Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty

Mark J. Machina

Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty

By MARK J. MACHINA

University of California, San Diego

I. INTRODUCTION

IN THE LAST DECADE, the economic theory of choice under uncertainty has gone from one of the most settled branches of economics to one of the most unsettled. Although the debate encompasses several topics, it revolves around a single issue: the continued supremacy of the classical "expected utility" model of individual choice under uncertainty in light of a growing body of evidence that individuals do *not* maximize expected utility and the development of a number of alternative "non-expected utility" models of individual decision making.[1]

[1] For surveys of these empirical findings and theoretical developments, see Kenneth MacCrimmon and Stig Larsson (1979), Machina (1983a, 1983b, 1987), Robert Sugden (1986), Martin Weber and Colin Camerer (1987), and Peter Fishburn (1988).

Before these alternative models will be adopted by economists, however, researchers in this area will have to accomplish three goals. The first, which can be termed the *empirical goal*, is to show that non-expected utility models *fit the data better* than the standard expected utility model. Because these models are typically generalizations (that is, weakenings) of the expected utility model, it is not enough that they simply be compatible with more observations. To be successful, they must weaken the expected utility hypothesis in a manner that captures the types of systematic violations of expected utility that *have* been cataloged, while retaining those of its empirical properties that *have not* (or at least not yet) been refuted.

The second objective can be termed

the *theoretical goal.* This is to show that non-expected utility models of individual decision making can be used *to conduct analyses of standard economic decisions under uncertainty,* such as insurance, gambling, investment, or search, in a manner that at least approximates the elegance and power of expected utility analysis. Unless and until economists are able to use these new models as engines of inquiry into basic economic questions, they—and the laboratory evidence that has inspired them—will remain on a shelf.

Researchers have come a long way toward attaining each of these two goals. The data on expected utility's key empirical property of "linearity in the probabilities" has been very uniform, exhibiting a systematic form of departure from this property that has been captured by several non-expected utility models and that continues to be observed in experimental investigations.[2] Although it has not been as extensive as the empirical work in the area, the theoretical application of non-expected utility models, both to standard economic questions as well as to theoretical issues that *cannot* be handled by the classical model, is also proceeding apace.[3]

However, there remains one more objective that must be attained prior to the general acceptance of non-expected utility models, which can be termed the *normative goal.* This is to counter the widely held belief that non-expected utility maximizers will behave in a *dynamically inconsistent manner* that is particularly

subject to systematic manipulation and exploitation. This last objective forms the topic of this paper.

Given attainment of the empirical goal and the theoretical goal, why would a *descriptive* economist ever worry about the normative goal? Descriptive psychologists, for example, would never reject an empirically well-grounded and theoretically fruitful theory of decision making (or perception, or belief formation, or memory, and so on) on the grounds that it was not "rational"—that, after all, is what distinguishes psychologists as behavioral scientists from, say, statisticians or philosophers.

Economists, on the other hand, do exhibit a considerable affinity for the property of "rationality," and we have been severely criticized for this in light of an onslaught of laboratory evidence. But there is a good reason for this attitude, which derives from the additional responsibility that economists as *social* scientists must bear. Whereas experimental psychologists can be satisfied as long as their models of individual behavior perform properly in the laboratory, *economists are responsible for the logical implications of their behavioral models when embedded in social settings.*

To take a related example, there is a lot of laboratory evidence indicating that individuals' preferences can be systematically intransitive in certain situations. Why haven't economists responded by simply dropping the assumption of transitivity from their standard model of the consumer? Because if you take such a naive model of intransitive preferences and put it in a cage with a classical economic agent, it will get eaten alive by a simple "money pump" argument. (See Section 3.2 for details of this argument, as well as references to some more sophisticated models of nontransitive preferences that are immune to it.) Unless and until economists observe such ex-

[2] Again, see the references in Footnote 1.

[3] For example, Beth Allen (1987), Soo Hong Chew (1983), Chew, Larry Epstein, and Itzhak Zilcha (1988), Vincent Crawford (1989), Eddie Dekel (1986, 1989), Epstein and Stanley Zin (1989), Fishburn (1984a), Jerry Green and Bruno Jullien (1988), Edi Karni (1987), Karni and Zvi Safra (1987), David Kreps and Evan Porteus (1979), Machina (1982a, 1982b, 1984), Uzi Segal (1987), and Larry Selden (1978, 1979).

plicit money pumping in the real world, they won't adopt models that imply it must exist.

In other words, economists will not, and *should* not, employ behavioral models that imply economically self-destructive behavior in the presence of other (greedy) economic agents. This is not to say that we cannot allow the individual to exhibit such tendencies to some extent, or some of the time. Models in which the form of advertising or packaging has an effect on consumers' utility of a product constitute perfectly good economic models, and we can be sure that they are being fruitfully used by real world marketing and packaging departments. But any model that is so unsophisticated to imply that the agent can be invariably and repeatedly bilked out of cash will be rejected by any positive economist aware of the fact that such continual bilking simply does not take place.

Herein lies the importance of the normative goal to the general acceptance of non-expected utility models. There is a widespread belief that, just as intransitive preferences allow you to be "money pumped," non-expected utility preferences make you susceptible to a similar form of ruinous exploitation, namely, that someone can get you willingly to "make book against yourself." If this were true, economists would be right to reject these models regardless of their laboratory performance or theoretical properties. The objective of this paper is to dispel this impression, and to demonstrate that non-expected utility models are capable of generating behavior that is dynamically consistent and is not "manipulable" in any sense that is not also exhibited by preferences under certainty.

The following section develops the framework for these arguments by discussing the key distinction between expected utility and non-expected utility preferences, namely the property of *separability across mutually exclusive events*. Section 3 reviews the two classes of arguments purporting to be able to trick non-expected utility maximizers into "making book against themselves." The first class of arguments—those involving "static choice"—can be dismissed quite easily, while the second class—involving "dynamic choice"—are seemingly more formidable. Section 4 offers a critique of this latter group of arguments. In particular, I will show that each of them relies upon a hidden assumption concerning how decision makers behave in dynamic choice situations, namely *consequentialism* in the sense of Peter Hammond (1988a, 1988b). However, the property of consequentialism, though automatically satisfied by expected utility maximizers, is essentially a dynamic version of the very separability that non-expected utility maximizers reject, and is accordingly *inappropriate* to impose on such agents. In Section 5, I show that when the assumption of consequentialism is dropped and non-expected utility preferences are extended to dynamic choice settings in the same manner that economists would extend nonseparable preferences across time, commodities, or any other economic dimension, non-expected utility maximizers will be dynamically consistent, and are not "manipulable" in any sense not shared by nonseparable preferences over commodities. Section 6 provides a discussion of several aspects of the process of modeling nonseparable (i.e., non-expected utility) preferences under uncertainty.[4]

[4] Although the treatment in this paper will be self-contained, readers wishing a historical overview of the debate on "rational" choice under uncertainty should refer to Jacob Marschak (1950), Paul Samuelson (1952a), Herman Wold (1952), Maurice Allais (1953a, 1953b), Daniel Ellsberg (1954, 1961), R. Duncan Luce and Howard Raiffa (1957), Raiffa (1961, 1968), Jacques Drèze (1974), Paul Weller (1978), Edward McClennen (1983, 1988a, 1988b, 1989), Graham Loomes and Sugden (1984a, 1986, 1987), David Bell (1985), Thore Johnsen and John Donaldson (1985), Amartya Sen (1985), Sugden (1985), Mena-

2. EXPECTED UTILITY VERSUS NON-EXPECTED UTILITY PREFERENCES

2.1 LOTTERIES, PREFERENCES, AND PREFERENCE FUNCTIONS

An individual making a one-shot or "static" decision under uncertainty can be viewed as having to choose out of a set of alternative risky prospects or *lotteries*. Algebraically, we can represent a lottery (or more formally, a *single-stage* or *simple lottery*) by the notation $\tilde{X} = (x_1,p_1; \ldots ; x_n,p_n)$, where p_i denotes the probability of obtaining the outcome x_i. Depending upon the setting, the outcomes $(x_1, \ldots, x_n)$ could represent alternative final wealth levels, alternative changes from the individual's current wealth level, or alternative nonmonetary outcomes. We will adopt the convention that the probabilities $(p_1, \ldots, p_n)$ in any lottery $\tilde{X} = (x_1,p_1; \ldots ; x_n,p_n)$ are all positive and sum to unity, so if "zero final wealth" or "zero change in wealth" is a possible outcome, then both it and its probability should be explicitly represented. We do not require the outcomes $(x_1, \ldots, x_n)$ to be distinct, since different events could lead to the same monetary or nonmonetary outcome, although we will identify $(\ldots ; x, p; x, q; \ldots)$ and $(\ldots ; x, p + q; \ldots)$ as the same lottery. Graphically, we can represent such lotteries as in Figure 1, where the circle is known as a *chance node*.
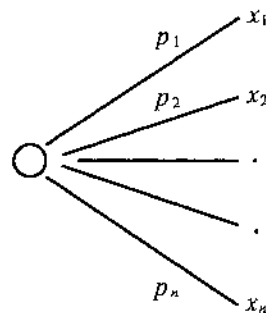

Figure 1. Graphical Representation of a Single-Stage Lottery

As in standard consumer theory, we assume that the individual has a preference ordering over this set of lotteries, so that if $\tilde{X} = (x_1,p_1; \ldots ; x_n, p_n)$ and $\tilde{Y} = (y_1, q_1; \ldots ; y_m, q_m)$ are two lotteries, we have either

$\tilde{X} \sim \tilde{Y}$ ($\tilde{X}$ is indifferent to $\tilde{Y}$)
$\tilde{X} \succ \tilde{Y}$ ($\tilde{X}$ is strictly preferred to $\tilde{Y}$), or
$\tilde{X} \prec \tilde{Y}$ ($\tilde{X}$ is strictly *less* preferred than $\tilde{Y}$).

Provided it satisfies the appropriate notion of continuity (e.g., Jean-Michel Grandmont 1972), this preference ordering can be represented by a *preference function* $V(\cdot)$, in the sense that

$\tilde{X} \sim \tilde{Y}$ if and only if $V(\tilde{X}) = V(\tilde{Y})$,
$\tilde{X} \succ \tilde{Y}$ if and only if $V(\tilde{X}) > V(\tilde{Y})$, and
$\tilde{X} \prec \tilde{Y}$ if and only if $V(\tilde{X}) < V(\tilde{Y})$.

The left-hand lottery in Figure 2 represents a *two-stage* or *compound* lottery of the general form $(\tilde{X}_1, p_1; \ldots ; \tilde{X}_n, p_n)$, that is, a lottery whose "outcomes" are themselves lotteries (termed *sublotteries*). Although the successive chance nodes in a compound lottery are resolved sequentially rather than simultaneously, we assume that this process does not require an economically significant amount of time and that the individual has no other economic activities or decisions (including consumption/savings decisions) to undertake in the meantime, so that he has no reason to prefer single-stage over compound lotteries on grounds of

hem Yaari (1985), Isaac Levi (1986), Glenn Shafer (1986), Paul Anand (1987), Green (1987), Hammond (1988a, 1988b, 1989), Richard Jeffrey (1988), Karni and Safra (1988a), Teddy Seidenfeld (1988a), and Chew and Epstein (1989a, 1989b), as well as the exchanges in *Econometrica* (October 1952), *Quarterly Journal of Economics* (February, May, and November 1963), *Statistical Science* (November 1986), *Economics and Philosophy* (October 1988), and *Journal of Behavioral Decision Making* (July 1989). Prominent changes of opinion in this debate have included Jacob Marschak (1938, p. 320 vs. 1951), William Baumol (1951 vs. 1958), and Samuelson (1950, 1983, pp. 509–18, vs. 1952b, 1988).
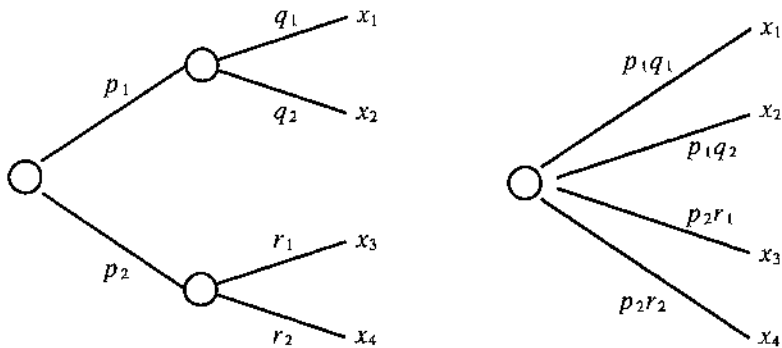
Figure 2. A Compound Lottery and Its Probabilistically Equivalent Single-Stage Lottery

impatience and/or planning benefits alone.[5]

Because the two-stage lottery in Figure 2 yields a $p_1 q_1$ probability of obtaining the outcome x_1, a $p_1 q_2$ probability of the outcome x_2, and so on, we say that it is *probabilistically equivalent* to the right-hand, single-stage lottery in the figure. Given the timing assumptions of the previous paragraph, we shall assume that the individual is always indifferent between any compound lottery and its probabilistically equivalent single-stage lottery, an assumption known as the *reduction of compound lotteries axiom*. By determining the probabilistically equivalent single-stage counterpart of each compound lottery, we can accordingly extend the individual's preference ranking and preference function from the set of all single-stage lotteries to the set of all compound lotteries.

Under these assumptions, the individual's behavior in a one-shot choice situation is fully determined: Given an opportunity set of simple and/or compound lotteries from which to choose, he will choose the lottery that is the most preferred, or equivalently, that yields the maximum value of the preference function $V(\cdot)$. "Nature" (or "chance," or "lady luck") then determines the outcome the individual will receive, according to the probabilities specified in the chosen lottery.

2.2 Expected Utility Preferences Over Lotteries

In regular consumer theory, we often make assumptions concerning the functional form of an individual's preference function over commodity bundles (e.g., Cobb-Douglas, CES, Leontief). In choice under uncertainty, the *expected utility hypothesis* is essentially that the preference function over lotteries takes, or can be monotonically transformed to take, the form:

$$V(\tilde{X}) = V(x_1, p_1; \ldots ; x_n, p_n)$$
$$= \sum_{i=1}^{n} U(x_i) p_i \qquad (1)$$
$$= U(x_1) p_1 + \cdots + U(x_n) p_n$$

where $U(\cdot)$ is termed the individual's *von Neumann-Morgenstern utility function*. Under the reduction of compound lotteries axiom, the expected utility of both the single-stage and the compound lottery of Figure 2 is $U(x_1) p_1 q_1 + U(x_2) p_1 q_2 + U(x_3) p_2 r_1 + U(x_4) p_2 r_2$. As al-

---

[5] For discussions of the applicability of expected utility theory when delays in the resolution of uncertainty *are* economically significant, see Jan Mossin (1969), Michael Spence and Richard Zeckhauser (1972), Drèze and Franco Modigliani (1972), Kreps and Porteus (1979), Michael Rossman and Selden (1979), Epstein (1980), and Machina (1984).
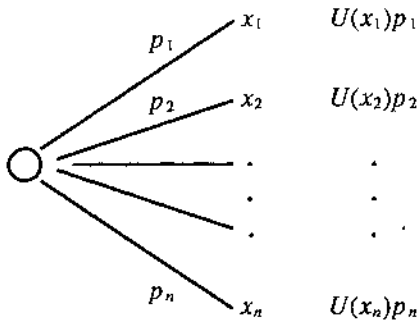
Figure 3. Contribution of Each Outcome/ Probability Pair to the Expected Utility of a Single-Stage Lottery
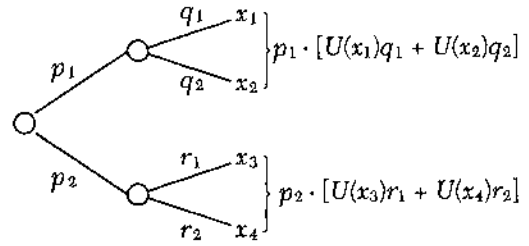

Figure 4. Contribution of Each Sublottery to the Expected Utility of a Compound Lottery

luded to above, economists have accumulated a considerable body of elegant and powerful theorems linking properties of expected utility maximizers' von Neumann-Morgenstern utility functions to their attitudes, and hence behavior, toward risk.

Separability Across Mutually Exclusive Events

The characteristic feature of the expected utility preference function (equation 1) is that it is "linear in the probabilities," which implies that expected utility preferences exhibit what can be termed *separability across mutually exclusive events*. This general attribute can be broken down into two specific properties: *replacement separability* and *mixture separability*. Replacement separability follows from the additive structure of the expected utility preference function $\sum_{i=1}^{n} U(x_i)p_i$ and the fact that the contribution of each outcome/probability pair (x_i, p_i) to this sum is independent of the other outcome/probability pairs, as illustrated in Figure 3. Thus, if an individual would prefer to replace the pair (x_1, p_1) by (y_1, p_1) in the figure, or in other words, if the lottery $(y_1, p_1; x_2, p_2; \ldots; x_n, p_n)$ were preferred

to $(x_1, p_1; x_2, p_2; \ldots; x_n, p_n)$, then he would also prefer to replace (x_1, p_1) by (y_1, p_1) in any *other* lottery of the form $(x_1, p_1; x_2^*, p_2^*; \ldots; x_m^*, p_m^*)$.[6]

The property of mixture separability follows from the fact that the contribution of each outcome/probability pair to expected utility can be interpreted as the utility of its outcome $U(x_i)$ times its probability p_i. Because the conditions $U(y_1)p_1 > U(x_1)p_1$ and $U(y_1) > U(x_1)$ are equivalent, an expected utility maximizer will prefer $(y_1, p_1; x_2, p_2; \ldots; x_n, p_n)$ over $(x_1, p_1; x_2, p_2; \ldots; x_n, p_n)$— that is, prefer a probability mixture of y_1 with $(x_2, \ldots, x_n)$ over the same probability mixture of x_1 with $(x_2, \ldots, x_n)$—if and only if he would prefer y_1 to x_1 in an *outright* choice over these two sure outcomes.[7]

These two separability properties also extend to mutually exclusive *sublotteries* in a compound lottery. From Figure 4 it is clear that the contribution of each sublottery to the expected utility of a compound lottery is independent of the other sublottery (or sublotteries) in the compound lottery. Thus, if an expected utility maximizer preferred replacing the

---

[6] It is important to note that this will be true even when one or more of the outcomes $(x_2, \ldots, x_n)$ or $(x_2^*, \ldots, x_m^*)$ take on the same value as x_1 or y_1.

[7] Once again, this will be true even when one or more of the outcomes $(x_2, \ldots, x_n)$ take on the same value as x_1 or y_1.

upper sublottery in this figure by some other sublottery, for example, if

were preferred to
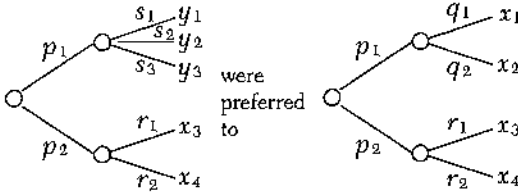
then he would prefer to make this replacement for any *other* configuration of the lower sublottery in the figure, which is replacement separability over sublotteries. It is also clear from Figure 4 that the contribution of each sublottery to expected utility can be interpreted as the expected utility of the sublottery itself (displayed in brackets) times its probability of occurrence (p_1 or p_2). Because the condition $[U(y_1)s_1 + U(y_2)s_2 + U(y_3)s_3] \cdot p_1 > [U(x_1)q_1 + U(x_2)q_2] \cdot p_1$ is again equivalent to $[U(y_1)s_1 + U(y_2)s_2 + U(y_3)s_3] > [U(x_1)q_1 + U(x_2)q_2]$, an expected utility maximizer will exhibit the previously displayed preferences over compound lotteries if and only if

were preferred to
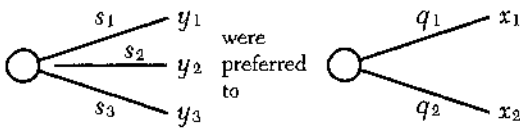
in an *outright* choice between these two single-stage lotteries, which is precisely mixture separability over sublotteries.

The properties of replacement and mixture separability across mutually exclusive outcomes and/or sublotteries can be combined and summarized in the following principle, known as the *independence axiom:*

The lottery $\bar{X}$ is preferred (indifferent) to $\bar{Y}$ if and only if $(\bar{X}, p; \bar{Z}, 1 - p)$ is preferred (indifferent) to $(\bar{Y}, p; \bar{Z},$ $1 - p)$ for all lotteries $\bar{Z}$ and all positive probabilities p.[8]

As several researchers have shown, the independence axiom, and hence the above pair of separability properties, is equivalent to the property that the individual preference function (provided it exists) takes the expected utility form

$$V(x_1, p_1; \dots; x_n, p_n) \equiv \sum_{i=1}^{n} U(x_i)p_i.$$ (See, for example, Marschak 1950; Samuelson 1952b; or Israel Herstein and John Milnor 1953, as well as the treatments in Luce and Raiffa 1957; Morris DeGroot 1970; Fishburn 1982a, and Kreps 1988).

2.3 Non-Expected Utility Preferences Over Lotteries

Observed Violations of Separability

When the outcomes x_1, y_1, and so on represent alternative wealth levels, replacing any individual outcome x_1 by a greater outcome y_1 always leads to a first-order stochastically dominating distribution (e.g., James Quirk and Rubin Saposnik 1962), and the properties of replacement separability and mixture separability *over individual outcomes* are nothing more than implications of the widely accepted property of first-order stochastic dominance preference over monetary lotteries. However, it is crucial to note that the expected utility properties of replacement and mixture separability *over sublotteries* are strictly *stronger* properties, and *do not* logically follow from first-order stochastic dominance preference.[9]

---

[8] This axiom implies mixture separability directly as stated. A double application yields $(\bar{X}, p; \bar{Z}, 1 - p) > (\sim) (\bar{Y}, p; \bar{Z}, 1 - p) \Leftrightarrow \bar{X} > (\sim) \bar{Y} \Leftrightarrow (\bar{X}, p; \bar{Z}^*, 1 - p) > (\sim) (\bar{Y}, p; \bar{Z}^*, 1 - p)$, which is replacement separability.

[9] To make an analogy with regular consumer theory, consider preferences over commodity bundles of the form (a, b, c, d). In this context, "dominance" (preference for vector-dominating bundles) implies

As mentioned in the Introduction, there is a growing body of empirical evidence demonstrating that individuals' preferences over lotteries in fact *do not* exhibit separability over sublotteries. One of the earliest examples of this is the well-known *Allais paradox*. This consists of the following pair of decision problems:

$$a_1: \{1.00 \text{ chance of } \$1,000,000$$

versus

$$a_2: \begin{cases} .10 \text{ chance of } \$5,000,000 \\ .89 \text{ chance of } \$1,000,000 \\ .01 \text{ chance of } \$0 \end{cases}$$

and

$$a_3: \begin{cases} .10 \text{ chance of } \$5,000,000 \\ .90 \text{ chance of } \$0 \end{cases}$$

versus

$$a_4: \begin{cases} .11 \text{ chance of } \$1,000,000 \\ .89 \text{ chance of } \$0. \end{cases}$$

Researchers such as Allais (1953a), Donald Morrison (1967), Raiffa (1968), Paul Slovic and Tversky (1974) and others have given this problem to hundreds and hundreds of subjects, and the modal if not majority choice in these studies is invariably for a_1 in the first pair and a_3 in the second pair.

There are several ways to see why the typical choices of a_1 and a_3 in this example violate the expected utility hypothesis. Algebraically, a preference for a_1 over a_2 implies

---

that if (a^*, b, c, d) is preferred to (a, b, c, d) then (a^*, b', c', d') will be preferred to (a, b', c', d') for any other $b', c',$ and d', which is replacement separability *over individual commodities*. Dominance does *not* imply that if (a^*, b^*, c, d) happens to be preferred to (a, b, c, d) then (a^*, b^*, c', d') will be preferred to (a, b, c', d') for any other c' and d', which would be replacement separability *over subbundles*. For more on the distinction between dominance and independence (i.e., separability) in choice under uncertainty, see Thomas Marschak (1986).

$$U(\$1M) > .10 \cdot U(\$5M) + .89 \cdot U(\$1M) + .10 \cdot U(\$0)$$

or equivalently

$$.11 \cdot U(\$1M) + .89 \cdot U(\$0) > .10 \cdot U(\$5M) + .90 \cdot U(\$0)$$

so that a_4 ought to be preferred to a_3. For purposes of the present discussion, however, it is more illuminating to invoke the reduction of compound lotteries axiom and rewrite these four prospects as:
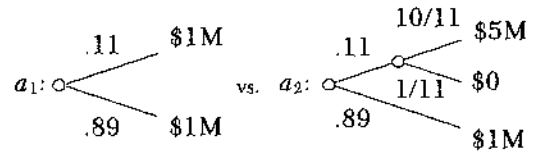
and
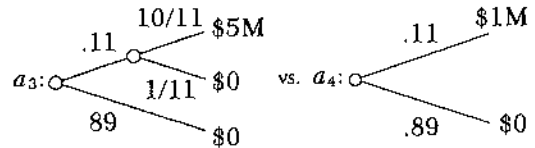
where $\$1M = \$1,000,000$. Viewed in this manner, the typical preferences in the Allais paradox are seen to violate replacement separability over sublotteries, because a choice of a_1 over a_2 indicates a preference for replacing the upper sublottery in a_2 by a sure $\$1M$ when the *lower* branch yields $\$1M$, but a choice of a_3 over a_4 indicates an unwillingness to make this same replacement when the lower branch yields $\$0$.

This Allais paradox is not an isolated example, but rather, a member of a whole class of similar violations of replacement separability. Such violations have been observed by MacCrimmon (1968), Herbert Moskowitz (1974), MacCrimmon and Larsson (1979), Daniel Kahneman and Tversky (1979), and Chew and William Waller (1986). In each

of these studies, the predominant form of departure from replacement separability corresponded to that of the Allais paradox, namely, a swing in preference from more risky to less risky sublotteries in one branch of a compound lottery as the sublottery in the *other* branch improves in the sense of first-order stochastic dominance, and this general pattern has come to be known as the "common consequence effect."

Another class of systematic violations, this time of mixture separability, involves pairs of the form:

$$b_1 : \begin{cases} p & \text{chance of } \$X \\ 1 - p & \text{chance of } \$0 \end{cases}$$

versus

$$b_2 : \begin{cases} q & \text{chance of } \$Y \\ 1 - q & \text{chance of } \$0 \end{cases}$$
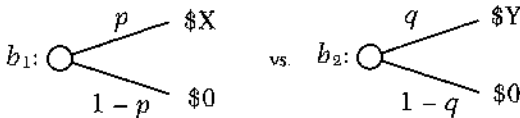
and

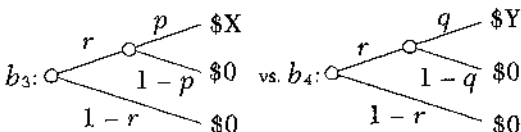$$b_3 : \begin{cases} r\,p & \text{chance of } \$X \\ 1 - r\,p & \text{chance of } \$0 \end{cases}$$

versus

$$b_4 : \begin{cases} r\,q & \text{chance of } \$Y \\ 1 - r\,q & \text{chance of } \$0 \end{cases}$$

where $p > q$, $0 < X < Y$ and $0 < r < 1$. Because we can invoke the reduction of compound lotteries axiom to write these four prospects as:

and

mixture separability over sublotteries implies choices of either b_1 in the first pair and b_3 in the second pair, or else b_2 in the first pair and b_4 in the second pair.[10] However, studies by Allais (1953a), Tversky (1975), Ole Hagen (1979), MacCrimmon and Larsson (1979), and Chew and Waller (1986) have found a systematic tendency for choices to depart from these predictions in the direction of preferring b_1 and b_4, a phenomenon known as the "common ratio effect." Thus Kahneman and Tversky (1979), for example, found that 80 percent of their experimental subjects preferred a sure gain of 3,000 Israeli pounds to a .80 chance of winning 4,000, but 65 percent preferred a .20 chance of winning 4,000 to a .25 chance of winning 3,000 ($p = 1$, $q = .8$, $X = 3,000$, $Y = 4,000$ and $\alpha = \frac{1}{4}$). In a study of laboratory rats choosing over gambles that involved substantial variation in their actual daily food intake, Raymond Battalio, John Kagel, and Don MacDonald (1985) also found this same pattern of violation. The reader is referred to the references in Footnote 1 for discussions of these and other systematic violations of the expected utility hypothesis of linearity in the probabilities, or in other words, of violations of separability across mutually exclusive events.

Non-Expected Utility Models of Preferences

Researchers have responded to this growing body of evidence by developing, analyzing, and testing *nonlinear* ("non-expected utility") functional forms for individual preference functions over lotteries. Some examples of non-expected utility preference functions, and researchers who have studied them, are listed in

---

[10] Algebraically, the choices b_1 and b_3 or else b_2 and b_4 are equivalent to the conditions $[U(X) - U(0)]p > [U(Y) - U(0)]q$ or $[U(X) - U(0)]p < [U(Y) - U(0)]q$ respectively.

TABLE 1

Examples of Non-Expected Utility Preference Functions

"Subjective Expected Utility"[11]/"Prospect Theory"

$$\sum_{i=1}^{n} v(x_i)\pi(p_i)$$

Ward Edwards (1955, 1962)
Kahneman and Tversky (1979)

"Subjectively Weighted Utility"

$$\left[\sum_{i=1}^{n} v(x_i)\pi(p_i)\right] \Big/ \left[\sum_{i=1}^{n} \pi(p_i)\right]$$

Uday Karmarkar (1978, 1979)

"Weighted Utility"

$$\left[\sum_{i=1}^{n} v(x_i)\tau(x_i)p_i\right] \Big/ \left[\sum_{i=1}^{n} \tau(x_i)p_i\right]$$

Chew and MacCrimmon (1979a, 1979b)
Chew (1983), Fishburn (1983)

"Anticipated Utility"

$$\sum_{i=1}^{n} v(x_i)\left[g\left(\sum_{j=1}^{i} p_j\right) - g\left(\sum_{j=1}^{i-1} p_j\right)\right]$$

John Quiggin (1982)

"General Quadratic"

$$\sum_{i=1}^{n} \sum_{j=1}^{n} T(x_i, x_j)p_i p_j$$

Chew, Epstein, and Segal (1988)

"Optimism/Pessimism"

$$\sum_{i=1}^{n} v(x_i)g(p_i; x_1, \ldots, x_n)$$

John Hey (1984)

"Ordinal Independence"

$$\sum_{i=1}^{n} h\left(x_i, \sum_{j=1}^{i} p_j\right)\left[g\left(\sum_{j=1}^{i} p_j\right) - g\left(\sum_{j=1}^{i-1} p_j\right)\right]$$

Segal (1984)
Green and Jullien (1988)

Table 1. Many (though not all) of these forms are flexible enough to exhibit the properties of stochastic dominance preference, risk aversion, and the types of observed violations of separability mentioned above, and many have proven to be highly useful both theoretically and empirically. Additional analyses of these and related forms can be found in Chew, Karni, and Safra (1987), Fishburn (1982b, 1984a, 1984b, 1988), Loomes and Sugden (1982, 1986, 1987), Ailsa Röell (1987), Ariel Rubinstein (1988), Segal (1987), and Yaari (1987). Again, the reader is referred to the refer-ences in Footnote 1 for general surveys of these models.

3. STATIC AND DYNAMIC ARGUMENTS AGAINST NON-EXPECTED UTILITY PREFERENCES

As mentioned in the Introduction, there are two classes of "making book" arguments that have been leveled against non-expected utility maximizers, namely those involving situations of "static choice" and those involving situations of "dynamic choice." Because these two types of arguments, and our responses to them, differ quite substantially, it is useful to pause and review the distinction between static and dynamic choice.

3.1 Static Versus Dynamic Choice Situations

As alluded to in Section 2.1, a deci-sion problem under uncertainty is said

[11] "Subjective expected utility," the name used for the preference function $\Sigma v(x_i)\pi(p_i)$ in the psychology literature of the 1950s and early 1960s, should be distinguished from the "subjective expected utility" approach of Leonard Savage (1954), which posited neither a weighting function $\pi(\cdot)$ nor objective proba-bilities, but instead derived the individual's "subjec-tive probabilities" over events from his preferences over state-payoff bundles (see Footnote 27).
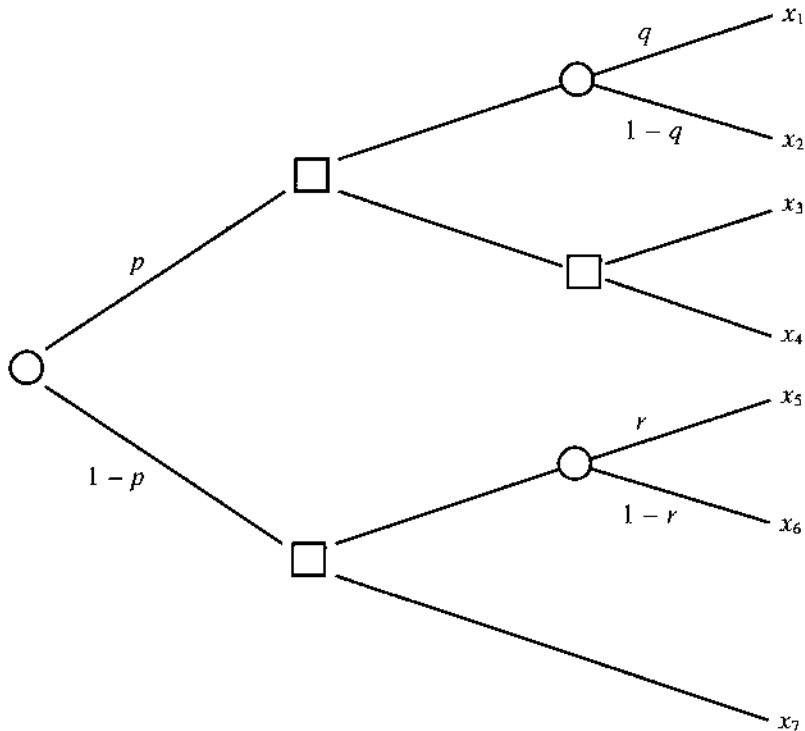
Figure 5. Example of a Decision Tree Representing a Dynamic Choice Problem

to involve *static choice* if the individual's final decision or decisions must be made (in the sense of *irrevocably* made) before any of the alternative lotteries (or stages of compound lotteries) are resolved. In other words, "nature" does not make any moves until the decision maker has irrevocably made all of his own moves.

On the other hand, a situation is said to involve *dynamic choice* if it involves decisions that are made *after* the resolution of some uncertainty. This could occur for a couple of reasons. One is simply that the individual may not have to (or even be able to) commit to a decision until after some uncertainty is resolved. Another reason might be that the available set of choices depends upon the outcome of the uncertainty. In any event, a dynamic choice situation will include at least some choices that the individ-

ual can (or must) postpone until after nature has made at least some of her "moves."

It is frequently convenient to represent the sequencing of choice and chance stages in a dynamic choice problem by the standard "decision tree" diagram (e.g., Harry Markowitz 1959, Ch. X; Raiffa 1968; Robert Schlaifer 1969). In the example of Figure 5, the individual begins at the left end or *root* of the tree, where he faces a round chance node, indicating that nature makes the first move by choosing either the upper or lower branch with the displayed probabilities. Each branch leads to a square *choice node* (or *decision node*), where the individual must make a decision that (in this example) will lead to either another chance node, another choice node, or directly to a final outcome. Although the

Figure 6. Decision Tree Representation of the Two Static Choices Problems of the Allais Paradox

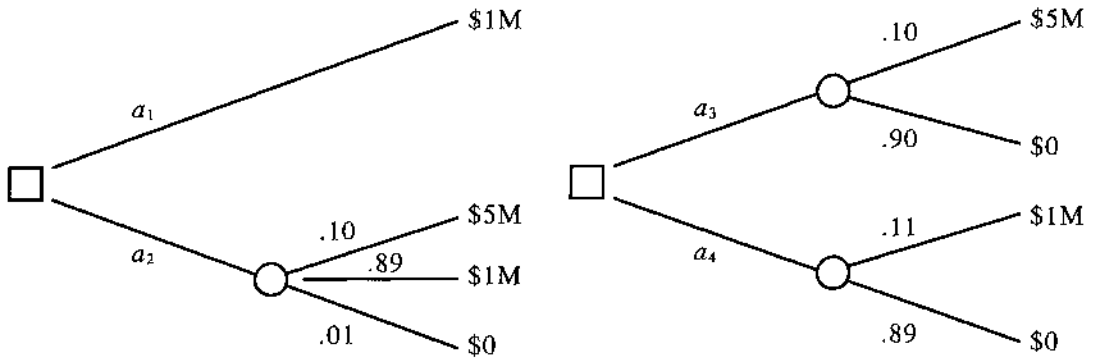individual is assumed to learn the outcome of each chance node before having to make any subsequent decisions, we once again assume that this process does not involve an economically relevant amount of time, and accordingly continue to assume the reduction of compound lotteries axiom, so that the individual has no preference for few-staged versus many-staged trees per se.

Of course, it is also possible to represent static decision problems with decision trees, as in Figure 6, which represents the two static choices problems offered in the Allais paradox. Given the above definitions, it is clear that a given decision tree represents a static choice problem if and only if no chance node is ever followed by a choice node, and represents a dynamic choice problem if and only if at least one chance node is followed by a choice node.

The relevance of the distinction between static and dynamic choice is that while static choice situations imply that all decisions will be irrevocably made before the resolution of any uncertainty, dynamic choice situations allow us to distinguish between an individual's *planned choices* for each decision node at the beginning of the decision problem (i.e., at the root of the tree) and his *actual choices* upon arriving at a given decision node.

It is this distinction that lies at the heart of the dynamic consistency issue.

3.2 STATIC ARGUMENTS AGAINST NON-EXPECTED UTILITY MAXIMIZERS

Before proceeding to the issue of dynamic consistency, however, it is useful to examine the classic *static* arguments against non-expected utility maximizers, if only to demonstrate how a properly designed non-expected utility model will be immune to each of them.

Intransitive Preferences Over Lotteries

The simplest of these arguments is against someone whose preferences over lotteries (or for that matter, any other objects) are intransitive. (See, for example, Donald Davidson, J. C. C. McKinsey, and Patrick Suppes 1955; Raiffa 1968, pp. 77–79; Yaari 1985; Frederic Schick 1986; or Anand 1987). Say the individual exhibits the following triple of strict preference rankings:

$$\tilde{Z} > \tilde{Y}, \quad \tilde{Y} > \tilde{X}, \quad \text{and} \quad \tilde{X} > \tilde{Z},$$

and say that he currently owns the (not yet resolved) lottery $\tilde{X}$. By continuity of preferences, there will exist some small positive ϵ such that $\tilde{X} - \epsilon > \tilde{Z}$, where $\tilde{X} - \epsilon$ denotes the lottery $\tilde{X}$ with the amount ϵ subtracted from each possible

payoff. To "make book" against such an individual, begin by offering him $\tilde{Y}$ in exchange for $\tilde{X}$. Given his preferences, he will accept it. Next, and before allowing nature to make her move (i.e., before allowing any of the lotteries to be resolved), offer $\tilde{Z}$ in exchange for $\tilde{Y}$ (once again, it will be accepted). Finally, offer $\tilde{X} - \epsilon$ in exchange for $\tilde{Z}$ (again, it will be accepted). Thus, an individual who started out owning $\tilde{X}$ has ended up owning $\tilde{X} - \epsilon$, or in other words, has been bilked out of the sure amount ϵ.

It is sometimes argued that one could continue this process and turn such an individual into a "money pump" which gushes out ϵ dollars at each push of the handle (each cycle), until he has delivered his entire net worth over to you. However, because the first cycle leaves the individual with $\tilde{X} - \epsilon$ rather than $\tilde{X}$, continuing this process requires that the individual's preferences contain an intransitive cycle involving $\tilde{X} - \epsilon$, another intransitive cycle involving $\tilde{X} - 2\epsilon$, and so on, which doesn't necessarily follow from the existence of the original cycle. However, this argument does claim to show that *to the extent* that preferences are intransitive, the individual can be exploited.

However, because each of the non-expected utility models of Table 1 represents preferences by a real-valued maximand $V(\cdot)$, and because it is impossible for three real numbers to satisfy $V(\tilde{Z}) > V(\tilde{Y}) > V(\tilde{X}) > V(\tilde{Z})$, these models will never exhibit intransitive cycles. Thus, they cannot be subjected to this form of "making book."[12]

Violations of First-Order Stochastic Dominance Preference

A second type of static "making book" argument *does* apply to some (though not all) of the non-expected utility preference functions in Table 1. Consider for example the "subjective expected utility" form

$$V(x_1, p_1; \ldots; x_n, p_n) \equiv \sum_{i=1}^{n} v(x_i)\pi(p_i).$$ If the function $\pi(\cdot)$ is not linear (i.e., if the form does not reduce to expected utility), there will exist probabilities $p_1, \ldots, p_m$, summing to unity, such that $\sum_{i=1}^{m} \pi(p_i) \neq \pi(1)$.

Say $\sum_{i=1}^{m} \pi(p_i) > \pi(1)$ (the reverse case follows similarly). In this case, there will exist outcome levels $x_1 < x_2 < \ldots < x_m < x^*$ such that

$$\sum_{i=1}^{m} v(x_i)\pi(p_i) > v(x^*)\pi(1).$$

Because this implies that the individual prefers the lottery $(x_1, p_1; \ldots; x_m, p_m)$ to a sure chance of x^*, we could accomplish in *single* trade that which in the previous argument took three trades, namely, bilking an individual holding x^* out of a strictly positive amount of money (at the very least $x^* - x_m$, and perhaps more).

Of course, the property of this preference function that allows us to do this is its violation of first-order stochastic dominance preference. Because the subjected expected utility form $\sum_{i=1}^{n} v(x_i)\pi(p_i)$ was for several years the most prominent example of a non-expected utility model in the literature (e.g., Edwards 1955, 1962; Luce and Suppes 1965; Tversky

---

[12] A final note on intransitivity: Even in the case when pairwise preferences *are* intransitive, Loomes and Sugden (1982, 1987), Fishburn (1984d; 1988, pp. 43–44), and Maya Bar-Hillel and Avshai Margalit (1988) have pointed out that the money pump argument implicitly assumes that the individual would ignore what his or her original holding had been when evaluating the second and third of these exchanges, and have argued that no one with intransitive preferences would want to ignore this information, and have developed "sophisticated" extensions of intransitive preferences to dynamic and/or nonpairwise choice situations that correspond closely to the approach presented in Section 5 of this paper and that are immune to the above money pump argument.

1967a, 1967b; Thomas Wallsten 1971),[13] this feature contributed to the general impression that *all* departures from linearity in the probabilities implied violations of first-order stochastic dominance preference. However, as noted in Section 2.3 (including Footnote 9), violations of the expected utility property of separability across mutually exclusive events and violations of first-order stochastic dominance preference are distinct concepts, and the more recent examples of non-expected utility models in Table 1 (the last five forms) will all exhibit first-order stochastic dominance preference provided (as with the expected utility model) that their component functions $v(\cdot)$, $\tau(\cdot)$, $g(\cdot)$ . . . satisfy the appropriate monotonicity conditions.[14] This property, along with the transitivity of these models, ensures that it is impossible to "make book" against them via any sequence of (well-specified) probability distributions in any static choice situation.

Incoherent Subjective Probabilities

A third type of static "making book" argument is against an individual whose underlying probabilistic beliefs or updating/calculating procedures are not *coherent*, that is, do not satisfy the regular laws of probability theory (that the proba-

bility of an event plus that of its complement sum to unity, that the probability of the union of disjoint events is the sum of their probabilities, Bayes' law for updating probabilities, etc.). First elucidated by Frank Ramsey (1926) and Bruno de Finetti (1937) and further developed by John Kemeny (1955), Abner Shimony (1955), R. Sherman Lehman (1955), David Freedman and Roger Purves (1969), and others, these arguments show that if the probabilities an agent assigns to some set of events (including their unions and intersections) are not coherent, and the agent is willing to accept betting odds on these events based on these "probabilities," then it will be possible to induce the agent to accept a set of bets that jointly imply that he cannot win money and that he has a positive probability of losing money. For the specifics of these arguments (which are beyond the scope of this paper), the reader is referred to the above sources as well as the discussions in Henry Kyburg and Howard Smokler (1980, pp. 3–22), Schick (1986), and Yaari (1985).

Do violations of linearity in the probabilities per se expose the non-expected utility preference functions of Table 1 to these types of arguments? No. These arguments involve the determination and manipulation of event probabilities and their use in the evaluation of functions from events to payoffs ("acts" or "bets"), whereas the preference functions in Table 1 are *defined* over well-specified, coherent probability distributions, such as the type presented in the Allais paradox and related examples. To make an analogy with regular consumer theory, it is clearly possible to "make book" against someone who cannot correctly add the amounts of commodities in different sized containers, or who cannot correctly multiply quantities by prices. However, if an individual with monotonic preferences *can* perform such operations cor-

---

[13] The treatment of first-order stochastic dominance preference in Kahneman and Tversky's "prospect theory," which consists of a combination of the subjective expected utility form $\Sigma v(x_i)\pi(p_i)$ with psychologically motivated "editing stages," is less straightforward. Their 1979 paper postulated that "dominated alternatives are detected and eliminated prior to the evaluation of prospects" in any given choice situation (Kahneman and Tversky 1979, pp. 275, 284). Their 1986 paper suggests that this elimination might be limited to cases where such dominance is "transparent," and that violations of dominance in nontransparent settings are consistent with the theory (Tversky and Kahneman 1986, Sect. IV).

[14] A necessary and sufficient condition for a general (differentiable) non-expected utility preference function to satisfy first-order stochastic dominance preference is given in Machina (1982a, Theorem 1).

rectly, it is not possible to make book against him just because these totals enter nonlinearly into their utility function. This is not to say that the issue of the formation and manipulation of subjective probabilities is unimportant (see, for example, Machina 1987, pp. 147–49, and the references cited therein), merely that it is a *different* issue from that of separable versus nonseparable (i.e., expected utility versus non-expected utility) preferences over well-defined probability distributions, which forms the topic of this paper.

3.3 DYNAMIC ARGUMENTS AGAINST NON-EXPECTED UTILITY MAXIMIZERS

We have seen that a properly designed non-expected utility model will be immune to each of the above static notions of "making book." However, the following class of arguments, involving situations of *dynamic choice,* seems to pose a more formidable challenge to non-expected utility models.

Argument for the Dynamic Inconsistency of Non-Expected Utility Maximizers

The simplest and probably most well known[15] of these arguments can be illustrated using the typical non-expected utility preferences expressed in the Allais paradox, namely a preference for lottery a_1 over a_2 and for a_3 over a_4. The first step in this argument consists of obtaining (and for good measure, writing down) the individual's preference ranking over the pair of options:

$\widetilde{W}$: 10/11 — \$5M   1/11 — \$0   versus   a sure \$1M
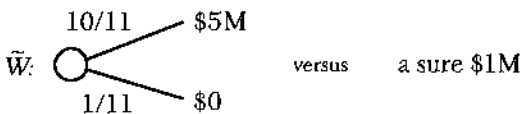
The second step consists of presenting the individual with the dynamic choice problems illustrated in Figure 7.

Consider the opportunity set of lotteries implied by the left-hand tree in this figure. A choice of the *upper* branch at the choice node, when combined with the probabilities at the initial chance node, would imply overall probabilities (.10, .89, .01) of receiving the outcomes (\$5M, \$1M, \$0), which is precisely lottery a_2 of the Allais paradox. A choice of the *lower* branch at the choice node would yield a sure chance of \$1M, which is, of course, prospect a_1.[16] Thus an individual with the typical Allais paradox preferences of a_1 over a_2 would, if he had to make this decision in advance, plan (or instruct his agent) to choose down at the choice node in the left-hand tree. A similar calculation demonstrates that the opportunity set for the right-hand tree consists of the Allais paradox prospects (a_3, a_4), so that an individual who preferred a_3 over a_4 would plan to choose *up* at the choice node in the right-hand tree.

At this point (goes the argument), an individual with Allais-type preferences will be in trouble. Say his preference in the first step had been for the lottery $\widetilde{W}$ over the sure \$1M. If nature were to choose up at the initial chance node in the left-hand tree, the individual would be facing *precisely* this first-step choice, and would accordingly reverse his original plan of choosing down (to obtain the \$1M) in favor of choosing up (to obtain $\widetilde{W}$). On the other hand, say his first-step choice had been for the sure \$1M over $\widetilde{W}$. In that case, he would undertake a similar *volte-face* should he arrive at the choice node in the right-hand tree. In other words, his behavior in one or the

---

[15] See, for example, Markowitz (1959, pp. 218–24), and Raiffa (1968, pp. 82–83) who refers to a similar argument by Robert Schlaifer.

[16] Because it generates the same opportunity set of lotteries, this tree is said to be *strategically equivalent* to the left-hand tree in Figure 6 (e.g., Irving LaValle and Fishburn 1987).
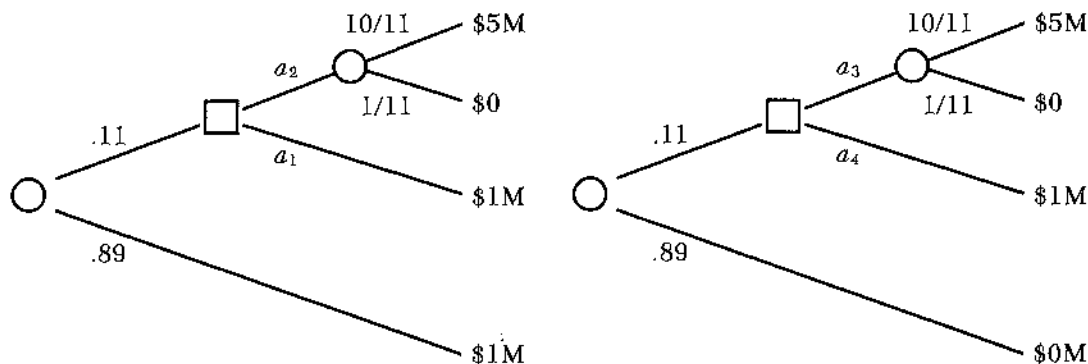
Figure 7. Dynamic Choice Problems that Generate the Same Opportunity Sets as the Allais Paradox Problems

other of these two trees will be *dynamically inconsistent,* in the sense that his actual choice upon arriving at the decision node would differ from his planned choice for that node.

It is important to note that this argument does not depend upon the specific lotteries and choices of the Allais paradox, but can be constructed out of *any* violation of replacement separability, mixture separability, or the independence axiom, or in other words, any departure from expected utility preferences.[17] Thus, unlike the static arguments of the previous section—which we have seen do not apply to "properly designed" non-expected utility models—this argument seems to demonstrate that non-expected utility maximizers are *generically* incapable of behaving consistently in even the simplest of planning situations.

Classical Argument for Making Book Against Non-Expected Utility Maximizers

Several researchers have shown how the above dynamic inconsistency argument can be adopted to "make book" against (that is, extract a sure payment from) a non-expected utility maximizer (e.g., Raiffa 1968, pp. 83–85; Yaari 1985; Shafer 1986; Green 1987; Kim Border 1987; Seidenfeld 1988a).[18] As noted in Footnote 17, any non-expected utility maximizer will exhibit the preferences

$$\hat{X} > \tilde{Y} \qquad \text{but}$$
$$[\tilde{Y}, p; \hat{Z}, (1 - p)] > [\hat{X}, p; \hat{Z}, (1 - p)]$$

for at least some lotteries $\hat{X}, \tilde{Y}, \hat{Z}$ and probability p. By continuity of preferences, there will exist some small positive ϵ such that

$$\hat{X} - \epsilon > \tilde{Y} \qquad \text{but}$$
$$[\tilde{Y}, p; \hat{Z} - \epsilon, (1 - p)] > [\hat{X}, p; \hat{Z}, (1 - p)]$$

where (as in Section 3.2) $\hat{X} - \epsilon, \hat{Z} - \epsilon,$ and so on denote these lotteries with the amount ϵ subtracted from each possible payoff.

Say that there is some event E with

---

[17] Any violation of replacement separability, mixture separability, or the independence axiom implies that the individual will exhibit preference rankings of the form $\hat{X} > \tilde{Y}$ but $(\tilde{Y}, p; \hat{Z}, 1 - p) > (\hat{X}, p; \hat{Z}, 1 - p)$ for some lotteries $\hat{X}, \tilde{Y}, \hat{Z}$ and probability p. Construct a decision tree whose initial chance node leads down with probability $1 - p$ to $\hat{Z},$ and up with probability p to a decision node offering a choice of $\hat{X}$ or $\tilde{Y}$. According to the argument, the individual would *plan* on choosing $\tilde{Y}$ at the choice node but would *in fact* choose $\hat{X}$ should he actually arrive there.

[18] We refer to this as the "classical" making book argument to distinguish it from the one discussed in Section 6.5.

probability p and that the individual currently owns the event-contingent prospect $(\tilde{X}$ if $E;\ \tilde{Z}$ if $\sim E)$, which implies that he owns a compound lottery of the form $[\tilde{X},\ p;\ \tilde{Z},\ (1-p)]$. In exchange for this initial holding, offer the individual the event-contingent prospect $(\tilde{Y}$ if $E;\ \tilde{Z} - \epsilon$ if $\sim E)$. Because this prospect implies the lottery $[\tilde{Y},\ p;\ \tilde{Z} - \epsilon,\ (1-p)]$, which is preferred to $[\tilde{X},\ p;\ \tilde{Z},\ (1-p)]$, the offer will be accepted. Now let the first stage of uncertainty be resolved. If the event $\sim E$ occurs, the terms of your exchange imply that you receive the lottery $\tilde{Z}$ and must pay the individual $\tilde{Z} - \epsilon$, so you have made a gain of ϵ. If the event E occurs, you receive the lottery $\tilde{X}$ and must pay the individual $\tilde{Y}$. However, from the above-displayed ranking it follows that the individual would be happy to accept a payment of $\tilde{X} - \epsilon$ instead of $\tilde{Y}$, so you again make a gain of ϵ. When the dust clears, the individual who initially held the prospect $(\tilde{X}$ if $E;\ \tilde{Z}$ if $\sim E)$ has come out of this process as if he held $(\tilde{X} - \epsilon$ if $E;\ \tilde{Z} - \epsilon$ if $\sim E)$, and you (who entered the story with no resources at all) have come out with $(\epsilon$ if $E;\ \epsilon$ if $\sim E)$. In other words, you have found a way to bilk a generic non-expected utility maximizer out of a sure ϵ.[19]

Argument That Non-Expected Utility Maximizers Could Be Averse to Information

It is almost a truism that advanced resolution of uncertainty in a dynamic choice problem, that is, prior knowledge of what nature's "moves" will be, could

[19] Note that this argument relies upon the individual holding the proper initial probability distribution. Green (1987) has shown that if a non-expected utility maximizer's preferences are at least *quasi-convex* in the probabilties (so that a randomization of two indifferent distributions is never strictly preferred), then this making book argument will not work on any individual with a *nonstochastic* initial wealth holding.

never make an individual worse off ex ante, and in general, will be strictly preferred. However, a final argument, put forth by Peter Wakker (1988) and Ronald Hilton (1989) (see also Kevin Keasey 1984; Loomes and Sugden 1984b), purports to demonstrate that "sophisticated" non-expected utility maximizers can actually be made *worse off* by receipt of this type of information in dynamic choice settings. Although we shall illustrate this argument with the specific lotteries and preferences in the Allais paradox, it is important to note that, once again, such an example can be constructed out of any departure from expected utility preferences.

Say that the individual has the typical Allais preferences of a_1 over a_2 and a_3 over a_4, and that he prefers the lottery

$$\tilde{W}: \quad \overset{\displaystyle 10/11 \diagup \$5M}{\underset{\displaystyle 1/11 \diagdown \$0}{\bigcirc}} \quad \text{to} \quad \text{a sure } \$1M.$$
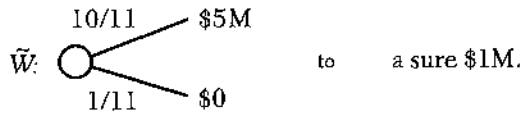
Let E be an event with probability .11, and consider the four event-contingent prospects:

	E (.11 chance)	~E (.89 chance)
a_1	\$1M	\$1M
a_2	$\tilde{W}$	\$1M
a_3	$\tilde{W}$	\$0
a_4	\$1M	\$0

It is straightforward to verify that these prospects generate the four correspondingly named lotteries from the Allais paradox.

If the individual must choose one of these four event-contingent prospects *before* learning whether the event E has or has not occurred, he will choose a_1, because a_1 first order stochastically dominates a_4, a_2 stochastically dominates a_3, and (by assumption) a_1 is strictly preferred to a_2. Thus, in the "no informa-

tion" case he will choose a_1, or in other words, the event-contingent bundle ($1M if E, $1M if $\sim E$).

But say the individual was informed that he would be given knowledge of whether E or $\sim E$ occurs before having to choose. If he were to learn that E occurred, he would be facing a choice of either $\tilde{W}$ (by choosing either a_2 or a_3) or a sure $1M (by choosing a_1 or a_4), and by the above-displayed preferences, he knows that he would choose to obtain $\tilde{W}$. On the other hand, if he were to learn that $\sim E$ occurred, he would clearly choose to obtain $1M rather than the alternative of $0.

Thus (alleges the argument), a "sophisticated" non-expected utility maximizer will realize that, if given prior information as to the occurrence/nonoccurrence of E, he would end up consuming the state-contingent prospect ($\tilde{W}$ if E, $1M if $\sim E$), which reduces to lottery a_2 from the Allais paradox. If *not* given this information, he would consume $1M with certainty (the lottery a_1). But because a_1 is preferred to a_2, the individual would accordingly *rather not* have the information![20]

4. CRITIQUE OF THE DYNAMIC ARGUMENTS

We have seen that while each of the *static* arguments against non-expected utility maximizers can be deflected by a

---

[20] As noted above, this argument does not rely upon the specifics of the Allais paradox but can be applied to any departure from expected utility. Recall that any such departure implies some preferences of the form $\tilde{X} > \tilde{Y}$ but $(\tilde{Y}, p; \tilde{Z}, 1 - p) > (\tilde{X}, p; \tilde{Z}, 1 - p)$. Let the event E have probability p and consider the event-contingent prospects $(\tilde{Y}$ if $E; \tilde{Z}$ if $\sim E)$ and $(\tilde{X}$ if $E; \tilde{Z}$ if $\sim E)$. If the individual had to choose without information, he would choose $(\tilde{Y}$ if $E; \tilde{Z}$ if $\sim E)$. If he could learn whether E occurred before making his choice, he would "realize" that he would choose to obtain $\tilde{X}$ if E occurred, and hence end up with the less preferred prospect $(\tilde{X}$ if $E; \tilde{Z}$ if $\sim E)$. Thus (goes the argument), he feels that the information would make him worse off.

properly designed non-expected utility model, the class of *dynamic* arguments is apparently more formidable. In this section we shall take a closer, more rigorous look at these arguments, discover that they each rely upon a hidden assumption concerning behavior in dynamic choice situations, and argue that this assumption is an inappropriate one to impose on non-expected utility maximizers.

4.1 The Hidden Assumption in These Arguments: Consequentialism

Consider an individual at the root of the dynamic decision problem illustrated in Figure 8. How would he act in such a situation? The classical economic model of choice assumes that, as in any decision problem, he would:

(a) determine the opportunity set implied by the situation,
(b) identify the most preferred element of this set, and
(c) adopt the strategy that leads to this most preferred element.

We have already noted that because such situations involve uncertainty, the elements of the opportunity set are not the alternative outcomes, but rather probability distributions over these outcomes. The opportunity set of distributions implied by this tree, and the strategies that generate each of them, are listed in Table 2. In the table, strategies are denoted by vectors of the form (U, U), (U, L), and so on, which specify how the agent would act at each of the two choice nodes, where U or L denote a choice of the upper or lower branch.

Because the third and fourth distributions on this list are stochastically dominated by the first and second distributions respectively, it is clear that neither of them could be optimal, so the individual must decide which of the two lotteries
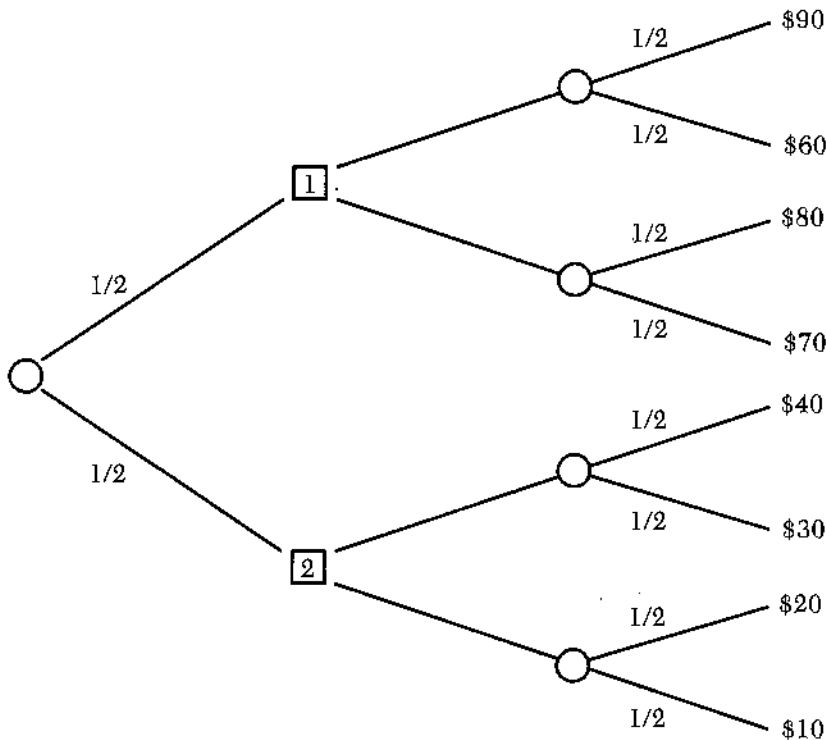
Figure 8. Example of a Decision Tree Representing a Dynamic Choice Problem

($90, ¼; $60, ¼; $40, ¼; $30, ¼) or
($80, ¼; $70, ¼; $40, ¼; $30, ¼)

is the more preferred. Say the former is
more preferred, in other words, say that

$$V(\$90, ¼; \$60, ¼; \$40, ¼; \$30, ¼)$$
$$> V(\$80, ¼; \$70, ¼; \$40, ¼; \$30, ¼)$$

so that the individual would plan on
choosing up at choice node 1 and up at
choice node 2, which is the strategy
(U, U).

So much for their *initial* choice of prob-
ability distribution and associated strat-
egy. We now turn to the question of how
individuals would behave if given the
chance to *reconsider* their plans in the
middle of a decision tree. Say that the
individual adopted the strategy (U,U)
and that nature chose "up" at the initial
chance node, so that he is now at choice
node 1 in the figure. Would a recalcula-
tion of his optimal strategy at this point
lead to a revision of his original plans?

TABLE 2

OPPORTUNITY SET OF PROBABILITY DISTRIBUTIONS FOR FIGURE 8

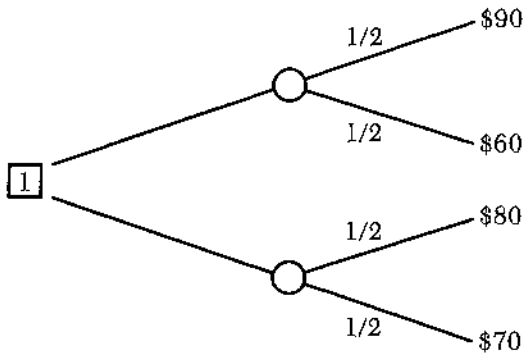| Strategy	Probability Distribution	Preference Function Value
(U, U)	($90, 1/4; $60, 1/4; $40, 1/4; $30, 1/4)	$V(\$90, 1/4; \$60, 1/4; \$40, 1/4; \$30, 1/4)$
(L, U)	($80, 1/4; $70, 1/4; $40, 1/4; $30, 1/4)	$V(\$80, 1/4; \$70, 1/4; \$40, 1/4; \$30, 1/4)$
(U, L)	($90, 1/4; $60, 1/4; $20, 1/4; $10, 1/4)	$V(\$90, 1/4; \$60, 1/4; \$20, 1/4; \$10, 1/4)$
(L, L)	($80, 1/4; $70, 1/4; $20, 1/4; $10, 1/4)	$V(\$80, 1/4; \$70, 1/4; \$20, 1/4; \$10, 1/4)$

Figure 9. Decision Tree Resulting from the Consequentialist Procedure of "Snipping Off" Just Before Choice Node 1 in Figure 8

TABLE 3

Opportunity Set at Choice Node 1 of Figure 8 Under the Assumption of Consequentialism

Strategy	Probability Distribution	Preference Function Value
(U)	($90, 1/2; $60, 1/2)	V($90, 1/2; $60, 1/2)
(L)	($80, 1/2; $70, 1/2)	V($80, 1/2; $70, 1/2)

On the one hand, it is certainly true that he has received new information, namely nature's choice at the initial chance node, and new information is often a cause for recalculations that lead to revisions in plans. On the other hand, because his plans at choice node 1 were *conditional on precisely this circumstance* (i.e., on nature choosing up), we might feel that a recalculation of his optimal strategy at this point ought to yield the same result. I want to argue that the key to the dynamic consistency issue lies in the manner in which such a recalculation is assumed to be undertaken.

Like the individual's initial calculation, any recalculation in the middle of a decision tree must involve specifying an opportunity set of alternatives and a maximand or preference function over this set. One thing that seems obvious is that, whatever the opportunity set is, it now consists of only two elements. But there are two different approaches to specifying just what these two elements, and the individual's ranking of them (i.e., their preference function values), are.

Roughly speaking, the *consequentialist* approach to this decision consists of "snipping" the decision tree at (that is, just before) the current choice node, throwing the rest of the tree away, and recalculating by applying the *original*

preference ordering (or equivalently, the original preference function) to alternative possible *continuations* of the tree. In other words, the individual would act as if he had *started out* with the tree in Figure 9, or equivalently, the opportunity set of Table 3, and make a choice based on his preference ranking over the lotteries

($90, ½; $60, ½)   vs.   ($80, ½; $70, ½)

that is, make the comparison

$$V(\$90, \tfrac{1}{2}; \$60, \tfrac{1}{2}) \lesseqgtr V(\$80, \tfrac{1}{2}; \$70, \tfrac{1}{2})$$

where $V(\cdot)$ is his original preference function. The philosophy behind this approach is that the uncertainty that was involved in the rest of the tree, as represented by the probabilities at the snipped-off chance nodes and the planned choices at the snipped-off choice nodes, is now irrelevant and should be treated as if it never existed.[21] In other words, the only determinants of how decisions should be made in the continuation of a decision tree are the *original* preference ordering over probability distributions and the attributes of the *continuation* of the tree.

What would it mean *not* to be consequentialist? Roughly speaking, it means that the individual would *not* snip off the

---

[21] In this case, the uncertainty involved in the snipped-off part of the tree was the ¼ chance of receiving $40 and the ¼ chance of receiving $30 which, having chosen the strategy (U, U), the individual bore at the time the initial chance node was resolved.

rest of the tree, but would instead take this past uncertainty (that is, the risks he has borne) into account in a manner consistent with his original preferences. But before formally presenting this alternative approach and arguing that it is more appropriate than consequentialism in the case of non-expected utility preferences, it is useful to go back and demonstrate how each of the three dynamic arguments against non-expected utility maximizers implicitly relies upon consequentialism to achieve its conclusion.

Consequentialism in the Dynamic Inconsistency Argument

The assumption of consequentialism is implicitly invoked in the dynamic consistency argument at the point where we asked how a non-expected utility maximizer would behave if nature were to lead them to the choice nodes in either of the Figure 7 decision trees. The argument maintained that at this point, the individual would basically snip off the unrealized lower branch, ignore both what its probability had been as well as what outcome it would have led to, and choose exactly as he would have in the first-step decision problem, that is to say, as if he were *starting out* with a choice between the lotteries

$\tilde{W}$: 10/11 → \$5M   versus   a sure \$1M.
      1/11 → \$0
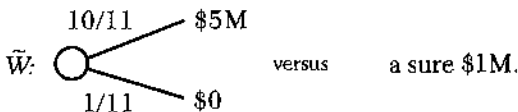
In other words, the argument has imposed the assumption that the individual's behavior at the choice node in each of the trees is fully determined by his answer to the first-step question.

Consequentialism in the Classical Making Book Argument

Consequentialism is invoked in the classical making book argument in its as-

sertion of how a non-expected utility maximizer owning the event-contingent prospect ($\tilde{Y}$ if E; $\tilde{Z} - \epsilon$ if $\sim E$) would behave if the event E should occur and he were then offered the lottery $\tilde{X} - \epsilon$ in exchange for his holding of $\tilde{Y}$. Once again, the argument asserts that his choice in this case would be determined by the original displayed preference ranking $\tilde{X} - \epsilon > \tilde{Y}$, that is, as if he were to snip off the "$\sim E$ branch" of the prospect and choose as if he had *started out* with a decision between the lottery $\tilde{Y}$ and the offer of $\tilde{X} - \epsilon$.

Consequentialism in the Aversion to Information Argument

The aversion to information argument invokes consequentialism when asserting how a "sophisticated" non-expected utility maximizer would predict his own behavior should the event E occur. In this case (goes the argument), the individual would ignore ("snip off") what would have happened in the event $\sim E$, consider himself back to a de novo choice between the lottery $\tilde{W}$ and a sure \$1M, and choose on the basis of his *original* ranking of these two prospects.

4.2 Consequentialism Is Inappropriate When Preferences Are Nonseparable

There is no question that, given their (implicit or explicit) assumption of consequentialism, these three dynamic arguments succeed in making individuals with non-expected utility preferences look rather foolish. However, the thrust of my critique of these arguments is that it is *inappropriate* to impose the property of consequentialism on non-expected utility maximizers. I want to argue that consequentialism is essentially a dynamic version of the very separability that non-expected utility maximizers reject, and that assuming it in this context is much like assuming, say, that agents with inter-

temporally nonseparable preferences would neglect their consumption histories when making subsequent decisions.

Parental Example

We can motivate this critique by an example involving the ultimate normative authority: one's Mom.[22] In this case, Mom has a single indivisible item—a "treat"—which she can give to either daughter Abigail or son Benjamin. Assume that she is indifferent between Abigail getting the treat and Benjamin getting the treat, and strongly prefers either of these outcomes to the case where neither child gets it. However, in a violation of the precepts of expected utility theory, Mom *strictly prefers* a coin flip over either of these sure outcomes, and in particular, strictly prefers ½ : ½ to any other pair of probabilities.

This random allocation procedure would be straightforward, except that Benjie, who cut his teeth on Raiffa's classic *Decision Analysis*, behaves as follows:

> Before the coin is flipped, he requests a confirmation from Mom that, yes, she does strictly prefer a 50:50 lottery over giving the treat to Abigail. He gets her to put this in writing.
>
> Had he won the flip, he would have claimed the treat.
>
> As it turns out, he loses the flip. But as Mom is about to give the treat to Abigail, he reminds Mom of her "preference" for flipping a coin over giving it to Abigail (producing her signed

statement), and demands that she flip again.

What would *your* Mom do if you tried to pull a stunt like this? She would undoubtedly say "You had your chance!" and refuse to flip the coin again. This is precisely what Mom does.

What is happening in this example? The set of possible outcomes is given by

(A, B) = (Abigail receives the treat, Benjamin receives the treat)[23]

Because Mom strictly prefers a 50:50 lottery to either *A* or *B*, she has the nonexpected utility (that is, nonseparable) preference ordering:
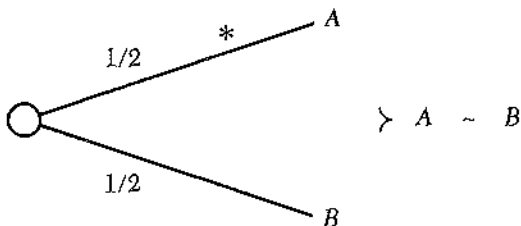
and so chooses the first of these (the lottery) over either sure outcome. The coin having landed in favor of Abigail (that is, nature having chosen the upper branch), Benjamin has tried to impose consequentialism on Mom by:

> snipping the tree at the point *
>
> throwing the rest of the tree away, and
>
> applying her *original* preference ranking to the *continuation* of the tree,

thereby trying to get her to replace:
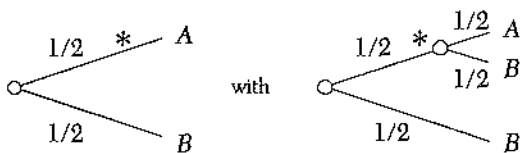
---

[23] The argument that this is *not* the set of appropriate outcomes is discussed in Section 6.6.

By replying "You had your chance," Mom is reminding Benjamin of the existence of the snipped-off branch (the original ½ probability of B) and that her preferences are *not* separable, so the fact that nature *could* have gone down that branch *still matters.* Mom is rejecting the property of consequentialism—and, in my opinion, rightly so.

What Is Mom (Or Any Other Nonseparable Agent) Telling Us?

Mom's original preference for (A, ½; B, ½) over (among other lotteries) (A, ¼; B, ¾), that is (under the reduction of compound lotteries axiom), her preference for the prospect

tells us that *conditional on having borne, but not realized, a ½ probability of B* (i.e., conditional on being at the point *), *she strictly prefers the outcome A to the lottery (A, ½; B, ½).* This being her attitude ex ante, Mom is only being dynamically consistent in maintaining the same attitude ex post, by refusing to flip the coin again at the point *.

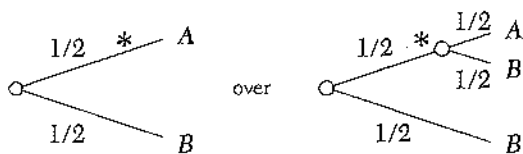This, in a nutshell, is why it is inappropriate to impose the property of consequentialism on non-expected utility maximizers. We have seen in Section 2 that such agents have nonseparable preferences across alternative events (or decision tree branches), so that their ex ante attitudes toward what happens along one branch may well depend upon what would have happened along the other branches. By forcing them to "snip" the rest of the tree once they are halfway down a branch and then act as if these other branches had not existed (i.e., act as if they were *starting out* with what

remains of the tree), consequentialism is essentially imposing separability upon them ex post, and it is no surprise that doing so would make their behavior look inconsistent (or worse), or that a non-expected utility maximizer like Mom would reject the validity of this procedure.

Analogy with Nonseparable Intertemporal Preferences

In order to best exploit our economists' intuition in this regard, it is useful to consider how the notions of nonseparability, consequentialism, and dynamic inconsistency would appear in a framework of choice over intertemporal consumption streams rather than uncertain prospects. The skeptical reader has my promise that, at the end of this brief excursion, I will be quite explicit about what I feel this intertemporal analogy *does* and *does not* have to say about the case of choice under uncertainty.

Consider the following preferences over consumption streams:

$$(\text{Star Wars I}) > (\text{Star Wars II})$$

$$\text{and (Star Wars I, Star Wars I)}$$
$$< (\text{Star Wars I, Star Wars II})$$

where "(Star Wars I)" denotes seeing that film once this evening, "(Star Wars I, Star Wars I)" denotes seeing it twice in succession, and so on. Because this individual's preferences for seeing Star Wars I versus Star Wars II depends upon whether or not he has just seen the former movie, his preferences over consumption streams are not intertemporally separable.[24] The analogue of "consequentialism" in such a context would be the property that, if given a

---

[24] This corresponds to a violation of mixture separability from Section 2.2. Replacing the top line of the displayed preferences by something like

(Jaws, Star Wars I) > (Jaws, Star Wars II)

would be a violation of replacement separability.

chance to revise at some point in time, the individual would

> snip the consumption stream at (that is, just before) the current point in time,
>
> throw the earlier part away, and
>
> recalculate by applying his *original* preference ranking to the alternative possible *continuations* of their consumption stream.

Given this, the analogue of Section 3.3's dynamic inconsistency argument would run as follows: "I note from your displayed preferences that if given a choice between seeing Star Wars I and Star Wars II, you would choose the former—let's just jot that fact down and set it aside for a moment. Now, let me offer you a different choice: We could either go to the Bijou Theater and see the double feature of Star Wars I and Star Wars II, or else we could go to the Paradise Cinema, which is showing only Star Wars I, but we could sit through it twice. The showings are at 7:00 P.M. and 9:00 P.M. at each theater."

Given your displayed preferences, you reply that we should go to the Bijou and see the pair of movies. But at the 8:45 intermission, I say, "Ah, but *now* we are back to a choice between seeing Star Wars II (by staying here) or else seeing Star Wars I (by going across town to the Paradise). Since I see from this piece of paper that you would rather see Star Wars I than Star Wars II, it follows that you would want to change theaters at this point" (or "that you have planned inconsistently," or "that you would be willing to pay me ε in order to change theaters," etc.).

Your response in this situation, which would in fact be the analogue of Mom's, is that the preferences written down on the paper were for an *outright* choice, and that your displayed preferences

made it clear from the start that *conditional on having already seen* Star Wars I, you would rather see Star Wars II than see Star Wars I again. A more formal statement of this might be as follows: "The intertemporal analogue of consequentialism, which states that agents would neglect their consumption histories when recalculating in the middle of an intertemporal choice situation, is clearly inappropriate to impose on an individual who has intertemporally nonseparable preferences. For such an agent, the fact that past consumption is gone in the sense of *consumed* does not mean that it is gone in the sense of *irrelevant.*"[25]

As noted above, it is important to be very explicit about what this analogy does and does not have to say about nonseparable preferences over alternative events. This analogy does *not* demonstrate that because it can be reasonable to have nonseparable preferences over time, it must, ipso facto, be reasonable to have nonseparable preferences over events. That, of course, is a non sequitur. Rather, it seeks to highlight the point that *if an individual informs you from the start that his preferences are nonseparable* (over time, over events, or over any other economic dimension), then it is inappropriate to impose separability ex post by explicitly or implicitly invoking consequentialism, and it is hardly surprising that doing so would

----

[25] The analogue of Section 3.3's "making book" argument would be similar: Start with an individual who owns tickets to (Star Wars I, Star Wars I) at the Paradise. By the second line of the above-displayed preferences, he would pay you ε for tickets to (Star Wars I, Star Wars II) at the Bijou. Then, at intermission, invoke the first line of the displayed preferences to get him to pay you ε to switch back to the Paradise and see Star Wars I. He has ended up seeing Star Wars I twice, which was his original holding, but has paid you 2ε in the process. Of course, the above objection to the dynamic inconsistency argument applies to this one as well.

lead to predictions of nonsensical behavior.

5. DYNAMICALLY CONSISTENT NON-EXPECTED UTILITY MAXIMIZERS

The analogy with intertemporal choice can fulfill an additional purpose, namely to motivate the manner in which nonseparable (that is, non-expected utility) preferences are *appropriately* extended to dynamic choice settings and how such an extension will lead to dynamically consistent choice behavior.

5.1 Extending Nonseparable Intertemporal Preferences to Dynamic Choice Situations

Consider an individual with the nonseparable preference ranking

> (pizza, pizza, salad, salad) >
> (pizza, pizza, pizza, pizza) >
> (pizza, pizza) > (salad, salad)

or equivalently

> V(pizza, pizza, salad, salad) >
> V(pizza, pizza, pizza, pizza) >
> V(pizza, pizza) > V(salad, salad)

where "(pizza, pizza)" denotes eating two (small) pizzas sequentially this evening, "(pizza, pizza, salad, salad)" denotes eating the two pizzas and then eating two salads, and so on. If such an individual could choose one of these four streams, he would clearly choose (pizza, pizza, salad, salad), and start eating pizza.

Now say he has finished his two pizzas, and were asked to reconsider his choice between a pair of salads at this point versus a third and fourth pizza. How would we properly represent the mathematics of his recalculation? The key point is to keep in mind that while the first two pizzas are gone in the sense of having been consumed, they are not gone in the sense

of *irrelevant*. Accordingly we could either

> Plug each of the *entire* time streams into the individual's *original* preference function $V(\cdot)$ (using bars to indicate the items that have already been consumed), so that the recalculation consisted of the comparison:

$$V(\overline{\text{pizza}}, \overline{\text{pizza}}, \text{salad}, \text{salad}) \text{ versus}$$
$$V(\overline{\text{pizza}}, \overline{\text{pizza}}, \text{pizza}, \text{pizza})$$

or equivalently:

> Plug the *continuation* of each time stream into the individual's *conditional* preference function $V_{\text{pizza, pizza}}(\cdot)$, so that the recalculation consisted of the comparison:

$$V_{\text{pizza, pizza}}(\text{salad, salad}) \text{ versus}$$
$$V_{\text{pizza, pizza}}(\text{pizza, pizza})$$

where we define

$$V_{\text{pizza, pizza}}(x, y) \equiv V(\text{pizza, pizza}, x, y).$$

Of course, these alternative notational procedures are equivalent, and both will imply that this nonseparable agent will stick with his original plan of moving on to the two salads rather than eating two more pizzas (i.e., will be "dynamically consistent"). In either case, the rest of the original displayed preference ranking, that is the ordering

> (pizza, pizza) > (salad, salad)
> or equivalently
> V(pizza, pizza) > V(salad, salad)

is quite irrelevant, given that at this point in the problem the individual has already eaten two pizzas. I, and I expect most economists, would maintain that the above procedure is the natural way to extend the individual's original nonseparable preferences over time streams to the dynamic choice problem generated by giving him the opportunity to reconsider halfway through the problem.
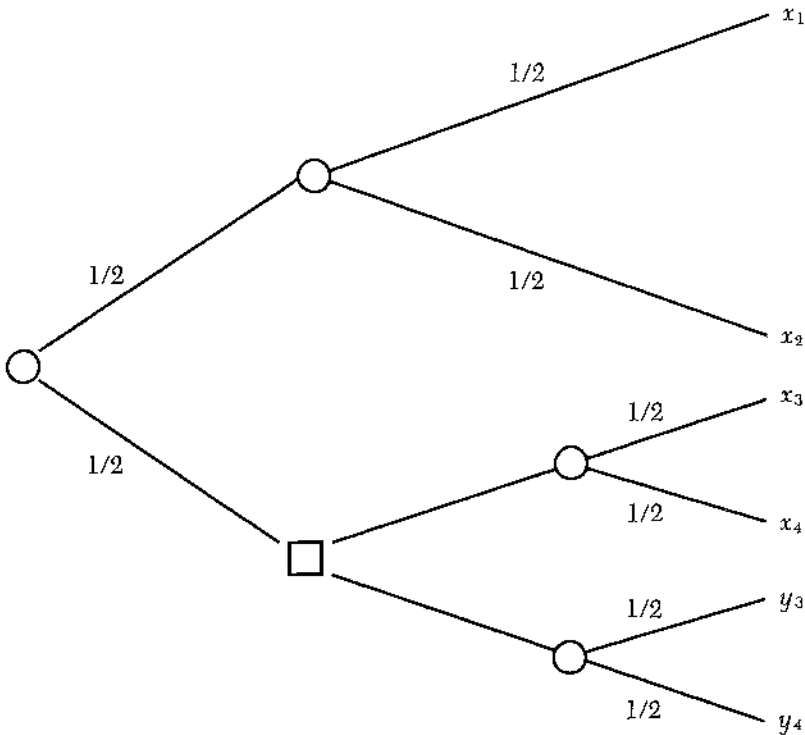
Figure 10. Decision Tree Illustrating a Dynamic Choice Situation

5.2 Extending Non-Expected Utility Preferences to Dynamic Choice Situations

The proper way to extend non-expected utility preferences to dynamic choice settings is completely analogous. Consider, for example, an agent with the non-expected utility preferences

$$(x_1, \tfrac{1}{4}; x_2, \tfrac{1}{4}; x_3, \tfrac{1}{4}; x_4, \tfrac{1}{4})$$
$$> (x_1, \tfrac{1}{4}; x_2, \tfrac{1}{4}; y_3, \tfrac{1}{4}; y_4, \tfrac{1}{4})$$
$$> (y_3, \tfrac{1}{2}; y_4, \tfrac{1}{2}) > (x_3, \tfrac{1}{2}; x_4, \tfrac{1}{2})$$

or equivalently:

$$V(x_1, \tfrac{1}{4}; x_2, \tfrac{1}{4}; x_3, \tfrac{1}{4}; x_4, \tfrac{1}{4})$$
$$> V(x_1, \tfrac{1}{4}; x_2, \tfrac{1}{4}; y_3, \tfrac{1}{4}; y_4, \tfrac{1}{4})$$
$$> V(y_3, \tfrac{1}{2}; y_4, \tfrac{1}{2}) > V(x_3, \tfrac{1}{2}; x_4, \tfrac{1}{2})$$

who faces the decision tree in Fig. 10. From the top line of his displayed preferences, he will clearly make plans to

choose up should he reach the choice node.

Say that nature chooses down at the initial chance node. What is the appropriate way to represent his recalculation? As with intertemporal nonseparability, the key thing is to remember that an agent with non-expected utility/nonseparable preferences feels (both ex ante and ex post) that risk which is borne but not realized, that is, the $\tfrac{1}{4}$ probabilities of having obtained x_1 or x_2, is gone in the sense of having been *consumed* (or "borne"), rather than gone in the sense of *irrelevant*. Accordingly, we could either:

Plug each of the *entire* probability distributions into the *original* preference function $V(\cdot)$ (using bars to indicate the risk that has already been borne) and make the comparison:

TABLE 4

OPPORTUNITY SET AT CHOICE NODE 1 OF FIGURE 8 WITHOUT THE ASSUMPTION OF CONSEQUENTIALISM

Strategy	Probability Distribution	Preference Function Value
$(U, \bar{U})$	($90, 1/4; $60, 1/4; $\overline{\$40, 1/4}; \overline{\$30, 1/4}$)	$V(\$90, 1/4; \$60, 1/4; \overline{\$40, 1/4}; \overline{\$30, 1/4})$
$(L, \bar{U})$	($80, 1/4; $70, 1/4; $\overline{\$40, 1/4}; \overline{\$30, 1/4}$)	$V(\$80, 1/4; \$70, 1/4; \overline{\$40, 1/4}; \overline{\$30, 1/4})$

$V(\overline{x_1}, \overline{1/4}; \overline{x_2}, \overline{1/4}; x_3, 1/4; x_4, 1/4)$ versus
$\quad V(\overline{x_1}, \overline{1/4}; \overline{x_2}, \overline{1/4}; y_3, 1/4; y_4, 1/4)$

or equivalently:

Plug the *continuation* of each branch into the *conditional* preference function $V_{x1, 1/4; x2, 1/4}(\cdot)$, and compare:

$V_{x1, 1/4; x2, 1/4}(x_3, 1/2; x_4, 1/2)$ versus
$V_{x1, 1/4, x2, 1/4}(y_3, 1/2; y_4, 1/2)$

where we define
$V_{x1, p1; x2, p2}(\check{Z})$
$\equiv V[x_1, p_1; x_2, p_2; \check{Z}, (1 - p_1 - p_2)].$

As in the intertemporal case, it is clear that these alternative notational procedures are equivalent, and that both will imply that the agent will stick with his original plan of choosing up at the choice node. Once again, the rest of the original displayed preference ranking, that is the ordering

$(y_3, 1/2; y_4, 1/2) > (x_3, 1/2; x_4, 1/2)$
or equivalently
$V(y_3, 1/2; y_4, 1/2) > V(x_3, 1/2; x_4, 1/2)$

is irrelevant, because by the time the individual is at the choice node, he or she has already "consumed" (that is, has already borne) some risk.

Going back to our "test example" of Figure 8, what does this approach imply about the opportunity set, preference function values, and ex post choice of an individual sitting at choice node 1 in that figure? Because the individual would clearly have planned on choosing up at choice node 2, we know that by the time

he is at node 1 he will have borne a 1/4 chance of receiving $40 and a 1/4 chance of receiving $30. The remaining opportunity set and its associated preference function values under this approach are accordingly *not* as given in Table 3, but rather, the first two elements of the *original* opportunity set and original preference function values from Table 2, using bars to indicate the risk that has already been borne (see Table 4). Under our original assumption that

$V(\$90, 1/4; \$60, 1/4; \$40, 1/4; \$30, 1/4)$
$> V(\$80, 1/4; \$70, 1/4; \$40, 1/4; \$30, 1/4)$

(displayed just prior to Table 2), it follows that the individual will indeed choose the upper branch upon arriving at choice node 1, just as originally planned.[26]

Finally, it is useful to mention an alternative way of operationalizing the notion of "borne risk." This is by means of the familiar Arrow-Debreu notion of a "contingent commodity," that is, a ticket that entitles the bearer to a particular outcome conditional upon the realization of a given state of nature (e.g., Kenneth Arrow 1964; Gerard Debreu 1959, Chapter 7). In this framework, "having borne the risk of an outcome x" is equivalent to having held, at the time the uncertainty was resolved, a ticket entitling you

---

[26] Antecedents of the formal approach presented here can be found in Anand (1987, pp. 209–10), Donaldson and Selden (1981, p. 213), Johnsen and Donaldson (1985, p. 1453, who, however, reject it), Machina (1981, p. 173), Loomes and Sugden (1984a, 1986), Yaari (1985), and McClennen (1988a, 1988b, 1989, who terms it "resolute choice").

to x in a state that ended up not being realized. The reader may verify that an individual who has non-expected utility (i.e., nonseparable) preferences over state-contingent outcomes, and who treats borne risk in the manner proposed in this section (that is, continues to take it into account ex post), will satisfy all of the same consistency properties.[27]

Immunity to the Dynamic Inconsistency Argument

Besides corresponding more closely to our treatment of nonseparable preferences in other economic settings, this manner of extending non-expected utility preferences to dynamic choice situations is immune to each of the dynamic arguments against non-expected utility maximizers presented in Section 3.3. As illustrated in the previous paragraph, the resolution of uncertainty that takes such an agent to a particular decision node does not lead to a *new* opportunity set of probability distributions and associated preference function values, but rather to that subset of the *original* opportunity set (and preference function values) that corresponds to the choices still available, with bars used to denote those risks that have already been borne. No matter what moves nature has made, the element of the original opportunity set that was most preferred (and hence corresponded to the ex ante plan)[28] will

always *remain* in this subset.[29] Because it will continue to have a preference function value higher than any of the remaining elements, the agent will always stick with their original plans. In other words, such agents will be dynamically consistent.

Immunity to the Classical Making Book Argument

Unlike the dynamic consistency argument, the classical making book argument involves introducing a *new* opportunity to the agent halfway through a dynamic choice problem. This consists of the offer to swap the prospect $\tilde{X} - \epsilon$ for $\tilde{Y}$ after the agent has borne a $1 - p$ chance of the prospect $\tilde{Z} - \epsilon$. Under our approach, the agent would evaluate this offer by comparing the distributions

$$(\tilde{X} - \epsilon, p; \overline{\tilde{Z} - \epsilon}, \overline{1 - p}) \text{ versus}$$
$$(\tilde{Y}, p; \overline{\tilde{Z} - \epsilon}, \overline{1 - p})$$

or equivalently, the preference function levels

$$V(\tilde{X} - \epsilon, p; \overline{\tilde{Z} - \epsilon}, \overline{1 - p}) \text{ versus}$$
$$V(\tilde{Y}, p; \overline{\tilde{Z} - \epsilon}, \overline{1 - p}).$$

But because we know that the individual has the ranking

$$(\tilde{Y}, p; \tilde{Z}, 1 - p)$$
$$> (\tilde{X}, p; \tilde{Z} - \epsilon, 1 - p)$$

we get

$$V(\tilde{Y}, p; \overline{\tilde{Z} - \epsilon}, \overline{1 - p})$$

---

[27] More formally, let $S = (s_1, \ldots, s_n)$ be a set of mutually exclusive and exhaustive states of nature, $C = (c_1, \ldots, c_n)$ be a "state-payoff bundle," which gives a consumption level of c_i should state i occur, and $V(c_1, \ldots, c_n)$ be the individual's preference function, which under the assumption of first-order stochastic dominance preference will be increasing in each argument. In this framework, the individual would take borne risk into account by drawing bars over the c_i levels for those states that end up not being realized.

[28] In the previous example, this was the distribution ($90, ¼; $60, ¼; $40, ¼; $30, ¼), which corresponded to the plan $(U, \bar{U})$.

[29] Why is this true? Let $\tilde{X}^*$ be the optimal distribution in the original opportunity set. As illustrated in the previous paragraph, the distributions still available at a particular choice node correspond to the set of possible strategies following from that node, combined with the risk that has been borne by this point, where that risk is determined by the original optimal strategy along each of the unresolved branches. Because $\tilde{X}^*$ is precisely the distribution that corresponds to having followed the optimal strategy along these unrealized branches, combined with making the optimal choices from this point on, it will be an element of this remaining subset of distributions.

$$> V(\tilde{X}, p; \overline{\tilde{Z}, 1 - p})$$
$$> V(\tilde{X} - \epsilon, p; \overline{\tilde{Z} - \epsilon, 1 - p})$$

so the individual will refuse this offer, and the argument cannot proceed beyond this point. More generally, an individual would only accept a new offer in the middle of a decision problem if, *when combined with the risk he has already borne*, it yields a higher preference function value than his original optimal distribution (and a fortiori, than his original holding). In terms of the contingent commodity version of this approach mentioned above, this implies that the individual cannot end up having held a collection of tickets (i.e., the ticket for the state that actually occurs plus the expired tickets for each of the other states) that is dominated by their original holding of tickets.

Immunity to the Aversion to Information Argument

The point at which the aversion to information argument fails against this approach is in its prediction that the Allais-type individual in the example would choose (or would foresee himself choosing) the lottery $\tilde{W}$ should he learn that the event E has occurred. Such a choice would indeed follow from the ranking $\tilde{W} > \$1M$ if the individual were consequentialist. However, because our individual would have borne a .89 chance of obtaining $\$1M$ by the time he learns that E has occurred, the appropriate comparison is not between $\tilde{W}$ and $\$1M$ but rather between the prospects

$$(.11, \tilde{W}; \overline{.89, \$1M}) \text{ versus}$$
$$(.11, \$1M; \overline{.89, \$1M})$$

or equivalently, the preference function values

$$V(.11, \tilde{W}; \overline{.89, \$1M}) \text{ versus}$$
$$V(.11, \$1M; \overline{.89, \$1M}).$$

Because these distributions are the same as the Allais choices a_2 and a_1 respectively, it follows that our individual would prefer the latter prospect, and hence stick with his planned choice of the $\$1M$ should E occur. Given this, information on whether or not E occurs will not cause him to depart from his original choice of $a_1 = (\$1M$ if E, $\$1M$ if $\sim E)$, and hence not have a negative value.

On the other hand, because he would choose a_1 whether or not he knew whether E occurred, advance information will not have any *positive* value in this example either. Does this mean that our approach is inferior to expected utility on the grounds that it assigns "too small" a value to information?

No. For one thing, the reader can verify that an individual with expected utility preferences will *also* assign exactly zero value to advance knowledge in this example. The reason is that, given the opportunity set $(\$1M, \tilde{W})$ should event E occur, and the opportunity set $(\$1M, \$0)$ should $\sim E$ occur, the set of event-contingent prospects (a_1, a_2, a_3, a_4) in this example forms a *complete event-contingent opportunity set*, in the sense that it allows for all combinations of choices from the available set $(\$1M, \tilde{W})$ under E and from the available set $(\$1M, \$0)$ under $\sim E$. In such circumstances, advance information concerning the event *ought* to have exactly zero value for either expected utility maximizers or (nonconsequentialist) non-expected utility maximizers.

Say, however, that an expected utility maximizer and a nonconsequentialist non-expected utility maximizer agreed that $a_1 = (\$1M$ if E, $\$1M$ if $\sim E)$ was preferred over the other three choices, and consider what would happen if this option were eliminated, so that the set (a_2, a_3, a_4) no longer formed a complete event-contingent opportunity set. Because a_2 stochastically dominates a_3,

which is in turn preferred to a_4, both of these individuals would pick a_2 if forced to make a choice before learning which event occurred. However, both of them would assign *positive* value to information regarding the occurrence or nonoccurrence of E in this case, because it would allow them an alternative way of attaining the preferred event-contingent option ($1M if E, $1M if $\sim E$), namely by choosing a_4 should E occur and a_2 should $\sim E$ occur.

In other words, expected utility maximizers and nonconsequentialist non-expected utility maximizers share the features that they will never assign a *negative* value to information, will assign a zero value to information when they face a complete event-contingent opportunity set, and will assign a positive (or at least nonnegative) value to information in the absence of a complete event-contingent opportunity set.

A Three-Way Classification of Decision Makers

We can summarize the above discussion with the following classification of decision makers in terms of their underlying preferences over lotteries and the manner in which they behave in dynamic choice situations:

α-*people: Expected utility maximizers.* Because their underlying preferences are separable anyway, their behavior in the middle of a decision tree will be the same whether or not they snip off the unrealized branches. Such individuals are therefore de facto consequentialist as well as dynamically consistent.

β-*people: Non-expected utility maximizers who are consequentialist.* These are the type of non-expected utility maximizers portrayed in the dynamic

arguments of Section 3.3. Such individuals *are not* dynamically consistent.

γ-*people: Non-expected utility maximizers who are not consequentialist.* These are the type of non-expected utility maximizers described in this section. Such individuals *are* dynamically consistent.[30]

6. MODELING NONSEPARABLE PREFERENCES UNDER UNCERTAINTY

The model presented in the previous section is quite different from the usual (that is, consequentialist) portrayal of non-expected utility maximizers, and the thoughtful reader will have undoubtedly anticipated several potential difficulties with and/or objections to this approach. The purpose of this section is to address and respond to these. Although the specific issues vary widely, the unifying theme behind my responses will be that the modeling of such agents involves no problems beyond those implied by the modeling of nonseparable preferences across time, commodities, or any other economic dimensions.

6.1 WHEN DO YOU START THE PROCESS AND WHAT IF YOU CAN'T OBSERVE THE PAST?

According to our approach, if the risk an individual bears in some event E affects his ex ante ranking of lotteries under some alternative event E^*, then it will have the same effect on his ex post ranking of these lotteries should the event E^* in fact occur. In the limit, this means that preferences over today's choice of lotteries are affected by what would have

---

[30] My friends Peter Hammond and Ned McClennen remind me that α, β, and γ were used to denote individuals of successively lower mental capacity in Aldous Huxley's *Brave New World*. Pay no attention to them.

occurred in each of the alternative un-
foldings of an individual's life. Doesn't
this make the modeling of non-expected
utility maximizers analytically intracta-
ble?

Not if we remind ourselves of the sim-
ple manner in which this is handled in
the case of nonseparable preferences
over consumption streams. In the inter-
temporally nonseparable case, what the
individual consumed yesterday (or last
month, or in his childhood) technically
affects his preferences over today's con-
sumption choices. However, this is han-
dled by simply subsuming past consump-
tion into the individual's current
preferences, so that what we refer to as
his preference ranking over consumption
streams beginning with the current pe-
riod are really his *conditional* prefer-
ences for such streams, given whatever
his past consumption happened to be.
It's true that without specific knowledge
of his past consumption stream we cannot
know the individual's *exact* preferences
from now on, but most economically use-
ful features of intertemporal preferences,
such as diminishing marginal rates of sub-
stitution across time-dated commodities,
will be inherited by these conditional
preferences. The only caveat about sub-
suming past consumption into current
preferences in this manner is that it must
be done *only once*, and at the *start* of
the problem at hand.

The case of nonseparable preferences
over events is identical. The effect of past
risks that have been borne can be sub-
sumed into the individual's preferences
(or preference function) over today's
lotteries. Economically meaningful prop-
erties of preferences such as quasi-con-
vexity or quasi-concavity in the probabili-
ties, or risk aversion, will be inherited.
Again, the only restriction is that this be
done *only once*, and at the *start* of the
decision problem.

6.2 How Far Do You Drag Along Unrealized Outcomes?

Of course, having subsumed all pre-
viously borne risk into preferences at the
outset, all risk borne during the *remain-
der* of a decision problem must be explic-
itly represented. In the simple examples
of Section 5.2, this was not very difficult.
But say the individual faces a very com-
plicated decision problem, with dozens
of choice and/or chance nodes (see the
examples of Raiffa 1968, pp. 12–13;
Schlaifer 1969, p. 92; Arnoldo Hax and
Karl Wiig 1977, p. 286). Are we really
supposed to carry the influence of each
of these alternative branches throughout
the whole problem?

In principle, yes. But, in principle, this
must also be done in any long intertem-
poral decision problem, or for that mat-
ter, any static multicommodity consump-
tion problem. After all, because there is
no normative justification for separability
across commodities, the individual's de-
mand for (say) peaches must in general
depend upon the prices of shoelaces,
phonograph needles, and so on. How-
ever, we live with this fact in the same
manner in each of these settings—by as-
suming (or hoping) that such cross-ef-
fects, while theoretically present, at least
approach zero as the commodities be-
come "farther and farther away" in some
appropriate sense. Thus, we assume that
the effect of consumption sufficiently far
in the past, or of a commodity that is
sufficiently unrelated, or of an unrealized
event that branched off sufficiently long
ago, has an effect that for all practical
purposes can be neglected. Lest this
sound as though I'm trying to invoke con-
sequentialism implicitly, recall that we
would *not* want to invoke this argument
for consumption in the very recent past,
for commodities (substitutes or comple-
ments) that are closely related, or for sig-

nificant (e.g., large probability or extreme outcome) risks that have recently been borne.

6.3 WHAT ABOUT THE SUNK COST FALLACY?

What's spent is spent and should have no effect on optimal plans from this point on.

This maxim, which is often termed the "sunk cost theorem," is taught in every principles of economics course. But if funds that actually *were* spent are irrelevant, doesn't it logically follow that outcomes that *didn't even occur* should also be irrelevant? No, it doesn't logically follow. Like so many of the other arguments in this literature, this argument is an example of taking a property that is a logical implication of separability and assuming that it must be true for the nonseparable case as well. To see this, consider what an example of the sunk cost principle looks like when it is actually formulated as a theorem (that is, as a formal statement capable of mathematical proof). Say a firm's maximization problem is given by:

$$\max_{L} R(L) - wL - wL_0$$

where $R(\cdot)$ is revenue (sales value of output) as a function of the input level L, w is the wage, and L_0 is some amount of labor that has already been contracted and paid for, but whose services were put to waste (the sunk cost of this problem) and hence does not enter into the revenue function. It is clear that the optimal value of L (call it L^*) will solve

$$R'(L^*) = w$$

and that this optimal value will not depend upon L_0. This is perhaps the simplest version of a formal sunk cost theorem.

But say the firm is a monopsonist in the labor market, so that its maximization problem is

$$\max_{L} R(L) - C(L + L_0),$$

where $C(L) \equiv L \cdot S^{-1}(L)$ is the total cost of inputs, given that the firm must raise its offered wage and move up the labor supply curve $S(\cdot)$ in order to obtain additional labor. In this case, the optimal value of L solves

$$R'(L^*) = C'(L^* + L_0),$$

which certainly *does* depend upon the sunk cost L_0.

The lesson, of course, is that sunk costs are irrelevant only when one's objective function is *separable* in the sunk cost variable—in this case, L_0—as with the first maximization problem but not the second. Thus, claiming that an agent with nonseparable (i.e., non-expected utility) preferences is committing the sunk cost fallacy reflects a misunderstanding of the mathematics of sunk costs.

6.4 WHAT ABOUT "FOLDING BACK"?

Besides α-people, β-people, and γ-people, there is another class of dynamically consistent non-expected utility maximizers—"δ-people"—who determine their optimal strategies in decision trees by a recursive process known in the decision theory literature as "folding back." Under this procedure, the individual begins by considering the "terminal choice nodes" of a decision tree (that is, those choice nodes not followed by any other choice node), determines the opportunity set of lotteries following from each such node, and uses his *original* preference function to determine the optimal choice out of that node. He then considers those choice nodes that are only followed by terminal choice nodes (or chance nodes) and repeats the process, subject to the previously determined

path out of each terminal choice node. This procedure of "folding back" then continues to earlier and earlier choice nodes, until the path out of each choice node in the tree has been determined. In the decision tree of Figure 5, for example, the individual would determine his optimal choice at the rightmost choice node on the basis of the comparison

$$V(x_3, 1) \lessgtr V(x_4, 1)$$

and would determine his optimal choice at the lower choice node on the basis of

$$V(x_5, r; x_6, 1 - r) \lessgtr V(x_7, 1).$$

In the event (say) that the left-hand quantity is greater in each of these cases, he would then determine his choice at the upper left choice node on the basis of the comparison

$$V(x_1, q; x_2, 1 - q) \lessgtr V(x_3, 1).$$

Such a model has been analyzed by Robert Weber (1982), Bell (1985), Karni and Safra (1986, 1988a, 1988b, 1988c), Gordon Hazen (1987), and Seidenfeld (1988a), who have applied it to various types of auctions and economic search problems.[31]

It is clear that δ-people will be dynamically consistent whether or not their preference functions are linear in the probabilities, because a repetition of this procedure halfway through any decision tree would produce the same optimal choice at each of the subsequent nodes. Because such individuals are also conse-

quentialist to boot,[32] why not adopt δ-people rather than γ-people as the model of the dynamically consistent non-expected utility maximizers we have been looking for?

One reason is that such a procedure implies some "undesirable" properties of behavior.[33] One implication, noted by Keeney and Winkler (1985), LaValle and Kenneth Wapman (1986), and Hammond (1988c), is that it can lead to different choices in strategically equivalent decision trees (that is, trees that imply the same opportunity sets of lotteries), which can in turn lead to nonindifference between such trees. To see this, take a δ-person with the typical Allais paradox preferences of

$$a_1 > a_2 \quad \text{and} \quad a_3 > a_4$$

who also happens to prefer

a sure \$1M   to $\widetilde{W}$:
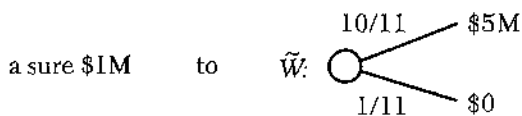
and consider the decision tree of Figure 11.[34]

Such an individual would act as follows: Because a_3 is strictly preferred to a_4, he would choose up in the event that he arrives at the upper choice node in the figure. Because a sure \$1M is strictly preferred to $\widetilde{W}$, he would choose down should he arrive at the lower choice node. This implies that his choice at the initial decision node is between a_3 if he

---

[32] This follows because the individual's action at each choice node depends only upon his *original* preference ordering and the continuation of the tree from that point on.

[33] At least, undesirable from the perspective of the approach proposed in this paper.

[34] In the opposite case when the individual preferred $\widetilde{W}$ to a sure \$1M, replace each occurrence of \$0 in Figure 11 by \$1M, replace a_3 and a_4 by a_2 and a_1, and then apply essentially the same argument as in the following paragraph.
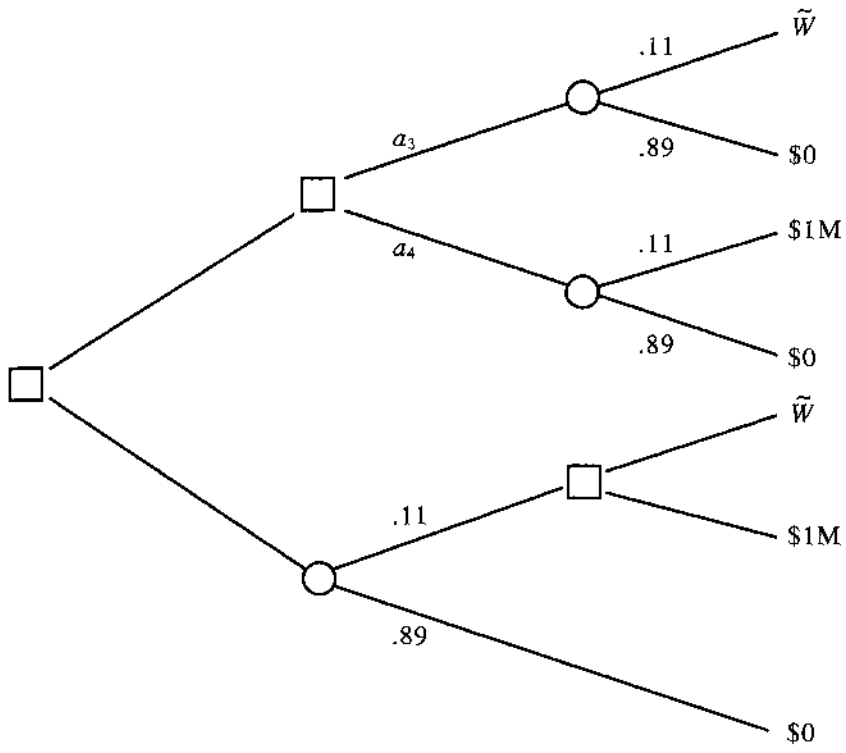
Figure 11. Choice Between a Pair of Strategically Equivalent Subtrees

chooses up or $a_4 = (\$1M, .11; \$0, .89)$ if he chooses down. Because we know that $a_3 > a_4$, such an individual would *strictly* prefer to choose up at the initial choice node in the figure. But because the upper and lower subtrees in this figure each imply an opportunity set of $(a_3, a_4) = [(\tilde{W}, .11; \$0, .89), (\$1M, .11; \$0, .89)]$, these subtrees are strategically equivalent, so that an individual who was indifferent between strategically equivalent trees (or subtrees) ought to be *indifferent* at the initial choice node. The reader can verify that if we replace the payoffs $\tilde{W}$ and 0 in the topmost sublottery by $\tilde{W} - \epsilon$ and $-\epsilon$ for some small enough ϵ, we get an example in which the individual would actually forgo the lottery $(\tilde{W}, .11; \$0, .89)$ (obtainable by choosing down at the initial and bottom choice nodes) in order to receive the *stochastically dominated* lottery $(\tilde{W} - \epsilon, .11; -\epsilon,$ .89) (by choosing up at the initial and top choice nodes).

A second undesirable implication is that such individuals are subject to a possible aversion to costless information in decision trees. Indeed, a reexamination of the aversion to information argument of Section 3.3 reveals that its notion of a "sophisticated" decision maker, that is, a consequentialist who anticipates his future decisions and takes them into account at earlier stages, is *precisely* what we have defined to be a δ-person.[35]

However, the more fundamental objection to this approach is that, as a formal optimization tool, *folding back is*

---

[35] Karni has pointed out that, in terms of the game-theoretic formulation of Footnote 31, aversion to information can be viewed as a strategy of precommitment on the part of the "first-stage agent" against subsequent moves on the part of the "latter-stage agents."

appropriate only when the objective function is separable across the various subdecisions of a problem, and this is simply not true for an individual with general nonseparable (that is, general non-expected utility) preferences who is facing a dynamic choice situation.

To see this in the context of nonseparable intertemporal preferences, say that my preferences over restaurant meals (written as consumption streams) are

(steak, sherbet) > (fish, banana split)
> (pizza, sherbet) > (all other meals)

and say that I "fold back"—that is, I work backward by choosing my dessert first (ignoring information about the entrées) and then choose the entrée subject to my choice of dessert. If my choice for dessert is sherbet, this procedure will lead me to a suboptimal choice of pizza should the evening's selection of entrées turn out to be (fish, pizza). But if my dessert choice is a banana split, this procedure will lead me to a suboptimal choice of fish should the selection of entrées turn out to be (steak, fish).

The problem of course is that my preferred dessert cannot be determined independently of the rest of the meal, or in other words, that my preferences over desserts and entrées are not separable. If we go back and look at the structure of objective functions in intertemporal problems where folding back *can* legitimately be applied, such as the "stagecoach problem" (e.g., Harvey Wagner 1985, pp. 265–70) or intertemporal maximization of discounted profits or discounted utility, we see that they are each separable across the different stages of the problem (i.e., different time periods), which implies that the optimal continuation of a stagecoach route (or profit or consumption stream) from a given point is independent of the particular route (or stream) that led up to that point. In other words, folding back, like consequential-

ism, is inherently inappropriate to apply to nonseparable preferences, be they over multicommodity consumption bundles, time streams of consumption, or lotteries.

6.5 HIDDEN NODES AND BRANCHES: COORDINATION IN THE FACE OF ZERO PROBABILITY EVENTS

We have seen how the model of non-expected utility preferences developed in Section 5 is immune to the "classical" making book argument. However, this approach is subject to a different type of procedure, which would appear to be another form of "making book."[36]

Consider, for example, a non-expected utility maximizer with the preferences

$$(\tilde{X}, \tfrac{1}{2}; \tilde{X}, \tfrac{1}{2})$$
$$\sim (\tilde{Y}, \tfrac{1}{2}; \tilde{Y}, \tfrac{1}{2}) > (\tilde{Y}, \tfrac{1}{2}; \tilde{X}, \tfrac{1}{2})$$

for some pair of lotteries $\tilde{X}$ and $\tilde{Y}$. By continuity, it follows that

$$(\tilde{X} - \epsilon, \tfrac{1}{2}; \tilde{X}, \tfrac{1}{2}) > (\tilde{Y}, \tfrac{1}{2}; \tilde{X}, \tfrac{1}{2})$$
$$\text{and} \quad (\tilde{Y}, \tfrac{1}{2}; \tilde{Y} - \epsilon, \tfrac{1}{2}) > (\tilde{Y}, \tfrac{1}{2}; \tilde{X}, \tfrac{1}{2})$$

for some small positive ϵ. Say the individual is currently endowed with the event-contingent prospect ($\tilde{X}$ if heads; $\tilde{Y}$ if tails) for some fair coin. Because he would be indifferent between this and the reversed prospect ($\tilde{Y}$ if heads; $\tilde{X}$ if tails), a "manipulator" who owned this latter prospect could presumably convince the individual to swap his original endowment for this new one. Say the two of them make this exchange.

Now let the coin be flipped. If it lands heads, the individual receives the lottery $\tilde{Y}$, having borne a 50 percent chance of having received the lottery $\tilde{X}$. But the ranking $(\tilde{X} - \epsilon, \tfrac{1}{2}; \tilde{X}, \tfrac{1}{2}) > (\tilde{Y}, \tfrac{1}{2}; \tilde{X}, \tfrac{1}{2})$ implies that under these circumstances (that is, having borne a 50 percent chance of $\tilde{X}$), he would rather pos-

---

[36] This procedure is based on examples by Uzi Segal and Eddie Dekel.

TABLE 5

A New Example of "Making Book"?

	Individual's Holdings		Manipulator's Holdings	
	Heads	Tails	Heads	Tails
Original endowment	$\tilde{X}$	$\tilde{Y}$	$\tilde{Y}$	$\tilde{X}$
After preflip trade	$\tilde{Y}$	$\tilde{X}$	$\tilde{X}$	$\cdot\tilde{Y}$
After postflip trade	$\tilde{X} - \varepsilon$	$\tilde{Y} - \varepsilon$	$\tilde{Y} + \varepsilon$	$\tilde{X} + \varepsilon$

sess $\tilde{X} - \epsilon$ than $\tilde{Y}$. In other words, should heads come up, the individual would now be willing to trade his holding of $\tilde{Y}$ for the lottery $\tilde{X} - \epsilon$. A similar argument shows that in the event of a tail, the individual would be willing to trade his holding of $\tilde{X}$ for the lottery $\tilde{Y} - \epsilon$. Table 5 illustrates the respective holdings of the individual and the manipulator at each stage of this process.

Have we finally succeeded in "making book" after all? To see what is happening

in this example, consider the decision tree representation of this procedure in Figure 12. The initial choice node represents the preflip offer, where the individual may either remain with his endowment ($\tilde{X}$ if heads; $\tilde{Y}$ if tails) by choosing down or else accept the preflip offer of ($\tilde{Y}$ if heads; $\tilde{X}$ if tails) by choosing up. As indicated by the arrow out of the initial choice node, the individual accepts the preflip offer and awaits the flip of the coin at the upper chance node, under the impression that he will now consume $\tilde{Y}$ if it lands heads and $\tilde{X}$ if it lands tails.

Each of the broken squares and lines in the figure indicates a "hidden" node or branch, that is, a postflip trade opportunity that is not revealed to the individual unless and until he actually gets to that particular point. Say the coin in the upper chance node lands heads, so that the individual, expecting to receive $\tilde{Y}$ having borne a 50 percent chance of $\tilde{X}$,
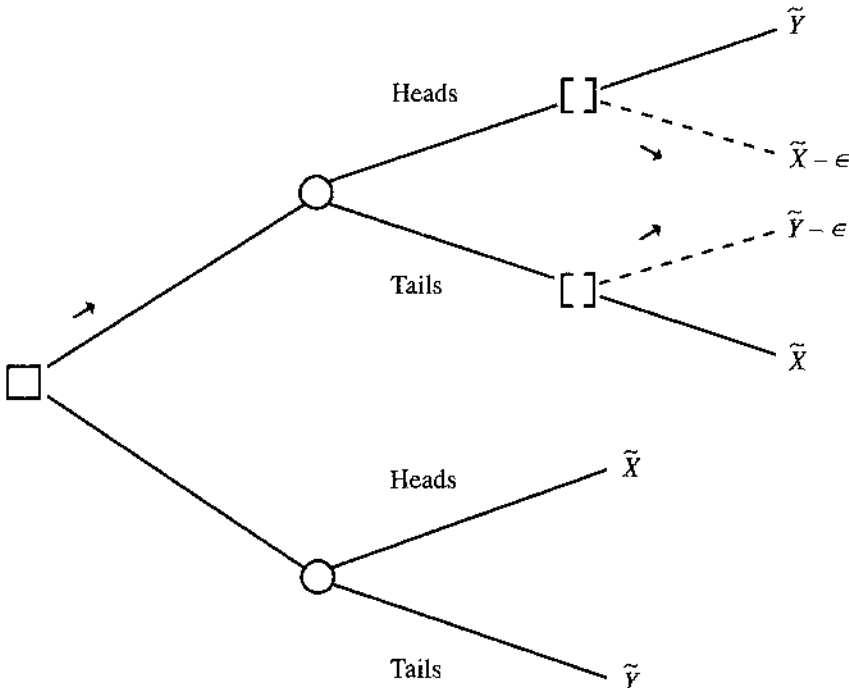

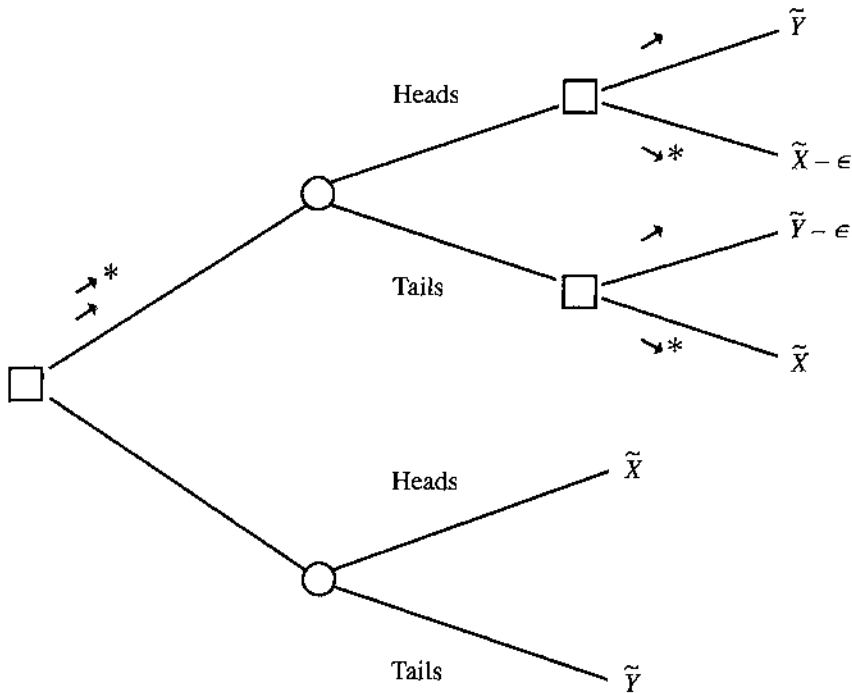*Figure* 12. Decision Tree Representation of the Procedure in Table 5

Figure 13. Individual's Choices Given Prior Knowledge of Both Postflip Offers

is suddenly offered $\tilde{X} - \epsilon$ (in other words, the upper hidden choice node is revealed). Because the individual is not (and at this point, never will be) aware of the *other* hidden choice node, he is under the impression that he really would have received $\tilde{X}$ had the coin landed tails, and it is on this basis that he opts for the postflip swap to $\tilde{X} - \epsilon$, as indicated by the arrow in the figure. A corresponding situation would occur if the coin had landed tails.

To see that this procedure depends upon the manipulator keeping these nodes hidden, consider how the individual would have acted had he been aware of these nodes from the start, that is, if he had faced the decision tree in Figure 13. Given his original displayed preferences, the individual would have been indifferent between the strategy indicated by the starred arrows, which would yield the distribution ($\tilde{X} - \epsilon$, ½; $\tilde{X}$, ½), or the strategy indicated by the unstarred

arrows, which would yield the distribution ($\tilde{Y}$, ½; $\tilde{Y} - \epsilon$, ½). If, for example, he adopted the starred strategy, he would plan on making the postflip exchange if heads should come up, but not if tails comes up.

In other words, given full knowledge of the set of available options (nodes and branches) in the tree, the individual would willingly forego ϵ to make exactly *one* of the postflip swaps (i.e., would pay $\epsilon/2$ in expected value terms). The reason the manipulator can get him to pay ϵ no matter how the coin lands is by hiding the fact that such an offer would have been available to the individual even if the flip had turned out the other way.

This leads to the question of precisely which comparison should be used in judging whether the individual has acted irrationally. Do we compare the outcomes under the individual's endowment ($\tilde{X}$ if heads, $\tilde{Y}$ if tails) with the outcomes

that would actually arise from the procedure (that is, $\bar{X} - \epsilon$ if heads and $\tilde{Y} - \epsilon$ if tails)? Or do we make a separate comparison for each state of nature, taking into account the different information that the individual has in each state? In the latter case, this would mean the *separate* comparisons:

For heads: Choosing $\bar{X} - \epsilon$ over $\tilde{Y}$ having (so he thought) borne a 50 percent chance of $\bar{X}$

For tails: Choosing $\tilde{Y} - \epsilon$ over $\bar{X}$ having (so he thought) borne a 50 percent chance of $\tilde{Y}$.

Given the individual's original preference rankings $(\bar{X} - \epsilon, \frac{1}{2}; \bar{X}, \frac{1}{2}) > (\tilde{Y}, \frac{1}{2}; \bar{X}, \frac{1}{2})$ and $(\tilde{Y}, \frac{1}{2}; \tilde{Y} - \epsilon, \frac{1}{2}) > (\tilde{Y}, \frac{1}{2}; \bar{X}, \frac{1}{2})$ and given the information he has in each event, each of these decisions seems to be fully rational.

How does this example differ from the "classical" making book argument of Section 3.3? That argument also involved an unforeseen ex post choice, specifically, the offer of $\bar{X} - \epsilon$ in exchange for $\tilde{Y}$ should the event E occur. But because it involved only a *single* unforeseen choice, the individual's response to it was made with correct information as to what otherwise would have happened (recall that no choice at all was offered in the other event). The present argument relies upon unforeseen choices in *each* event, so that in either case, the individual makes a choice based on incorrect information regarding his opportunities (and hence the risk he would actually bear) in the opposite event.

As with the other issues considered in this section, the problem of coordination in the light of unforeseen events is an inherent feature of nonseparable preferences, under uncertainty or certainty. To see this, consider the following example: You are throwing a dinner party and you want the wine and food to match, so that your preferences are

(beef, red wine) $\sim$ (fish, white wine)
$>$ (beef, white wine) $\sim$ (fish, red wine)

Say you call the local butcher, who says that there is no chance of getting any fish today, although they do have beef. You also call the local wine shop, and are told that there is no chance of getting any red wine today, although they do have some nice whites. While beef with white wine is not your favorite dinner, it is better than nothing. Accordingly, you send your cook to the butcher and your butler to the wine shop. Both are aware of your preferences, as well as the information you have received from the two shops. Upon arriving at the butcher, your cook finds that there has been an unexpected shipment of fish. Given his knowledge of your preferences and the fact that the wine shop had only white wine, your cook would be willing to pay an additional ϵ to purchase the fish. But in the meantime, your butler has arrived at the wine shop, to find (you guessed it) that a surprise shipment of red wine has come in, and acts similarly. By trying to respect your nonseparable commodity preferences in response to these unforeseen events, and unaware that there was an unforeseen event in *each* store, your agents have between themselves made you worse off. If we were to assume an initial stage where you had an endowment of (fish, red wine) that you traded for (beef, white wine), this story could be turned into a full-fledged example of "making book."

The problem, of course, is that decision makers with nonseparable preferences under uncertainty *or* certainty require full knowledge of the opportunities available in each component of the problem (for both heads and tails, or both the entrée and the wine) in order to make the proper choice in any one component. When this information is not known, or is incorrectly known, suboptimal and

even dominated choices can be made.

How to handle this problem? One possibility is to react to the sudden appearance of an unforeseen opportunity (a postflip choice, or the availability of fish) by realizing that if one's model of the opportunity set in *this* component was not correct, then it might not be correct in the other component. In formal terminology, you must "recondition on a zero probability event." How an individual would (or should) do this, however, is an open question (see, for example, Kreps 1989).

Another approach would be to try to plan ahead for all possible "unforeseen" contingencies ex ante, by including them in the original model of the decision tree. In the limit, however, this may mean specifying an intractable number of alternatives out of every choice node.

Perhaps the best approach is simply to recognize the interdependence of component-wise decisions under nonseparable preferences, and to remain mindful of its stronger informational requirements. In other words, when the manipulator offered the original preflip swap, the individual should have recognized his cross-state dependence and asked, "Do you have any *other* offers in store for me?" We have seen that a truthful answer to this question would prevent the manipulator from being able to "make book." Lying and saying "no" (and being believed) would have the same effect as a lie would in an exchange under certainty, namely that the individual could end up being exploited. An answer of "maybe" ought to lead the individual to incorporate formally the possibility of these other offers by the addition of chance nodes with branches leading out to these possible choice nodes.

The upshot, as mentioned above, is that making decisions with nonseparable preferences under uncertainty has greater informational requirements than

it does with separable preferences, just as it does under certainty (would you think of ordering the wine at a restaurant before looking at the list of entrées?). As with the other issues discussed in this section, I would suggest that the way we live with this fact be analogous to the way we live with it under certainty, namely by restricting the interaction effects we consider to the ones that we feel are the most important, but then explicitly taking them into account in the manner outlined in Section 5.

6.6 The Operational Definition of Consequences

A final objection to this approach, and in some sense to the formal modeling of nonexpected utility preferences in general, is the allegation that observed "violations" of the expected utility hypothesis are in fact not violations at all, but merely examples of improper definition of the basic outcomes or consequences. We can motivate this discussion by means of a simple example, based upon the ideas and examples of Samuelson (1952a), Drèze (1974), Machina (1981), and Sen (1985).

Say I face the four event-contingent prospects

	.03 chance	.97 chance
c_1	hamburger	bicycle
c_2	movie	bicycle
c_3	hamburger	WK
c_4	movie	WK

In this example, "movie" means seeing a new romance film featuring my favorite movie star, with whom I am deeply in love (not platonically); "bicycle" is a new ten-speed bike; and "hamburger" is a gourmet hamburger at a local restaurant. If I had to choose between c_1 and c_2, I would certainly hope that I win the bicycle (and given the odds, I probably will),

but if I don't, let's assume that I would rather see the movie, so I prefer c_2 over c_1.

How should I rank the options c_3 versus c_4? I hope that no reader who claims to subscribe to the precepts of expected utility theory feels he or she needs to know what "WK" stands for, because by replacement separability (or the independence axiom) I ought to prefer c_4 over c_3 *regardless* of the nature of this outcome.

In fact, "WK" stands for "week on a secluded tropical island with selfsame movie star." Given this, I reason as follows: I have a 97 percent chance of bliss. On the other hand, what would I feel like doing in the unlikely and unlucky event that I *don't* win that week on the island? Do I really want to be sitting (alone) in a (cold) movie theater, watching the object of my affection sail off into the sunset with *somebody else?* No, I would rather be eating a hamburger in that event, and accordingly, I would prefer c_3 over c_4.

Such preferences don't seem irrational. Do they violate the independence axiom, or equivalently, separability? To me, the answer is "Yes—in fact, by definition." However, defenders of expected utility might argue as follows:

"A preference of c_1 over c_2 and c_4 over c_3 in this example does not violate separability across alternative outcomes, because what we call 'movie' in prospect c_4 is in fact a *different* outcome from 'movie' in c_2. Specifically, 'movie' in c_2 can be described as 'seeing the movie,' but 'movie' in c_4 is 'seeing the movie *while disappointed.*' Because these are clearly different outcomes, let us call the latter one 'movie*' to distinguish it from the former, in which case a preference for c_1 over c_2 and c_4 over c_3 in this example is completely consistent with expected utility."

I certainly agree that my *attitude* toward seeing the movie is much different in c_2 from my attitude in c_4. But is it legitimate to claim that I have not violated expected utility on the grounds that they are *different outcomes?* As alluded to in Footnote 23, a similar objection might be made that A = "Abigail getting the treat outright" and A^* = "Abigail getting the treat when Benjamin also had a fair chance at it" are really different outcomes in the "Mom" example of Section 4.2.

Once again, consider how this kind of argument would sound in the more familiar context of standard consumer theory. In the following table, the options d_1 through d_4 are not event-contingent prospects, but rather standard commodity bundles involving the *simultaneous* consumption of a beverage and a condiment, so that d_1 denotes drinking milk *with* tea, and so on. Now it happens that I like lemon with my tea but milk with my coffee, so that I would prefer d_2 over d_1 but prefer d_3 over d_4.

d_1	milk	tea
d_2	lemon	tea
d_3	milk	coffee
d_4	lemon	coffee

Personally, I would simply describe my preferences as nonseparable and be done with it. But say someone tried to make the following argument:

"Your preference for d_2 over d_1 and d_3 over d_4 in this example does not really violate separability between beverage and condiment, because what we call 'lemon' in bundle d_4 is in fact a *different* commodity from 'lemon' in d_2. Specifically, 'lemon' in d_2 can be described as 'drinking lemon with your beverage,' but it is clear that 'lemon' in d_4 means 'drinking lemon *while disgusted.*' We should accordingly call the latter 'lemon*' to distinguish it from

the former, in which case your preference for d_1 over d_2 and d_4 over d_3 is completely consistent with separability."

Once again, I agree that my *attitude* toward adding lemon, even the sign of its marginal utility, is different in the bundle d_2 from my attitude in the bundle d_4. But the key question is, *why* is my attitude different? Is it really because it is a *different commodity* in these two bundles, or is it because it is *combined* with different commodities in these two bundles? I think that most economists would say the latter, and pronounce a verdict of nonseparability. After all, classical consumer theory presumes that we can define each "commodity" (such as milk or lemon juice) on the basis of its attributes alone, independent of which other items one consumer or another happens to like (or dislike) mixing it with. On the other hand, perhaps the maxim *de gustibus non disputandum est* compels us to respect the wishes (including the definitions) of any consumer who feels that "lemon juice that is put in tea" and "lemon juice that is put in coffee" really are two different commodities.

The situation is similar for choice under uncertainty. Ideally, a fair application of the expected utility model should proceed as follows:

(a) We begin by agreeing upon a set of "consequences" $(x_1, x_2, \ldots)$.

(b) The axioms then imply that we can assign individual utility levels $U(x_1)$, $U(x_2)$, and so on, to each of these consequences, with the property that

(c) The expectation of these utility levels can be used to evaluate the relative desirability of any probability distribution over these consequences.

If a proponent of expected utility should decide in retrospect that receiving the amount x_1 in some lottery $\bar{Y}$ is in fact a different consequence from receiving x_1 in some other lottery $\bar{Z}$ (or that "movie" in the prospect b_2 is a different consequence from "movie" in the prospect b_4), then perhaps we really should grant him the right to go back to step (a) and start over. However, it is important to realize that invoking such a right is tantamount to defending the expected utility model by rendering it irrefutable, because, for example, individuals with the typical preferences in the Allais paradox could also prove their "consistency" with expected utility by claiming that the $0 prize in the prospect a_2 is really a different "consequence" from the $0 prize in a_3, and so on, and the dynamic arguments of Section 3.3 would be helpless against this defense.

I have engaged in the above "definition of a consequence" argument with several proponents of expected utility whom I greatly respect. Experience shows that the only mutually acceptable way out of it is to adopt the following three-part "compromise":

1. The properties of separability/nonseparability must always be discussed with reference to a *given* set of consequences (that is, with respect to a particular level of description of the consequences).

2. For my part, I will grant that separability may well be rational provided the descriptions of the consequences are sufficiently deep to incorporate any relevant emotional states, such as *disappointment* (e.g., at having won $0 when you might have won $5 million), *regret* (at having forgone a sure chance of $1 million and then landing a 1 percent chance of $0), *jealousy* (over your favorite movie star), *feelings of unfairness* (that Benjamin won the treat in an unfair flip), and so on.

3. However, for their part, proponents of expected utility must grant that this level of description *may be below the usual level at which economists typically operate with or can observe,* so that preferences over observables such as monetary outcome levels, which child ends up receiving the treat, and so on, could legitimately be nonseparable. In other words, the various non-expected utility models of Table 1 could legitimately represent risk preferences when the consequences consist of monetary outcome levels.

These ideas are not new. As Samuelson (1952a, pp. 676–77) noted some time ago, separability across alternative consequences

> must always be applied to a definite set of entities—e.g., (1) single-event money prizes, (2) single-event vectors of goods, (3) single-event money prizes *cum* gaming and suspense feelings . . . [Separability] then has implications and restrictions upon choices among such entities; but, strictly speaking, it need not impose restrictions upon some different (and perhaps simpler) set of entities.
>
> In what dimensional space are we "really" operating? If every time you find my axiom falsified, I tell you to go to a space of still higher dimensions, you can legitimately regard my theories as irrefutable and meaningless . . . From my own direct and indirect observations, I am satisfied that a large fraction of the sociology of gambling and risk taking will never significantly be discernible in terms of money prizes alone, as distinct from elements of suspense . . .

The above compromise tries simultaneously to acknowledge (*a*) the normative appeal of separability at some deep enough level of consequence description, (*b*) normative reasons why preferences might be nonseparable at the level of description typically used by economists, and accordingly (*c*) the potential value of non-expected utility models for descriptive economists who can only observe (or only work in terms of) the usual

economic variables. Along with the critique of Section 4 and the dynamic model of Section 5, it is offered as a contribution to what I have termed the "normative goal" in the campaign for the general acceptance and use of non-expected utility models.

References

ALLAIS, MAURICE. "Fondements d'une Théorie Positive des Choix Comportant un Risque et Critique des Postulats et Axiomes de L'Ecole Americaine," *Econométrie*, 1953a, *40*, pp. 257–332. (English translation: "The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School," in ALLAIS AND HAGEN 1979, pp. 27–145.)

———. "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica*, Oct. 1953b, *21*(4), pp. 503–46. (Summarized version of ALLAIS 1953a)

———. "The So-called Allais Paradox and Rational Decisions Under Uncertainty," in ALLAIS AND HAGEN 1979, pp. 437–681.

ALLAIS, MAURICE AND HAGEN, OLE, eds. *Expected utility hypotheses and the Allais Paradox.* Dordrecht, Holland: D. Reidel Publishing Co., 1979.

ALLEN, BETH. "Smooth Preferences and the Approximate Expected Utility Hypothesis," *J. Econ. Theory*, Apr. 1987, *41*(2), pp. 340–55.

ANAND, PAUL. "Are the Preference Axioms Really Rational?" *Theory and Decision*, Sept. 1987, *23*(2), pp. 189–214.

ARROW, KENNETH. "The Role of Securities in the Optimal Allocation of Risk-Bearing," *Rev. Econ. Stud.*, Apr. 1964, *31*(1), pp. 91–96.

BAR-HILLEL, MAYA AND MARGALIT, AVISHAI. "How Vicious Are Cycles of Intransitive Choice?" *Theory and Decision*, Mar. 1988, *24*(2), pp. 119–45.

BATTALIO, RAYMOND; KAGEL, JOHN AND MACDONALD, DON. "Animals' Choices over Uncertain Outcomes," *Amer. Econ. Rev.*, Sept. 1985, *75*(4), pp. 597–613.

BAUMOL, WILLIAM. "The Neumann-Morgenstern Utility Index: An Ordinalist View," *J. Polit. Econ.*, Feb. 1951, *59*(1), pp. 61–66.

———. "The Cardinal Utility Which Is Ordinal," *Econ. J.*, Dec. 1958, *68*(272), pp. 665–72.

BELL, DAVID. "Disappointment in Decision Making Under Uncertainty," *Operations Research*, Jan./Feb. 1985, *33*(1), pp. 1–27. (Reprinted in BELL, RAIFFA, AND TVERSKY 1988, pp. 358–83)

BELL, DAVID; RAIFFA, HOWARD AND TVERSKY, AMOS, eds. *Decision making: Descriptive, normative, and prescriptive interactions.* Cambridge: Cambridge U. Press, 1988.

BINMORE, KENNETH. "Game Theory and the Social Contract: Mark II," manuscript, London School of Economics, 1988.

BORDER, KIM. "Stochastic Dominance and Choice

of Lotteries," manuscript, Division of Humanities and Social Sciences, California Institute of Technology, 1987.

BROOME, JOHN. "Equity in Risk Bearing," *Operations Research*, Mar./Apr. 1982, 30(2), pp. 412–14.

———. "Uncertainty and Fairness," *Econ. J.*, Sept. 1984, 94(375), pp. 624–32.

CHEW, SOO HONG. "A Generalization of the Quasilinear Mean with Applications to the Measurement of Inequality and Decision Theory Resolving the Allais Paradox," *Econometrica*, July 1983, 51(4), pp. 1065–92.

CHEW, SOO HONG AND EPSTEIN, LARRY. "The Structure of Preferences and Attitudes Towards the Timing of the Resolution of Uncertainty," *Int. Econ. Rev.*, Feb. 1989a, 30(1), pp. 103–17.

———. "Non-Expected Utility Preferences in a Temporal Framework with an Application to Consumption-Savings Behavior," *J. Econ. Theory*, 1989b.

CHEW, SOO HONG; EPSTEIN, LARRY AND SEGAL, UZI. "Mixture Symmetric Utility Theory," manuscript, Dept. of Economics, U. of Toronto, 1988.

CHEW, SOO HONG; EPSTEIN, LARRY AND ZILCHA, ITZHAK. "A Correspondence Theorem Between Expected Utility and Smooth Utility," *J. Econ. Theory*, Oct. 1988, 46(1), pp. 186–93.

CHEW, SOO HONG; KARNI, EDI AND SAFRA, ZVI. "Risk Aversion in the Theory of Expected Utility with Rank Dependent Probabilities," *J. Econ. Theory*, Aug. 1987, 42(2), pp. 370–81.

CHEW, SOO HONG AND MACCRIMMON, KENNETH. "Alpha-Nu Choice Theory: A Generalization of Expected Utility Theory," manuscript, Faculty of Commerce and Business Administration, U. of British Columbia, 1979a.

———. "Alpha Utility Theory, Lottery Composition, and the Allais Paradox," manuscript, Faculty of Commerce and Business Administration, U. of British Columbia, 1979b.

CHEW, SOO HONG AND WALLER, WILLIAM. "Empirical Tests of Weighted Utility Theory," *J. Math. Psychol.*, Mar. 1986, 30(1), pp. 55–72.

CRAWFORD, VINCENT. "Equilibrium Without Independence," *J. Econ. Theory*, forthcoming 1989.

DAVIDSON, DONALD; MCKINSEY, J. C. C. AND SUPPES, PATRICK. "Outlines of a Formal Theory of Value, I," *Philosophy of Sci.*, Apr. 1955, 22(2), pp. 140–60.

DEBREU, GERARD. *Theory of value.* New Haven, CT: Yale U. Press, 1959.

DEGROOT, MORRIS. *Optimal statistical decisions.* NY: McGraw-Hill, 1970.

DEKEL, EDDIE. "An Axiomatic Characterization of Preferences Under Uncertainty: Weakening the Independence Axiom," *J. Econ. Theory*, Dec. 1986, 40(2), pp. 304–18.

———. "Asset Demands Without the Independence Axiom," *Econometrica*, Jan. 1989, 57(1), pp. 163–69.

DIAMOND, PETER. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: Comment," *J. Polit. Econ.*, Oct. 1967, 75(5), pp. 765–66.

DONALDSON, JOHN AND SELDEN, LARRY. "Arrow-Debreu Preferences and the Reopening of Contingent Claims Markets," *Econ. Letters*, 1981, 8(3), pp. 209–16.

DRÈZE, JACQUES. "Axiomatic Theories of Choice, Cardinal Utility and Subjective Probability," in *Allocation under uncertainty: Equilibrium and optimality.* Ed.: JACQUES DRÈZE. London: Macmillan, 1974, pp. 3–23. (Reprinted in *Essays on economic decisions under uncertainty.* Ed.: JACQUES DRÈZE. Cambridge: Cambridge U. Press, 1987.)

DRÈZE, JACQUES AND MODIGLIANI, FRANCO. "Consumption Decisions Under Uncertainty," *J. Econ. Theory*, Dec. 1972, 5(3), pp. 308–35. (Reprinted in *Essays on economic decisions under uncertainty.* Ed.: JACQUES DRÈZE. Cambridge: Cambridge U. Press, 1987)

EDWARDS, WARD. "The Prediction of Decisions Among Bets," *J. Experimental Psychology*, Sept. 1955, 50(3), pp. 201–14.

———. "Subjective Probabilities Inferred from Decisions," *Psychological Review*, Mar. 1962, 69(2), pp. 109–35.

ELLSBERG, DANIEL. "Classic and Current Notions of 'Measurable Utility,'" *Econ. J.*, Sept. 1954, 64(255), pp. 528–56.

———. "Risk, Ambiguity, and the Savage Axioms," *Quart. J. Econ.*, Nov. 1961, 75(4), pp. 643–69.

EPSTEIN, LARRY. "Decision Making and the Temporal Resolution of Uncertainty," *Int. Econ. Rev.*, June 1980, 21(2), pp. 269–83.

EPSTEIN, LARRY AND ZIN, STANLEY. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, July 1989, 57(4), pp. 937–69.

DE FINETTI, BRUNO. "La Prévision: Ses Lois Logiques, Ses Sources Subjectives," *Annales de l'Institut Henri Poincaré*, 1937, 7(1), pp. 1–68. (English translation: "Foresight: Its Logical Laws, Its Subjective Sources" in KYBURG AND SMOKLER 1980)

FISHBURN, PETER. *The foundations of expected utility.* Dordrecht, Holland: D. Reidel Publishing Co., 1982a.

———. "Nontransitive Measurable Utility," *J. Math. Psychol.*, Aug. 1982b, 26(1), pp. 31–67.

———. "Transitive Measurable Utility," *J. Econ. Theory*, Dec. 1983, 31(2), pp. 293–317.

———. "SSB Utility Theory: An Economic Perspective," *Math. Soc. Sci.*, Aug. 1984a, 8(1), pp. 63–94.

———. "SSB Utility Theory and Decision-making under Uncertainty," *Math. Soc. Sci.*, Dec. 1984b, 8(30), pp. 253–85.

———. "Equity Axioms for Public Risk," *Operations Research*, July/Aug. 1984c, 32(4), pp. 901–08.

———. "Dominance in SSB Utility Theory," *J. Econ. Theory*, Oct. 1984d, 34(1), pp. 130–48.

———. *Nonlinear preference and utility theory.* Baltimore: Johns Hopkins U. Press, 1988.

FISHER, FRANKLIN AND ROTHENBERG, JEROME. "How Income Ought to Be Distributed: Paradox Lost," *J. Polit. Econ.*, Apr. 1961, 69(2), pp. 162–80.

———. "How Income Ought to Be Distributed: Par-

adox Enow," *J. Polit. Econ.*, Feb. 1962, *70*(1), pp. 88–93.

FREEDMAN, DAVID AND PURVES, ROGER. "Bayes' Method for Bookies," *Ann. Math. Statistics*, Aug. 1969, *40*(4), pp. 1177–86.

GÄRDENFORS, PETER AND SAHLIN, NILS-ERIC, eds. *Decision, probability, and utility: Selected readings.* Cambridge: Cambridge U. Press, 1988.

GRANDMONT, JEAN-MICHEL. "Continuity Properties of a von Neumann–Morgenstern Utility," *J. Econ. Theory*, Feb. 1972, *4*(1), pp. 45–57.

GREEN, JERRY. "'Making Book Against Oneself,' The Independence Axiom and Nonlinear Utility Theory," *Quart J. Econ.*, Nov. 1987, *102*(4), pp. 785–96.

GREEN, JERRY AND JULLIEN, BRUNO. "Ordinal Independence in Non-Linear Utility Theory," *J. of Risk and Uncertainty*, Dec. 1988, *1*(4), pp. 355–87.

HAGEN, OLE. "Towards a Positive Theory of Preferences Under Risk," in ALLAIS AND HAGEN 1979, pp. 271–302.

HAMMOND, PETER. "Ex-post Optimality as a Dynamically Consistent Objective for Collective Choice Under Uncertainty," in *Social choice and welfare.* Eds.: PRASANTA PATTANAIK AND MAURICE SALLES. Amsterdam: North-Holland Publishing Co., 1983, pp. 175–205.

———. "Consequentialism and the Independence Axiom," in *Risk, decision, and rationality.* Ed.: BERTRAND MUNIER. Dordrecht, Holland: D. Reidel Publishing Co., 1988a, pp. 503–15.

———. "Consequentialist Foundations for Expected Utility," *Theory and Decision*, July 1988b, 25(1), pp. 25–78.

———. "Orderly Decision Theory: A Comment on Professor Seidenfeld," *Econ. and Philosophy*, Oct. 1988c, *4*(2), pp. 292–97.

———. "Consistent Plans, Consequentialism, and Expected Utility," *Econometrica*, forthcoming 1989, 57(6).

HARSANYI, JOHN. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *J. Polit. Econ.*, Aug. 1955, *63*(4), pp. 309–21. (Reprinted in JOHN HARSANYI, *Essays on ethics, social behavior and scientific exploration.* Dordrecht, Holland: D. Reidel Publishing Co., 1976)

———. "Nonlinear Social Welfare Functions: Do Welfare Economists Have a Special Exemption from Bayesian Rationality?" *Theory and Decision*, Aug. 1975, *6*(3), pp. 311–32. (Reprinted in JOHN HARSANYI, *Essays on ethics, social behavior and scientific exploration.* Dordrecht, Holland: D. Reidel Publishing Co., 1976)

———. "Bayesian Decision Theory and Utilitarian Ethics," *Amer. Econ. Rev.*, May 1978, *68*(2), pp. 223–28.

HAX, ARNOLDO AND WIIG, KARL. "The Use of Decision Analysis in Capital Investment Problems," in *Conflicting objectives in decisions.* Eds.: DAVID BELL, RALPH KEENEY, AND HOWARD RAIFFA. NY: John Wiley and Sons, 1977, pp. 277–94.

HAZEN, GORDON. "Does Rolling Back Decision Trees Really Require the Independence Axiom?" *Manage. Sci.*, June 1987, *33*(6), pp. 807–09.

HELLER, WALTER; STARR, ROSS AND STARRETT, DAVID, eds. *Social choice and public decision making: Essays in honor of Kenneth J. Arrow*, Vol. 1. Cambridge: Cambridge U. Press, 1986.

HERSTEIN, ISRAEL AND MILNOR, JOHN. "An Axiomatic Approach to Measurable Utility," *Econometrica*, Apr. 1953, *21*(2), pp. 291–97.

HEY, JOHN. "The Economics of Optimism and Pessimism: A Definition and Some Applications," *Kyklos*, 1984, *37*(2), pp. 181–205.

HILTON, RONALD. "Failure of Blackwell's Theorem Under Machina's Generalization of Expected-Utility Analysis Without the Independence Axiom," *J. of Economic Behavior and Organization*, forthcoming 1989.

JEFFREY, RICHARD. "Biting the Bayesian Bullet: Zeckhauser's Problem," *Theory and Decision*, Sept. 1988, *25*(2), pp. 117–22.

JOHNSEN, THORE AND DONALDSON, JOHN. "The Structure of Intertemporal Preferences Under Uncertainty and Time Consistent Plans," *Econometrica*, Nov. 1985, *53*(6), pp. 1451–58.

KAHNEMAN, DANIEL AND TVERSKY, AMOS. "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, Mar. 1979, *47*(2), pp. 263–91. (Reprinted in GÄRDENFORS AND SAHLIN 1988.)

KARMARKAR, UDAY. "Subjectively Weighted Utility: A Descriptive Extension of the Expected Utility Model," *Organizational Behavior and Human Performance*, Feb. 1978, *21*(1), pp. 61–72.

———. "Subjectively Weighted Utility and the Allais Paradox," *Organizational Behavior and Human Performance*, Aug. 1979, *24*(1), pp. 67–72.

KARNI, EDI. "Generalized Expected Utility Analysis of Risk Aversion with State-Dependent Preferences," *Int. Econ. Rev.*, Feb. 1987, *28*(1), pp. 229–40.

KARNI, EDI AND SAFRA, ZVI. "Vickrey Auctions in the Theory of Expected Utility with Rank-Dependent Probabilities," *Econ. Letters*, 1986, *20*(1), pp. 15–18.

———. "'Preference Reversal' and the Observability of Preferences by Experimental Methods," *Econometrica*, May 1987, *55*(3), pp. 675–85.

———. "Behavioral Consistency in Sequential Decisions," manuscript, Dept. of Political Economy, Johns Hopkins U. 1988a.

———. "Behaviorally Consistent Optimal Stopping Rules," manuscript, Dept. of Economics, Tel-Aviv U., 1988b.

———. "Ascending Bid Auctions with Behaviorally Consistent Bidders," manuscript, Dept. of Economics, Tel-Aviv U., 1988c.

KEASEY, KEVIN. "Regret Theory and Information: A Note," *Econ. J.*, Sept. 1984, *94*(375), pp. 645–48.

KEENEY, RALPH. "Equity and Public Risk," *Operations Research*, May/June, 1980, *28*(3), pp. 527–34.

KEENEY, RALPH AND WINKLER, ROBERT. "Evaluating Decision Strategies for Equity of Public Risks," *Operations Research*, Sept./Oct. 1985, *33*(5), pp. 955–70.

KEMENY, JOHN. "Fair Bets and Inductive Probabili-

ties," *J. Symbolic Logic,* Sept. 1955, 20(3), pp. 263–73.

KREPS, DAVID. *Notes on the theory of choice.* Boulder, CO: Westview Press, 1988.

———. "Static Choice in the Presence of Unforeseen Contingencies," manuscript, Graduate School of Business, Stanford U., 1989.

KREPS, DAVID AND PORTEUS, EVAN. "Temporal von Neumann–Morgenstern and Induced Preferences," *J. Econ. Theory,* Feb. 1979, 20(1), pp. 81–109.

KYBURG, HENRY AND SMOKLER, HOWARD, eds. *Studies in subjective probability,* 2nd ed. Huntington, NY: Robert E. Krieger Publishing Co., 1980.

LAVALLE, IRVING AND FISHBURN, PETER. "Equivalent Decision Trees and Their Associated Strategy Sets," *Theory and Decision,* July 1987, 23(1), pp. 37–63.

LAVALLE, IRVING AND WAPMAN, KENNETH. "Rolling Back Decision Trees Requires the Independence Axiom!" *Management Science,* Mar. 1986, 32(3), pp. 382–85.

LEHMAN, R. SHERMAN. "On Confirmation and Rational Betting," *J. Symbolic Logic,* Sept. 1955, 20(3), pp. 251–62.

LEVI, ISAAC. "The Paradoxes of Allais and Ellsberg," *Economics and Philosophy,* Apr. 1986, 2(1), pp. 23–53.

LOOMES, GRAHAM AND SUGDEN, ROBERT. "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *Econ. J.,* Dec. 1982, 92(368), pp.805–24.

———. "The Importance of What Might Have Been," in *Progress in utility and risk theory.* Eds.: OLE HAGEN AND FRED WENSTØP. Dordrecht, Holland: D. Reidel Publishing Co., 1984a, pp. 219–35.

———. "Regret Theory and Information: A Reply," *Econ. J.,* Sept. 1984b, 94(375), pp. 649–50.

———. "Disappointment and Dynamic Consistency in Choice Under Uncertainty," *Rev. Econ. Stud.,* Apr. 1986, 53(2), pp. 271–82.

———. "Some Implications of a More General Form of Regret Theory," *J. Econ. Theory,* Apr. 1987, 41(2), pp. 270–87.

LUCE, R. DUNCAN AND RAIFFA, HOWARD. *Games and decisions: Introduction and critical survey.* NY: John Wiley and Sons, 1957.

LUCE, R. DUNCAN AND SUPPES, PATRICK. "Preference, Utility and Subjective Probability," in *Handbook of Mathematical Psychology,* Vol. III. Eds.: R. DUNCAN LUCE, ROBERT BUSH, AND EUGENE GALANTER. NY: John Wiley and Sons, 1965, pp. 249–410.

MACCRIMMON, KENNETH. "Descriptive and Normative Implications of the Decision-Theory Postulates," in *Risk and uncertainty: Proceedings of a conference held by the International Economic Association.* Eds.: KARL BORCH AND JAN MOSSIN. London: Macmillan and Co., 1968, pp.3–23.

MACCRIMMON, KENNETH AND LARSSON, STIG. "Utility Theory: Axioms Versus 'Paradoxes,'" in ALLAIS AND HAGEN 1979, pp. 333–409.

MACHINA, MARK. "'Rational' Decision Making Versus 'Rational' Decision Modelling? A Review of Allais and Hagen (eds.), *Expected Utility Hypotheses and the Allais Paradox,*" *J. of Mathematical Psychology,* Oct. 1981, 24(2), pp. 163–75.

———. "'Expected Utility' Analysis without the Independence Axiom," *Econometrica,* Mar. 1982a, 50(2), pp. 277–323.

———. "A Stronger Characterization of Declining Risk Aversion," *Econometrica,* July 1982b, 50(4), pp. 1069–79.

———. "Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom," in *Foundations of utility and risk theory with applications.* Eds.: BERNT STIGUM AND FRED WENSTØP. Dordrecht, Holland: D. Reidel Publishing Co., 1983a; pp. 263–95. (Reprinted in GÄRDENFORS AND SAHLIN 1988)

———. "The Economic Theory of Individual Behavior Toward Risk: Theory, Evidence and New Directions," manuscript, Institute for Mathematical Studies in the Social Sciences Technical Report, Stanford U., 1983b.

———. "Temporal Risk and the Nature of Induced Preferences," *J. Econ. Theory,* Aug. 1984, 33(2), pp. 199–231.

———. "Choice Under Uncertainty: Problems Solved and Unsolved," *J. Econ. Perspectives,* Summer 1987, 1(1), pp. 121–54.

MARKOWITZ, HARRY. *Portfolio selection: Efficient diversification of investments.* New Haven: Yale U. Press, 1959.

MARSCHAK, JACOB. "Money and the Theory of Assets," *Econometrica,* Oct. 1938, 6(4), pp. 311–25.

———. "Rational Behavior, Uncertain Prospects, and Expected Utility," *Econometrica,* Apr. 1950, 18(2), pp. 111–41. ("Errata," *Econometrica,* July 1950, 18(3), p. 312)

———. "Why 'Should' Statisticians and Businessmen Maximize 'Moral Expectation'?" in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability.* Ed.: JERZY NEYMAN. Berkeley: U. of California Press, 1951, pp. 493–506.

MARSCHAK, THOMAS. "Independence Versus Dominance in Personal Probability Axioms," in *Social choice and public decision making: Essays in honor of Kenneth J. Arrow,* Vol. I. Eds.: WALTER HELLER, ROSS STARR, AND DAVID STARRETT. Cambridge: Cambridge U. Press, 1986, pp. 129–71.

MCCLENNEN, EDWARD. "Sure-thing Doubts," in *Foundations of utility and risk theory with applications.* Eds.: BERNT STIGUM AND FRED WENSTØP. Dordrecht, Holland: D. Reidel Publishing Co., 1983, pp. 117–36. (Reprinted in GÄRDENFORS AND SAHLIN 1988)

———. "1986 Postscript to 'Sure-Thing Doubts'" in GÄRDENFORS and SAHLIN (1988a), pp. 180–82.

———. "Dynamic Choice and Rationality," in *Risk, decision, and rationality.* Ed.: BERTRAND MUNIER. Dordrecht, Holland: D. Reidel Publishing Co., 1988b, pp. 517–36.

———. "Ordering and Independence: A Comment on Professor Seidenfeld," *Economics and Philosophy,* Oct. 1988c, 4(2), pp. 298–308.

_____. *Rationality and dynamic choice: Foundational explorations*. Cambridge: Cambridge U. Press, 1989.

MORRISON, DONALD. "On the Consistency of Preferences in Allais' Paradox," *Behavioral Science*, Sept. 1967, *12*(5), pp. 373–83.

MOSKOWITZ, HERBERT. "Effects of Problem Representation and Feedback on Rational Behavior in Allais and Morlat-Type Problems," *Decision Sciences*, Apr. 1974, *5*(2), pp. 225–42.

MOSSIN, JAN. "A Note on Uncertainty and Preferences in a Temporal Context," *Amer. Econ. Rev.*, Mar. 1969, *59*(1), pp. 172–74.

QUIGGIN, JOHN. "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, Dec. 1982, *3*(4), pp. 323–43.

QUIRK, JAMES AND SAPOSNIK, RUBIN. "Admissibility and Measurable Utility Functions," *Rev. Econ. Stud.*, Feb. 1962, *29*(1), pp. 140–46.

RAIFFA, HOWARD. "Risk, Ambiguity, and the Savage Axioms: Comment," *Quart. J. Econ.*, Nov. 1961, *75*(4), pp. 690–94.

_____. *Decision analysis: Introductory lectures on choices under uncertainty*. Reading, MA: Addison-Wesley, 1968.

RAMSEY, FRANK. "Truth and Probability," in *Foundations of mathematics and other logical essays*. London: K. Paul, Trench, Trubner and Co., 1926. (NY: Harcourt, Brace and Co., 1931.) (Reprinted in GÄRDENFORS AND SAHLIN, 1988, pp. 19–47; and in KYBURG AND SMOKLER 1980)

RÖELL, AILSA. "Risk Aversion in Quiggin and Yaari's Rank-Order Model of Choice Under Uncertainty," *Econ. J.*, 1987, *97*(Supplement), pp. 143–59.

ROSSMAN, MICHAEL AND SELDEN, LARRY. "Notes on Indirect Cardinal Utility and Intertemporal Choice Under Uncertainty," manuscript, Graduate School of Business, Columbia U., 1979.

RUBINSTEIN, ARIEL. "Similarity and Decision Making Under Risk (Is There a Utility Theory Resolution to the Allais Paradox?)," *J. Econ. Theory*, Oct. 1988, *46*(1), pp. 145–53.

SAMUELSON, PAUL. "Probability and the Attempts to Measure Utility," *Economic Review*, July 1950, *1*(3), pp. 167–73. (Reprinted in STIGLITZ 1966)

_____. "Utility, Preference, and Probability (Abstract of a paper given at the conference, "Les fondements et applications de la théorie du risque en econométrie," Paris), 1952a. (Reprinted in STIGLITZ 1966)

_____. "Probability, Utility, and the Independence Axiom," *Econometrica*, Oct. 1952b, *20*(4), pp. 670–78. (Reprinted in STIGLITZ 1966)

_____. *Foundations of economic analysis*, Enlarged ed. Cambridge: Harvard U. Press, 1983.

_____. "How a Certain 'Internal Consistency' Entails the Expected Utility Dogma," *J. Risk and Uncertainty*, Dec. 1988, *1*(4), pp. 389–93.

SAVAGE, LEONARD. *Foundations of statistics*. NY: John Wiley and Sons, 1954. Revised and enlarged ed., NY: Dover Publications, 1972.

SCHICK, FREDERIC. "Dutch Bookies and Money Pumps," *J. of Philosophy*, Feb. 1986, *83*(2), pp. 112–19.

SCHLAIFER, ROBERT. *Analysis of decisions under uncertainty*. NY: McGraw-Hill, 1969.

SEGAL, UZI. "Nonlinear Decision Weights with the Independence Axiom," manuscript, Dept. of Economics, U. of California, Los Angeles, 1984.

_____. "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *Int. Econ. Rev.*, Feb. 1987, *28*(1), pp. 175–202.

SEIDENFELD, TEDDY. "Decision Theory Without 'Independence' or Without 'Ordering': What Is the Difference?" *Economics and Philosophy*, Oct. 1988a, *4*(2), pp. 267–90.

_____. "Rejoinder," *Economics and Philosophy*, Oct. 1988b, *4*(2), pp. 309–15.

SELDEN, LARRY. "A New Representation of Preferences over 'Certain × Uncertain' Consumption Pairs: The 'Ordinal Certainty Equivalent' Hypothesis," *Econometrica*, Sept. 1978, *46*(5), pp. 1045–60.

_____. "An OCE Analysis of the Effect of Uncertainty on Saving Under Risk Preference Independence," *Rev. Econ. Stud.*, Jan. 1979, *46*(1), pp. 73–82.

SELTEN, REINHARD. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International J. of Game Theory*, 1975, *4*(1/2), pp. 25–55.

SEN, AMARTYA. "Rationality and Uncertainty," *Theory and Decision*, Mar. 1985, *18*(2), pp. 109–27. (Reprinted in *Recent developments in the foundations of utility and risk theory*. Eds.: LUCIANO DABONI, ALDO MONTESANO, AND MARJI LINES. Dordrecht, Holland: D. Reidel Publishing Co., 1986, pp. 3–25.)

SHAFER, GLENN. "Savage Revisited," *Statistical Science*, Nov. 1986, *1*(4), pp. 463–501. (Reprinted in BELL, RAIFFA, AND TVERSKY 1988)

SHIMONY, ABNER. "Coherence and the Axioms of Confirmation," *J. of Symbolic Logic*, Mar. 1955, *20*(1), pp. 1–28.

SLOVIC, PAUL AND TVERSKY, AMOS. "Who Accepts Savage's Axiom?" *Behavioral Science*, Nov. 1974, *19*(6), pp. 368–73.

SPENCE, MICHAEL AND ZECKHAUSER, RICHARD. "The Effect of the Timing of Consumption Decisions and the Resolution of Lotteries on the Choice of Lotteries," *Econometrica*, Mar. 1972, *40*(2), pp. 401–03.

STIGLITZ, JOSEPH, ed. *The collected scientific papers of Paul A. Samuelson*, Vol. I. Cambridge: MIT Press, 1966.

STIGUM, BERNT AND WENSTØP, FRED, eds. *Foundations of utility and risk theory with applications*. Dordrecht, Holland: D. Reidel Publishing Co., 1983.

STROTZ, ROBERT. "How Income Ought to be Distributed: A Paradox in Distributive Ethics," *J. Polit. Econ.*, June 1958, *66*(3), pp. 189–205.

_____. "How Income Ought to Be Distributed: Paradox Regained," *J. Polit. Econ.*, June 1961, *69*(3), pp. 271–78.

SUGDEN, ROBERT. "Regret, Recrimination and Rationality," *Theory and Decision*, July 1985, *19*(1), pp. 77–99. (Reprinted in *Recent developments in the foundations of utility and risk theory*. Eds.:

Luciano Daboni, Aldo Montesano, and Marji Lines. Dordrecht, Holland: D. Reidel Publishing Co., 1986)

———. "New Developments in the Theory of Choice Under Uncertainty," *Bull. Econ. Res.*, Jan. 1986, 38(1), pp. 1–24. (Reprinted in *Surveys in the economics of uncertainty*. Eds.: John Hey and Peter Lambert. Oxford: Basil Blackwell Ltd., 1987)

Thaler, Richard and Shefrin, Harold. "An Economic Theory of Self-Control," *J. Polit. Econ.*, Apr. 1981, 89(2), pp. 392–406.

Tversky, Amos. "Utility Theory and Additivity Analysis of Risky Choices," *J. of Experimental Psychology*, Sept. 1967a, 75(1), pp. 27–36.

———. "Additivity, Utility and Subjective Probability," *J. of Mathematical Psychology*, June 1967b, 4(2), pp. 175–201.

———. "A Critique of Expected Utility Theory: Descriptive and Normative Considerations," *Erkenntnis*, 1975, 9(2), pp. 163–73.

Tversky, Amos and Kahneman, Daniel. "Rational Choice and the Framing of Decisions," *J. of Business*, Oct. 1986, 59(4, Part 2), pp. S251–78. (Reprinted in *Rational choice: The contrast between psychology and economics*. Eds.: Robin Hogarth and Melvin Reder. Chicago: U. of Chicago Press, 1987)

Wagner, Harvey. *Principles of operations research*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1985.

Wakker, Peter. "Nonexpected Utility as Aversion of Information," *J. of Behavioral Decision Making*, July/Sept. 1988, 1(3), pp. 169–75.

Wallsten, Thomas. "Subjectively Expected Utility Theory and Subjects' Probability Estimates: Use of Measurement-Free Techniques," *J. of Experimental Psychology*, Apr. 1971, 88(1), pp. 31–40.

Weber, Martin and Camerer, Colin. "Recent Developments in Modelling Preferences Under Risk," *OR Spektrum*, 1987, 9, pp. 129–51.

Weber, Robert. "The Allais Paradox, Dutch Auctions and Alpha-Utility Theory," manuscript, Kellogg School of Management, Northwestern U., 1982.

Weller, Paul. "Consistent Intertemporal Decision-making Under Uncertainty," *Rev. Econ. Stud.*, June, 1978, 45(2), pp. 263–66.

Wold, Herman. "Ordinal Preferences or Cardinal Utility?" *Econometrica*, Oct. 1952, 20(4), pp. 661–63, 664.

Yaari, Menahem. "On the Role of 'Dutch Books' in the Theory of Choice Under Risk," Nancy L. Schwartz Memorial Lecture, J. L. Kellogg Graduate School of Management, Northwestern U., 1985.

———. "The Dual Theory of Choice Under Risk," *Econometrica*, Jan. 1987, 55(1), pp. 95–115.