\cdot)$ required to insure nonmanipulability can be weakened to quasi convexity. Theorem 3 shows that quasi convexity is sufficient, and Theorem 4 shows that it is necessary.

THEOREM 3. Let $V(\cdot)$ be quasi-convex, and let C be an acceptable manipulation of G_w for some nonstochastic $w \in I$. Then G_w does not stochastically dominate F_C .

Proof. Let L be any lottery with a payoff function whose value is identically w . Let C consist of the simple manipulations, $M_1(L; \cdot), M_2(M_1(L; n_1); \cdot), \dots$. Since C is acceptable and since preferences are quasi convex, at least one of the subplotteries $M_1(L, n_1)(n_2)$ must be preferred to L . Similarly, at least one of $M_2(M_1(L; n_1); (n_2)(n_3))$ must be preferred to $M_1(L, n_1)(n_2)$ and hence better than L . As this process terminates, one of the subplotteries at the last stage $M_k(M_{k-1}(\dots (M_1(L; n_1); (n_2); \dots))(m_k))$ must be preferred to L . Because V is monotone in stochastic dominance, this subplottery must entail a positive probability of obtaining an outcome greater than w . As this subplottery is realized with positive probability, F_C , which is a mixture of it with other distributions, cannot be dominated by G_w . ||

THEOREM 4. Let $V(\cdot)$ fail to be quasi-convex. Then there exists $w \in I$ and an acceptable manipulation C of G_w such that F_C is stochastically dominated by G_w .

Proof. The theorem will be proved by exhibiting a manipulation with the desired property.

If $V(\cdot)$ is not quasi-convex, then there exists F, F^1, F^2 , and $\alpha \in (0,1)$ such that $F = \alpha F^1 + (1 - \alpha)F^2$ and $V(F^i) < V(F)$ for $i = 1,2$. By the continuity of $V(\cdot)$ we can take F, F^1 , and F^2 to have finite support, without loss of generality.

Let $w = c(F)$. In the first step of the manipulation of G_w , we replace it by a lottery L with two immediate successors n_1 and n_2 of its initial node, and such that $F_L = F, F_{L(n_1)} = F^1$, and $F_{L(n_2)} = F^2$. This can obviously be done by a lottery having all the successors of n_1 and n_2 be the terminal nodes, and with the probabilities of n_1 and n_2 equal to α and $(1 - \alpha)$, respectively.

Let $M(L;n_i)$ be a lottery whose induced distribution is concentrated on the certainty equivalent $c(F^i)$ of the distribution F^i . By hypothesis, $M(L;\cdot)$ is an acceptable manipulation of L . Hence the compound manipulation of G_w is acceptable and results in a two-point distribution, with probabilities α and $1 - \alpha$ of $c(F^1)$ and $c(F^2)$, respectively. ||

V. CONCLUDING COMMENTS

This paper has demonstrated the connection between the linearity of the preference functional, its quasi concavity, and related notions of the inconsistency of choice. We have seen that an individual whose initial wealth is nonstochastic and whose preference functional is quasi-convex cannot be manipulated to accept a random variable whose value is surely below this wealth level. Conversely, if it is known that for every nonstochastic initial wealth level there is no acceptable dynamic manipulation leading to a dominated result, then the individual must have a quasi-convex preference functional.

When the possible initial situations include stochastic as well as nonstochastic wealth, the corresponding results are stronger: any nonlinear preference can be manipulated from some initial stochastic wealth situation; and no linear preference can ever be manipulated. Moreover, if the realization of the initial stochastic wealth is an observable event upon which manipulations can be made contingent, then the result of a manipulation can lead to a reduction in payoff with probability one. If it is not observable, the result can be a distribution dominated in the sense of first-order stochastic dominance, rather than pointwise.

I believe that these results clarify the extent to which arguments about "making book against oneself" can be used to constrain the form of nonlinear utility functionals. The principal conclusion that preferences should be quasi-convex is another

independent argument for quasi convexity, bolstering those based on temporal risks (see Kreps and Porteus [1978] and Machina [1984] and those based on the ability of agents themselves to randomize their choices, and hence obtain any convex combination from among those offered to them.

HARVARD UNIVERSITY

REFERENCES

- deFinetti, B., *Theory of Probability*, 2 volumes (New York: Wiley, 1974).
- Freedman, D. A., and R. A. Purves, "Bayes' Method for Bookies," *Annals of Mathematical Statistics*, XL (1969), 117-88.
- Kreps, D., and E. Porteus, "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, XLVI (1978), 185-200.
- Machina, M., "Expected Utility Analysis without the Independence Axiom," *Econometrica*, L (1982), 277-323.
- , "The Economic Theory of Individual Behavior Toward Risk: Theory, Evidence and New Directions," IMSSS Technical Report No. 433, Stanford University, 1983.
- , "Temporal Risk and the Nature of Induced Preferences," *Journal of Economic Theory*, XXXIII (1984), 199-231.
- Samuelson, P. A., *Foundations of Economics Analysis*, enlarged edition (Cambridge, MA: Harvard University Press, 1983).
- Yaari, M., "On the Role of 'Dutch Books' in the Theory of Choice under Risk," Nancy L. Schwartz lecture, Northwestern University, April; addendum, Center for Research in Mathematical Economics and Game Theory, research memorandum no. 67, July 1985.

Copyright of Quarterly Journal of Economics is the property of MIT Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.