MAYA BAR-HILLEL AND AVISHAI MARGALIT

HOW VICIOUS ARE CYCLES OF INTRANSITIVE CHOICE? *

ABSTRACT. Transitivity is a compelling requirement of rational choice, and a transitivity axiom is included in all classical theories of both individual and group choice. Nonetheless, choice contexts exist in which choice might well be systematically intransitive. Moreover, this can occur even when the context is transparent, and the decision maker is reflective. The present paper catalogues such choice contexts, dividing them roughly into the following classes:

  1. Contexts where the intransitivity results from the employment of a choice rule which is justified on ethical or moral grounds (typically, choice by or on behalf of a group).

  2. Contexts where the intransitivity results from the employment of a choice rule that is justified on economic or pragmatic grounds (typically, multi-attribute choice).

  3. Contexts where the choice is intrinsically comparative, namely, where the utility from any chosen alternative depends intrinsically on the rejected alternative(s) as well (typically, certain competitive contexts).

  In the latter, independence from irrelevant alternatives may be violated, as well as transitivity. However, the classical money-pump argument against intransitive choice cycles is inapplicable to these contexts. We conclude that the requirement for transitivity, though powerful, is not always overriding.

1. TRANSITIVITY AND RATIONAL CHOICE

Standard theories of choice almost always include a transitivity axiom (however, see, e.g. Fishburn, 1982, 1984). This is motivated not only by mathematical convenience and elegance (Bell *et al.*, 1983), but primarily by the fact that transitivity is a compelling desideratum of choice and preference. Indeed, transitivity seems to be *constitutive* of rational choice. When queried, most people want their choices to be transitive, believe that they are, and are disturbed if presented with evidence that they aren't (e.g. Tversky, 1969). Moreover, "when one is made aware of intransitivities... he is willing to admit inconsistency and to realign his responses to yield to transitive ordering" (Luce and Raiffa, 1957, p. 25).[1] This being so, it is perhaps gratuitous to seek justification for transitivity. Perhaps no reason for being transitive is as intuitively compelling as the very requirement for transitivity itself. Nonetheless, such reasons exist, in the form of the penalties incurred by intransitive choice patterns: their

non-extendability, their *exploitability as a money pump*, and their *vulner-ability to agenda manipulation*.

To spell these out, consider a set of alternatives and a binary preference relationship defined on it. Unless the preference relationship is transitive, then even if A is preferred to B and B is preferred to C, it can not be inferred that A is preferred to C. In other words, the preference order is *not extendable* from the pairs (A,B) and (B,C) to the pair (A,C). Moreover, the relationship does not extend from pairs to larger sets of alternatives. For example, a set of three alternatives that is cyclically ordered does not include a "best" (i.e., maximally preferred) alternative, since every alternative is dominated by some other one.

Another consequence of intransitivity is that if the choice procedure is based on successive pair comparisons (see, e.g., Dummet, 1984), it will never terminate, even when the number of alternatives is finite. This characteristic, irritating enough in itself, can be exploited to turn the decision maker into a *money pump*, the argument being as follows. Suppose one prefers A to B, B to C, and C to A. Usually when one prefers A to B, one is willing to pay some amount, however small, to replace B by A. Similarly, one would be willing to pay to replace A by C, and C by B. Thus, a decision maker who started out with B could be seduced into paying for moving away from it, then later for returning back to it, thus finding himself where he was to begin with, but poorer of pocket. Moreover, such a cycle could, conceivably, be repeated.

Stop rules that are designed to guarantee a choice in a finite number of comparisons can extricate one from this loop, but at the price of possibly arbitrary choice. For example, if an alternative is eliminated as soon as another's preferability over it has been evinced, then which alternative will ultimately survive can be completely determined by *setting the agenda* (i.e., the order in which pairs of alternatives are considered) suitably. Thus, to assure the choice of B from the set $\{A,B,C\}$, say, the agenda need only present the pair (A,C) first.

Wherever these penalties attend intransitive choices (which, as we shall later show, they do not always do), they are at least a nuisance. This in itself is not sufficient, however, to establish that intransitive choice is necessarily irrational or wrong. In the present paper, we hope to show that in some settings, one wouldn't be embarrassed by intransitivities in one's choices even if one were perfectly aware of them, and even upon

extensive reflection. In some of these settings, the intrasitivity of choice is impelled by the very way that the alternatives interact among themselves. In other settings, the penalties for removing an intransitivity may outweigh the penalties of adhering to it. We do not wish to deny the force of transitivity as a normative requirement, merely to deny that this requirement is always normatively overriding.

Our examination of choice situations in which intransitive choice cycles are likely to be encountered will begin with social choice, where the possibility of intransitivity has been most extensively acknowledged and discussed. We shall then proceed to a series of examples that show that "reasonable intransitivity" is by no means confined to this context, and can well arise in individual choice as well.

2. INTRANSITIVITY IN SOCIAL CHOICE

In social choice, intransitivity has always been regarded somewhat differently than in individual choice. Since the 18th century, it has been recognized that the requirement of transitivity is inconsistent with such a highly attractive and eminently reasonable rule for combining individual preferences as the majority rule. This observation has been termed *the Condorcet paradox*.

Although the majority rule, since it is not immune to intransitivity, is susceptible to the penalties outlined earlier, in the context of social choice it is not customary to characterize it as "irrational". Indeed, in a paper entitled "The irrationality of transitivity in social choice", Fishburn (1970) proposed that it is transitivity that "is unreasonable and untenable as a general desideratum for social choice functions" (p. 119), precisely because it is not compatible with the majority rule.

3. BETWEEN GROUP AND INDIVIDUAL CHOICE

A discussion of the reasons for the marked difference in the way that transitivity is regarded in social choice versus individual choice would take us beyond the scope of the present paper. For present purposes, we only note that group choice and individual choice seem to invoke different intuitions. However, the attempt to distinguish between group choice and individual choice is complicated by the fact that some choices cannot be

crisply classified into one or the other. Some choices seem to lie in between, in that they share certain properties with the one, and others with the other. Indeed, the subjective acceptability of the majority rule – and hence of intransitivity – depends to some extent on the similarity between the kind of choice problems in which the rule might be employed and the problem of social choice. To hone our intuitions, we shall consider some examples.

3a. *The case of the benevolent dictator: Individual choice on behalf of a group.* Sometimes, an individual decision maker may be making a decision on behalf of a group of others, with no self-interest, or irrespective of self-interest. A case in point would be a benevolent dictator, who wishes to impose on his subjects the social program they would have, if enabled, chosen collectively themselves. The ethics of the case would seem to be those of group decision making. In other words, the dictator would seem to want to *mimic* the group. Therefore, if and when the majority rule would be justified for a group, it would be justified for the altruistic individual choosing on behalf of that group. This argues for a sort of *vicarious appeal* of the majority rule for this case, which is not altered by the fact that the deciding agency is an individual.

3b. *"When doctors disagree, who will decide?": Group choice on behalf of an individual.* Just as an individual may be called upon to make a decision on behalf of a group, so a group may be called upon to make a decision on behalf of an individual. A case in point is when a panel of doctors attempts to choose a medical procedure for some patient with an eye solely to that patient's welfare, or when a board of directors sets company policy with an eye only to improving the company's finances. Since the group in these cases is but a proxy for an individual, there is just one utility function that needs to be maximized. Hence, in these cases the choices ought, perhaps, to conform to the desiderata of individual choice, although the decision procedure is inevitably that of group choice.

3c. *Choice informed only by the opinions of others.* The appeal of the majority rule is not limited only to decisions made by, or on behalf of, groups. In some cases, the decision maker may find that all he has to

inform his own choice is knowledge of the (pairwise) choices others would have made. Although the ethics of group choice are largely irrelevant here, the majority rule might still seem attractive, in part, perhaps, because of the absence of an appealing alternative. Consider the following example.

Example 1: You want to buy yourself a personal computer. You are a novice in computer matters, and totally ignorant of the personal computer market. You decide to consult some prestigious consumer-guidance publications. Although ethically speaking this is not a group-choice situation, strategically speaking it is, since the problem of integrating the various recommendations is a non-trivial one, unless they are unanimous. You may have complete freedom to combine them whichever way you wish, but you may wish to combine them in a non-arbitrary, justifiable, reasonable manner. To abide by the majority recommendation certainly suggests itself as one possibility.

If you are lucky, the majority rule may point out one clear-cut favorite. But it may also entrap you in an intransitive cycle. It is possible that a majority of the publications you consult found model A superior to B, a majority found model B superior to C, and yet a majority also found C superior to A (the majorities would have to consist of different subsets of the publications for each comparison, but this could happen even if each individual publication orders the models it evaluates quite transitively). Such, for example, is the case in Table I, where the columns correspond to the ranking of A, B and C by the respective publications.

TABLE I

Cell entries give the ranking of the row options by the columns

A	1	3	2
B	2	1	3
C	3	2	1

Besides the majority rule there are, of course, other non-arbitrary, justifiable, reasonable ways of integrating the opinions of others (or any other set of individual transitive ranking). But Arrow's (1951) well-known

Impossibility Theorem proves that none of them are totally trouble-free, in a sense made clear by the theorem. Indeed, in some situations intransitivity may involve less trouble than that engendered by choice procedures alternative to majority voting. In the following sections, we shall discuss some of the merits and attractions of using majority rule, or something akin to it, to integrate individual transitive orderings even when this rule might lead to intransitive cycles.

4. COMPARATIVE CHOICE RULES

Hitherto, there has been little structure to the choice alternatives in our examples, and such structure as there was played no role in the decision rule used to guide choice. We turn now to consider choice between multi-dimensional alternatives (e.g., jobs that are characterized by pay, number of days off, job security, promotion possibilities, etc., houses that are characterized by price, location, size, age, etc.); as that is a context where intransitivities are particularly likely to occur, for reasons we shall indicate.

Multi-dimensional (or multi-attribute) alternatives can be evaluated either independently or comparatively, either separately or jointly. An example of independent evaluation is to attach a price ("cash equivalent"), or some other measure of worth, to each alternative separately, and between any two (or more) alternatives to choose the one that is valued highest. In a comparative evaluation procedure, on the other hand, pairs of alternatives would be directly compared, and the one which comes out ahead in this comparison (whether or not the difference is evaluated quantitatively) would be chosen.

For present purposes, a particularly interesting class of comparison rules is the *non-linear additive difference model*. A comparison rule is said to satisfy an *additive difference* model if the comparison of the multi-attribute alternatives proceeds by quantitatively evaluating the difference on each dimension separately, and then adding these evaluations up. If at least some of the intra-dimensional evaluation functions are not additive (i.e., if $f(a+b) \neq f(a) + f(b)$), then the model is also said to be *non-linear*. Whereas the independent pricing procedure guarantees transitivity, the non-linear additive difference model guarantees intransitivity (see proof and further discussion in Tversky, 1969).

Several authors (e.g., Fishburn, 1984; MacCrimmon and Larson, 1979; May, 1954) have noted that multi-dimensional choice problems are prime candidates for eliciting intransitive individual choice patterns. The additive difference model is often both psychologically and computationally simpler to apply to multi-dimensional objects than independent evaluation (Tversky, 1969). In addition, psychophysical functions (which the intra-dimensional difference evaluations really are) are typically non-linear.[2] This combination, therefore, provides the basic ingredients for intransitivity. Whether an independent pricing procedure or a comparative one is more appropriate, turns out – as we hope to show – to be dependent on the characteristics of the setting in which the choice is made, as well as on those of the decision maker. The misplaced use of a pricing rule, for instance, can disguise an underlying intransitive cycle no less than the misplaced use of a non-linear additive-difference rule can distort an underlying monotonic ranking. We shall see examples of both kinds below.

5. MAJORITY RULE IN MULTI-ATTRIBUTE CHOICE

A certain analogy exists between the attempts of a group to integrate the rankings given by each of its members into an overall ("societal") ranking, and the attempt of an individual to integrate the rankings on each of a number of dimensions into an overall ranking of the alternatives. The analogy is at least formal, and possibly more than that (see discussion in Section 11).

In any case, the majority rule is a possible decision rule in both cases. Indeed, the majority rule, applied to multi-attribute choice, is a special case of a non-linear additive-difference model. The number of dimensions on which one alternative is superior are summed. The non-linearity lies in the fact that any dimension on which one alternative dominates another has the same impact on the final summation, irrespective of how many of the other alternatives may be ranked in between the two.

Example 2: The following example is not a thought experiment but a real one, albeit one that was run rather informally, and elicited only hypothetical choices. May (1954) presented 62 college students with the

choice problem depicted in Table II (which is really just a concretization of Table I). The three alternatives were prospective spouses, and the three dimensions on which each matrimonial candidate was characterized were intelligence, looks, and wealth. Although only 17 of the 62 subjects ordered the candidates cyclically, all 17 of them chose A over B, B over C, and C over A, with none choosing the other possible cycle. May considered this to be evidence that these 17 subjects were basing their choices on the majority rule.

TABLE II

Characterization of three hypothetical marriage candidates (May, 1954)

Attributes		intelligence	looks	wealth
	A	very intelligent	plain looking	well off
Candi-dates	B	intelligent	very good looking	poor
	C	fairly intelligent	good looking	rich

To be sure, 45 other subjects managed to order the three candidates from most to least desirable without cyclicity. Apparently, they either consider-ed one single dimension overridingly important (for instance, 21 of the 45 transitive orderings were from most to least intelligent), or could trade off advantages and disadvantages across dimensions. Indeed, the difficul-ty in obtaining systematic violations of transitivity in such contexts arises because choice patterns that transparently violate normative desiderata are labile (though see Lindman and Lyons, 1978). People *can* make cross-dimensional tradeoffs, and are more inclined to do so if they realize that they are involved in an intransitive cycle from which they wish to extricate themselves.

Clearly, May's subjects found the majority rule less than over-whelmingly appealing as a rule for multi-attribute choice. Imagine, however, that you have to choose between just two options, ranked on,

say, a dozen dimensions. You consider these dimensions to be of equal importance, you find it hard to make cross-dimensional tradeoffs, and all you have is a rank order, not a cardinal measurement, of the two options on each dimension. You now find that one option ranks higher than the other on eleven of them; would it not then be enormously tempting to conclude that it is the better one? If so, you had better beware, because this same choice rule (i.e., if one option outranks another on at least eleven of twelve dimensions, choose it) can – given enough options (twelve, as a minimum) – bring about an intransitive cycle.

6. OTHER INTRANSITIVE RULES FOR MULTI-ATTRIBUTE CHOICE

6a. *The lexicographic semi-order.* Another choice rule for multi-attribute alternatives that can lead to intransitivities is the lexicographic semi-order (Tversky, 1969). According to this rule, the dimensions are ordered from most to least important, and considered in turn. If two alternatives differ by more than some threshhold value on the considered dimensions, the "better" one is preferred; otherwise, they are compared on the next dimension, etc. This rule is also a special case of a non-linear additive difference rule, obtained when the difference function is a step function – namely, when differences that are smaller than some threshhold value on one dimension are taken to be zero, whereas differences larger than this value are overriding, in the sense that they cannot be compensated for by what goes on in other dimensions. The following example presents a decision problem in which the lexicographic semi-order might be used as a choice rule. It is taken from Tversky (1969).

Example 3: A decision maker wishes to choose one of several candidates for a given job. Candidates are characterized only in terms of their intelligence (given by an IQ score) and previous experience (given by number of years) on a related job. The decision maker believes that intelligence is far more important for successful performance on this job than experience. He also knows that IQ tests are less than perfectly reliable. The decision maker prefers a more intelligent candidate, irrespective of previous experience. Of equally intelligent candidates, however, he prefers the more experienced one. Faced with pairs of candidates,

he adopts the following decision rule: Given two candidates, with different IQs and background experience, choose the one who is more intelligent; but if the two differ by less than three IQ points, consider them – for all practical purposes – equally intelligent, and choose the more experienced one.

This seemingly reasonable decision rule could lead the decision maker to order some candidates cyclically (e.g., if each candidate is represented as a pair of numbers [IQ, years of experience], then $[115,7] \succ [117,3] \succ [120,0] \succ [115,7]$). Nonetheless, in situations like this some people do spontaneously adopt the lexicographic semi-order. Selecting subjects for whom this seemed to be the case, Tversky (1969) was able to predict and obtain systematic violations of transitivity, in precisely the manner indicated by the lexicographic semi-order.

6b. *Favoring common dimensions.* The final example in this section also refers to a rule whose use has been experimentally demonstrated. Slovic and MacPhillamy (1974) required subjects to compare alternatives that were characterized by their scores on only two of three possible dimensions. "The results indicated that dimensions were weighted more heavily in the comparison when they were common than when they were unique" (p. 172), making this rule a non-linear additive-difference rule. How this tendency can lead to intransitive choice cycles is shown in the following example.

Example 4: Consider three students who are characterized, say, by their scores on two highschool subjects. X received an A in English and a B in Math, Y received a B in English and an A in History, and Z received an A in Math and a B in History, as depicted in Table III. The decision maker, who wishes to award a scholarship to the "best" student, is faced with two candidates at a time. If the three school subjects are taken to be equally diagnostic of a student's ability, and our decision maker shares the tendency found by Slovic and MacPhillamy, he will favor X in a comparison with Y (since X scored higher on their common subject, English, which will outweigh Y's advantage on the other, unique, test score). Similarly, he will favor Y in a comparison with Z – but he will also favor Z in a comparison with X, resulting in a cyclical ordering of the students.

TABLE III

Cell entries are the student's test scores

	English	Math	History
X	A	B	—
Y	B	—	A
Z	—	A	B

7. ON THE RATIONALITY OF APPROXIMATION SCHEMES

Tversky (1969) suggested that his subjects used the lexicographic semi-order as some kind of *approximation scheme* that simplified the information processing that is a prerequisite for multi-attribute choice, without their being aware that this rule might introduce intransitivity. Hence, he didn't regard the exhibited intransitivity as irrational.

Indeed, the use of simple comparative rules as *aids* in multi-attribute choice can be likened to the use of certain approximation methods (e.g., rounding) as aids in arithmetic, or to the use of limiting-cases theorems as aids in probability calculations. To ask whether it is rational to let one's choices be guided by such rules is akin to asking whether it is rational to let one's answers to computational problems be guided by some approximation. The "rationality" of such schemes depends on a suitable cost-benefit analysis.[3] To the extent that the cost of evaluating the alternatives independently outweighs the benefits of ordering them transitively, these rules may be instances of "bounded rationality" (Simon, 1957) rather than of irrationality. For a more elaborate discussion of this issue, see Tversky (1969).

8. INTRANSITIVITY AS A REFLECTION OF EXTERNAL CYCLICITY

In the previous three examples we argued that the alternatives in the choice sets could well have been independently evaluated, and that the decision maker evaluated them comparatively only because that was a

cognitively easier thing to do. In contrast, the alternatives in the following examples cannot – rather, *should* not – be independently evaluated, as their worth *vis-à-vis* the relevant decision is inherently dependent 'on the alternative(s) with which they are compared.

Before presenting the examples, we introduce a relationship we shall call *probabilistic prevalence*: a random variable A *probabilistically prevails* (*pp*s) over variable B iff $p(A > B) > 0.5$, namely, iff A has a better chance of exceeding B than B has of exceeding A. *pp* is an intransitive relationship, namely, there exist triads of random variables A, B and C, such that A *pp*s over B, B *pp*s over C, and C *pp*s over A. Consider, for example, the three uniform random variables A, B and C, represented in the rows of Table I, where the columns are taken to represent equiprobable states of the world. Clearly, C exceeds B more often than the reverse (i.e., in two cases out of three), as does B with respect to A and A with respect to C.[4]

Example 5: Consider now the following choice situation. Two players are presented with a pair of alternatives, which can be thought of as "strategies", and one is given first choice. The other then gets the remaining alternative, and the two are played against each other, the winner receiving some fixed predetermined prize.

Suppose the pairs of alternatives are any two rows of Table I, and they are played against each other by drawing at random among the equiprobable columns, the winner being the row containing the higher number for that column. Then a rational player given first choice should choose B if the pair is (A,B), should choose C if the pair is (B,C), and should choose A if the pair is (A,C).

From the perspective of an observer who sees only the choice sets and the choices, this is an intransitive choice pattern. Nonetheless, it is perfectly justified under the circumstances: the probabilistically prevalent alternative wins more often than not, hence it is the rational choice, even though *pp* is not a transitive relationship. To be sure, when B was chosen over A, C was chosen over B, and A was chosen over C, the very context of the choice itself changed. Each one of these three choices was made from a different choice set. However, this is hardly a peculiarity of this example. Indeed, Luce and Raiffa (1957) define transitivity of preferences as follows: "if A is preferred in the paired comparison (A,B) and B is

preferred in the paired comparison (B,C), then A is preferred in the paired comparison (A,C)" (p. 16). This pedantic definition is more careful – and appropriate – than the vaguer "if A is preferred to B, and B is preferred to C, then A is preferred to C", which does not specify the choice set in which the preference is expressed.

Example 5 is a formal one. Real games of this kind are rarely encountered, for the obvious reason that they would be too trivial. Where games do include pairs of "strategies" such that the choice of one would automatically confer an advantage on its player, symmetry between the players is often restored either by allocating the first choice through a lottery or through turn-taking, or by introducing a third strategy and requiring both players to choose simultaneously, or at least in ignorance of their opponent's choice. For example, the choice of a color in chess or of a field in a ball game is regulated by the first possibility, whereas the choice of a strategy in a game such as the children's game Rock, Scissors, Paper, is regulated by the second.[5] Nonetheless, the following is an attempt, due to Blyth (1972), to devise a semi-realistic, yet non-trivial, example.

Example 6: Having constructed three random variables A, B and C such that the probabilities $p(A \leq B)$, $P(B \leq C)$ and $P(C \leq A)$ are all greater than 1/2, Blyth suggested that A, B and C be interpreted as the distributions of "the respective times taken by three runners, A, B, C to run some course. [It follows that] A beats B, B beats C, and C beats A, each with probability exceeding 1/2" (p. 369).[6]

Blyth called this state of affairs "the non-transitivity paradox". This name notwithstanding, there is really no paradox here. The decision maker's preferences are perfectly coherent and well-ordered: he always prefers to bet on the runner more likely to win the race, therefore to win the bet for him. From this perspective, it is an objective truth that A is a better bet than B when these two are the competing runners, B is a better bet than C when these two are the competing runners, and C is a better bet than A when these two are the competing runners. Note that it is meaningless, or fallacious, to assert without qualification that "A is a better bet than B", or "B is a better bet than C"; "better", one should ask, "under what circumstances?" A is a worse bet than B if, for example, the race is between C and a choice of either A or B.

In the circumstances of Example 6, a decision maker exhibiting the above-mentioned cyclic choice pattern is perfectly justified and rational. Indeed, any other choice pattern would be flawed, because the cyclicity is in the external world, and the cyclic choice merely mirrors it.

It is of interest to note that even the penalties we outlined earlier as accruing to intransitive choice largely lose their sting here:

(i) *Money pump*. Most important, the decision maker cannot be turned into a money pump. For suppose it is known that A beats B, B beats C, and C beats A, each with probability exceeding $1/2$. Suppose further that the race our decision maker is betting on is a two-way race between A and B, and he has bet on A. Someone comes by with an offer to allow the decision maker, for a small fee, to switch his bet to C rather than A, since C is the better runner of the two in the sense of having a better chance of beating A than vice versa. The offer would be immediately turned down, since C isn't even participating in the race!

Suppose, however, that the offer is to replace A's competitor, B, by C *and* to allow the decision maker to switch his bet to C. If C's chances of beating A are greater than were A's chances of beating B, this would be an attractive offer. But this is not a money pump possibility, for the decision maker will soon settle on that one of the three possible pairwise races in which the favorite's odds to win are the highest – and there all transactions will terminate, since that consideration orders the pairwise races quite transitively. For example, if $P(A$ beats $B)=0.55$, $P(B$ beats $C)=0.6$, and $P(C$ beats $A)=0.65$, and if the original race was scheduled between A and B, our decision maker might pay something to replace B by C and switch bets, but will then pay no further, as his chances of winning the bet can no longer be improved.

(ii) *Non extendability*. The preference ordering by pp is, indeed, non-extendable. But the reason is not due to intransitivity. In other words, even when three random variables A, B and C *are* transitively ordered by pp, the ordering is still not extendable, in the sense that the pairwise best bet (i.e., best runner) might be the three-way worst one. Blyth terms this property of pp "the pairwise worst-best paradox", though it is more customarily referred to as *non independence of irrelevant alternatives*. It is demonstrated in the following example.

Example 7: Blyth constructed an example of three random variables A, B and C, such that $P(A<B)$, $P(B<C)$ and $P(A<C)$ are all greater than

1/2, yet $P(A = \min A,B,C) < P(B = \min A,B,C) < P(C = \min A,B,C)$, as in Figure 1.[7] Interpreting A, B and C as the distribution of running time of the runners A, B and C, we have an example where "[T]he respective times taken by runners A, B, C to run the course [are such that] A beats B, A beats C, and B beats C, each with probability exceeding 1/2 [and so with respect to pairwise races they are *transitively* ordered]; yet C is most likely (and A least likely) to win a race in which all three run" (pp. 370–371). Since "a race run in heats is a very different contest from a single race in which all runners compete together" (p. 371), non-extendability in this case cannot be blamed on intransitivity, but on the fact that the three-way winning probabilities are not sufficiently constrained by the two-way probabilities.
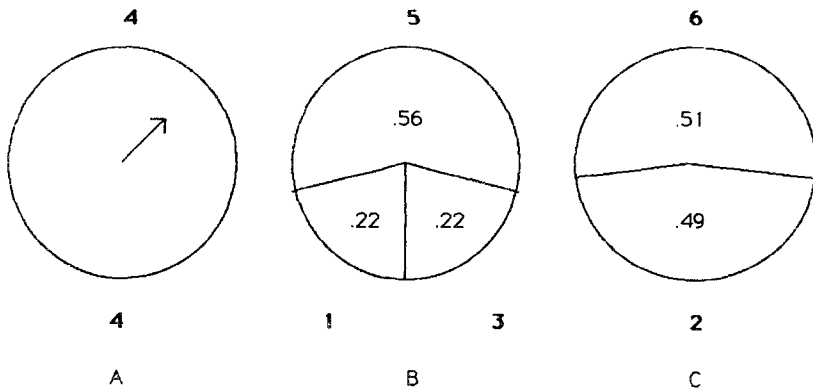
A spinner is turned, and a prize is given to the highest scorer.
$P(A<C)=0.51>0.5$ $P(B<C)=0.51+0.49\times0.22=0.6178>0.5$ $P(A<B)=0.56>0.5$
but
$P(A=\min A,B,C)=0.56\times0.51=0.2856<P(B=\min A,B,C)=0.22+0.22\times0.51=0.3322<$
$P(C=\min A,B,C)=0.49\times0.78=0.3822$

Fig. 1.

(iii) *Agenda setting.* This is the only penalty that still persists and is attributable to the cyclic choice pattern in this example. The "agenda" in this case is simply the racing schedule. Because of the cyclicity, if each runner participates in but one pairwise race, the tournament can be "fixed" in a way that favors one runner or is biased against another.

9. INTRANSITIVITY AS A REFLECTION OF INTERNAL CYCLICITY

In the same way as an intransitive choice cycle might mirror some objective external cyclicity (as it does in Example 6), it might mirror an internal cyclicity. Consider Example 8, which is also taken from Blyth (1972), following an example by Halmos.

Example 8: Suppose, Blyth proposes, that the consumption over times of "apple, blueberry, and cherry pie... provide the individual with random (as they surely would be) amounts of "satisfaction" [A, B and C, respectively], whose distribution he knows to satisfy: [A *pp*s over B, B *pp*s over C, and C *pp*s over A (i.e., apple pie is tastier than blueberry pie more often than not, blueberry pie is tastier than cherry pie more often than not, and cherry pie is tastier than apple pie more often than not)]. Why shouldn't he prefer A to B, and B to C, and C to A, on the grounds that the preferred member of each pair has probability exceeding 1/2 of giving him more satisfaction than the other?" (p. 369).

Since the cyclicity of A, B and C with respect to the *pp* relationship is just as genuine when these random variables are interpreted as internal satisfaction or subjective tastiness as when they are interpreted as physically measured running times, the intransitivity of the choices between pairs of these variables follows necessarily from the attempt to maximize the probability of selecting the best pie on the menu.

Nonetheless, on intuitive grounds, this example appears less convincing than the runners example. Even harder to swallow is the following Example 9, which is a taste-version of the "pairwise worst-best paradox". Yet, if we accept the legitimacy of Blyth's preference criterion (i.e., to maximize the probability of choosing the alternative most likely to be superior to rejected alternatives), we would have no grounds for rejecting the startling Example 9.

Example 9: "Suppose that apple, blueberry and cherry pie provide [the individual] with random amounts [A, B and C, respectively] of "satisfaction", whose distribution he knows to satisfy: [$P(A \geq B)$, $P(B \geq C)$ and $P(A \geq C)$ are all greater than 1/2 – which produces a *transitive* ordering of the three variables – yet $P(A = \max A, B, C) < P(B = \max A, B, C) < P(C = \max A, B, C)$].[8] Offered a choice of A or C in the cafeteria every day at lunch, he always chooses A, reasoning that this is more likely to give him the greater satisfaction. But one day they also

offer B, and he remarks "seeing that you also have B, I'll take C instead of A", reasoning that of the three, C has the best chance of giving him the greatest satisfaction" (p. 371).

How can one deal with these "paradoxes"? One possibility is to deny the legitimacy of Blyth's criterion. In this spirit, Pratt (1972) commented that "many of Blyth's arguments are of a form something like: preference criterion A [e.g., "maximize the probability that your chosen alternative will be better than the rejected one(s)"] contradicts axiom B [e.g., transitivity, in Examples 6 and 8, or independence, in Example 9], therefore axiom B is unacceptable. ... My [Pratt's] conclusion would rather be that preference criterion A is unacceptable, and that the "paradoxes" simply expose the weaknesses in superficially plausible preference criteria" (p. 378).

We find ourselves siding with Blyth's position against Pratt's that people have a "right" to maximize whatever criterion they want to maximize, even if consistent adherence to its maximization leads to intransitive cycles. We question, however, the *psychological* realism of the pie example. The problem lies in accepting the analogy between the gourmet of Examples 8 and 9, and the gambler of Example 6. According to the story told, the gambler's payoff, or utility, depends *exclusively* on who wins the race (i.e., on whether or not the gambler bet on the actual winner). Even its size is not dependent on the probability with which the race is won, or on the difference in times between the winner and the loser. This assumption is plausible enough for a betting situation, but is rather implausible for a consumer of pies, whose utility clearly depends – additionally, if not solely – on the taste experience engendered by the pie that he or she actually consumes.[9] Even if the pleasure in consuming a pie is enhanced by knowing that the rejected pie would not have been as tasty as the selected pie, it still would seem to depend on the taste of the selected pie in and of itself; Indeed, it stands to reason that with respect to pies, one would rather eat a delicious pie, which is not the best one on the tray, than a pretty horrible one, which is the best of the day's lot, just as one would probably rather see a good movie even if it isn't the best in town than a mediocre movie which is. In technical language, the utility of a selected pie depends not only (if at all) on its joint distribution with respect to rejected pies, but on its marginal distribution as well.

With respect to other things, however, the competitive element might play a more prominent role, and perhaps even be overriding. Someone might prefer to be the best-dressed person in a crowd of poor dressers than somewhat better dressed, but surrounded by people dressed yet better than oneself, or to be the brighest member in a mediocre group than a mediocre member of a brilliant group, etc. In strictly competitive situations, such as official competitions where winner takes all (as in the previous betting example), it is, of course, best just to be ahead of the others, wherever that position might be in some absolute terms. But even in a clearly competitive situation such as war, it might be better to lose a conventional but limited war than to win a global nuclear one.

These are psychological, not conceptual, observations. Other practical considerations are the implausibility – though not impossibility – of encountering actual pies (or, for that matter, actual runners) with the kind of joint distributions necessary for the "paradoxes". But it cannot be denied that if a person's psychological makeup is such that the only – or major – consideration for him is to maximize the chances of selecting an alternative better than the one rejected, then he is subject to Blyth's paradoxes. In any event, we argue that Blyth's examples appear paradoxical only because of their psychological implausibility.

As a somewhat more realistic and convincing example of internal cyclicity, we offer Example 10.

Example 10: Let A, B and C in Table IV be gambles whose payoffs, given in dollar amounts in the cells, depend on which of three equiprobable states of the world will obtain. From the perspective of any standard theory of choice, these gambles are identical, because they are identically distributed. Suppose, however, that the gambler is sensitive not only to the amounts gained, but also to the amounts foregone, and in any choice situation, compares what he actually gained to what he might have gained if he had chosen otherwise. Suppose that this sensitivity is subadditive with the difference between what is and what could have been. In particular, two opportunities to be $ 100 ahead make him happier than a single opportunity of being $ 200 ahead, or a single loss of $ 200 is more than compensated for by two separate gains of $ 100 (recall that we are dealing with equiprobable opportunities). Then he would choose A from the set $\{A,B\}$, B from the set $\{B,C\}$, and C from the set $\{A,C\}$.[10]

TABLE IV

Columns are three equiprobable states of
the world

Gambles	States of the world		
	1/3	1/3	1/3
A	$ 100	$ 200	$ 300
B	$ 300	$ 100	$ 200
C	$ 200	$ 300	$ 100

10. INTRANSITIVITY AND REGRET

Examples 8 and 10 implicitly incorporate a notion of regret or rejoicing:
the appeal of some selected alternative is enhanced if, in addition to what
it gives the decision maker in itself, it also minimizes expected regret
(where *regret* is wishing that you had chosen differently) – or, alternative-
ly, maximizes expected rejoicing. This notion has been explicitly incor-
porated into several recent attempts to model certain prevailing patterns
of choice under uncertainty that cannot be dealt with by standard theories
of choice.

Loomes and Sugden (1982) based their attempt on a distinction between
"choiceless" and "modified" utility. The previous is the utility attached
to an outcome if an individual experiences it *"without having chosen it"*
(p. 807, italics there), while modified utility is the utility attached to an
outcome that is experienced "as the result of an act of choice" (p. 808).
If the modified utility is greater than the choiceless utility, the difference
is experienced as *rejoicing*, whereas if it is lesser, the difference is
experienced as *regret*.

Sage and White (1983) based their attempt on the distinction between
"*A* is preferred to *B*" and "selecting *A* and rejecting *B* is preferrable to
selecting *B* and rejecting *A*" (p. 144). The underlying intuition is clearly
the same in both theories, though Loomes and Sugden saw theirs as "An
alternative theory of *rational* choice under uncertainty" (italics ours; title
theirs), whereas Sage and White put forth what they regarded as a
descriptive theory. Motivated by similar intuitions, Bell (1982) saw his

theory as the basis for a *prescriptive* theory. Loosely speaking, a prescriptive theory is addressed to the question: "How can real people – as opposed to imaginary, idealized, super-rational genuises without psyches – [be advised in how to] make choices in a way that does not do violence to their deep cognitive concerns" (Bell *et al.*, 1983, p. 2), and yet "reflect[s] their reflective judgments and beliefs" (p. 22).

In all these attempts, as well as in Fishburn (1982, 1984), the representation of the preference relation by the utility function is of the form: $A \leq B$ iff $u(A,B) \leq 0$, rather than the classical $A \leq B$ iff $u(A) \leq u(B)$. This captures the fundamental fact that in what we might call *inherently comparative choice*, pairs of alternatives interact so that their joint distribution is not the product of their marginal distributions, and the utility of the selected alternative depends on its joint distribution with respect to the rejected alternative, and cannot be reduced to a function of its marginal distribution alone. In such choice contexts intransitivity is not inevitable, but it remains a distinct possibility. Nonetheless, its existence does not always entail that the decision maker can be turned into a money pump.

The following example roughly demonstrates, using a lexicographic semi-order kind of decision rule, how the addition of an inherently comparative consideration such as regret to considerations which are not inherently comparative can lead to intransitivity.

Example 11: Consider the following three eventualities. (A) One day after you purchased 800 shares of a new stock, its value increased by a dollar a share. (B) Your stock portfolio, which you rarely review and have not touched in months, increased in value by $ 900 overnight. (C) One day after you sold 1000 of your 2000 shares of some stock, its value increased by a dollar a share.

Following these eventualities, you are richer by $ 800, $ 900, or $ 1000, respectively. Nevertheless, you may well be the kind of person for whom A is preferred to B, because of the rejoicing that accompanies A's more modest gain; B is preferred to C because of the regret that accompanies C's larger gain; but C is preferred to A, because the considerably larger gain in C more than offsets the regret associated with it.

In such cases, as in the earlier runners case (Example 6), the cyclic choice or preference cannot be exploited to turn the decision maker into a money pump, and for similar reasons.

11. WHEN TRANSITIVITY FAILS

Occasionally, when patterns of choice behavior are observed to violate certain consistency requirements, the violation can be attributed to factors such as changes in taste over time, ad hoc choice criteria, cognitive limitations, emotional (i.e., "irrational") interferences, etc. In the present paper, however, we have focused on choice patterns that violate transitivity as a result of the consistent and deliberate application of a well-formulated choice rule, and do so even in transparent contexts (Tversky and Kahneman, 1986).

Table V presents a classification of choice situations in which intransitivities are likely to be encountered. Since intransitivity only occurs when alternatives are evaluated comparatively, the primary distinction between these situations is whether or not the comparison is indigenous to the situation.

TABLE V

A classification of intransitivity-prone choices

A. Inherently comparative choice – probabilistic prevalence
    1. External contingencies attach a predetermined fixed prize to choice of alternative that "wins" in a pairwise competition (e.g., prize goes to the winner of a race).
    2. Internal experience attaches a cost to choice of alternative that "loses" a pairwise competition (e.g., regret is associated with outcome that is not the best possible).
B. Non-inherently comparative choice – integrating separate orderings
    1. Group choice – Choice procedure selected by ethical or political considerations. Any other procedure, comparative or non-comparative, is also subject to some drawbacks.
    2. Multi-dimension choice – Choice procedure selected by considerations of convenience and simplicity. Non-comparative procedure can also be used, but at an economic cost.

Comparative choice rules are indigenous only to *inherently comparative* choice situations, namely situations in which the utility attached to the choice of an alternative A depends on the alternative which was rejected in favor of A. The inherently comparative nature of a choice can be a reflection either of some externally arranged contingencies, or of the nature of some internal experience. External contingencies are often

beyond the decision maker's control, but internal experience is occasionally open to modification. Although inherently comparative choice situations cannot be exploited to turn the decision maker into a money pump, they nonetheless incur other costs, that are avoidable in choice situations that are not inherently comparative. So the decision maker enjoys certain advantages if he can alter an inherently comparative choice into a non-inherently comparative one. To the extent that a person can train himself, say, to be less swayed by regret considerations or competitiveness, he may find himself better off with respect to the more tangible payoffs that he actually ends up with.

To determine whether or not a choice situation is inherently comparative, one must ask whether the experience of having A does or does not depend on whatever was forgone by opting for A. In the non-inherently comparative choice situations that we presented as examples, a series of transitive (ordinal) rankings of some alternatives had to be combined into a global, overall (ordinal) ranking of these alternatives. The transitive rankings were provided either by independent decision makers (members of some group), or by a single decision maker who ranked the alternatives with respect to several different considerations or dimensions. Intransitive cycles in such non-inherently comparative situations arise only if a comparative choice rule – more specifically, a non-linear additive difference rule – is employed to choose between pairs of alternatives, although it is not indigenous to that situation. For a non-linear additive difference rule to be applicable, the alternatives must present the decision maker with a natural series of meaningful comparisons. In that respect, group choice is analogous to multi-dimensional choice. In both, there are n transitive orderings of the options (given – in the social choice setting – by the n members of the group, and – in the multi-attribute setting – by the ranking of the options on each of the n dimensions). Hence, the task of combining several separate rankings into a single overall ranking, especially if it is ethically or technically difficult to compare the rankings, is – as we noted earlier – one which is particularly susceptible to resulting in intransitive cycles. It is no coincidence that in multi-dimensional choice and group choice alternatives are often evaluated by comparative rules, even though they could well be also assessed independently in a meaningful way.

If the use of a comparative rule in inherently comparative situations

is justified – indeed, *necessitated* – by the facts of the matter (i.e., the nature of the interaction between the alternatives), the use of a comparative rule in our examples of non-comparative situations was justified by procedural considerations. In other words, ethical or pragmatical considerations favored a certain type of choice rule in and of itself, in a way that transcended or overruled, at least to some extent, considerations of its precise suitability to the structure of the choice alternatives at hand. If it is deemed important or desirable to arrive at one's decision by a particular procedure (i.e., if the medium is part of the message, so to speak), the price may be intransitivity.

Note that in the schematic examples depicted in Tables I, III and IV, the choice alternatives are, in some important sense, essentially equivalent. Up to a permutation of the columns, the rows are identical, and the columns have been assumed to represent voters or dimensions worthy of equal weighting, or equiprobable states of the world. The appeal of a comparative choice rule in these examples was not argued for on the basis of this equivalence, and little in the examples depended on this equivalence. But in these examples, an independent evaluation scheme that would end up assigning all the alternatives the same value, would fail to provide grounds for choosing among them. When choice options are recognized as equivalent, one can be said to be *picking* rather than choosing among them (Ullman-Margalit and Morgenbesser, 1977). But choosing arbitrarily, or even by lottery, is often deemed politically or morally unacceptable. Choosing by a comparative rule such as majority vote at least provides grounds, justification, for the choice. Thus, their tie-breaking ability is an additional attraction of comparative choice rules.

12. CHOICE, PREFERENCE, AND DECISION RULES

Preferences are not subject to public observation, though their existence is readily confirmed by introspection. Commonsense, or naive psychology, suggests that people's choices, at least in simple, transparent cases, are dictated by their preferences. However, were people to be in touch with their "true preferences" in every choice situation, there would be no need for the burgeoning discipline of decision analysis, which develops

and implements procedures for elicting or uncovering preferences, and
for guiding and justifying choices, in the aid of decision makers.

The revealed preference approach in economics (Samuelson, 1948)
basically embodies this naive assumption. Nonetheless, it is clearly flaw-
ed. Choice, for example, is extensional (i.e., is not description-depen-
dent), whereas preference is clearly intensional. Jane chose to go see a
movie at the Park rather than the Lane cinema. What is to be inferred
from this choice? That she prefers to see the movie showing at the Park?
That she prefers the premises at the Park cinema? That the screening time
at the Park was more to her liking? Perhaps her choice was determined
by the movie being shown, but she was in error about which cinema house
is screening which movie. Moreover, the actions or alternatives between
which one is to choose are often linked to the outcomes between which
one has a preference in a very complex and circuitous way. The prefer-
ences which a chess player has are overwhelmingly on final game
outcomes (he'd rather win than tie, and tie than lose), but the choices are
between permissible moves. He may need to resort to computational or
analytical aids to determine which move is optimal in terms of his ultimate
goal, but can we infer that he has preferences among moves because he
chooses among them?

Recently, the concept of preference has come under attack from the
study of so-called preference reversals (see, e.g., Slovic and Lichtenstein,
1983). A typical example of the phenomenon of preference reversal is a
pair of (positive binary) gambles, such that decision makers asked to
choose between the two prefer A to B, while at the same time they place
a higher cash-equivalence on B than on A. The pervasiveness, systematici-
ty and predictability of preference reversals have led many to the con-
clusion that preferences, such as they are, are better thought of as being
constructed by choice rather than revealed by it, and that choice proce-
dures are geared towards justification and conflict resolution rather than
towards optimization (Tversky *et al.*, 1986).

13. CONCLUDING REMARKS

If one is willing to abandon the assumption that choice is always identical
to, or revealing of, preference, one can argue that, though choice may
sometimes be intransitive, preferences, such as they are, are not. To be

sure, we offer no way of ascertaining what preferences really are indepen-
dently of choice, hence this claim would be difficult to substantiate.
Nonetheless, intuition and introspection suggest that the claim is tenable.
Consider the types of examples of intransitivity presented in the present
paper, as summarized in Table V.

1. Societies can hardly be said to have preferences at all, therefore the
intransitivity of the majority rule is clearly a property of that rule, rather
than revealing of any property of societal preferences.

2. Preference among multi-dimensional alternatives can be maintained
to be transitive, even though the adoption of a simplifying comparative
rule might conceal this fact. The intransitive ordering may simply approx-
imate the transitive one, by reversing the ordering between nearly equiva-
lent options. The true preference can be revealed at some additional
computational, or other, cost.

3. Ordering alternatives by their probability of giving you a certain
fixed desired (or undesired) outcome is a transitive ordering when the
competition is kept constant. The fact that when it is allowed to change,
probabilistic prevalence can lead to intransitivity is not, therefore, indica-
tive of intransitive preference (even under the classic preference
approach).

Tversky's 1969 paper was titled "Intransitivity of preferences". Per-
haps preferences – if and when they can be meaningfully said to exist –
are never intransitive, though choices – since they are determined by other
considerations besides preferences – may sometimes form intransitive
cycles.

NOTES

* We wish to thank the colleagues and students, too many to list, with whom we have
discussed the ideas in this paper, and in particular to acknowledge our heavy intellectual
debt to Amos Tversky, who helped and inspired us in all stages of this study.
[1] In the logic of relations, it is customary to distinguish between transitive, intransitive, and
non-transitive relations (e.g., Suppes, 1958). For our purposes, however, it is sufficient to
distinguish the first from the other two, which we shall label generically "intransitive".
[2] An example is when the difference between the attractiveness of a $60000 job and that
of a $50000 job looms smaller than the sum of the difference between the attractiveness
of a $60000 job and that of a $55000 job, and the difference between the attractiveness
of a $55000 job and that of a $50000 job.
[3] In the particular instance at hand, for example, ordering the candidates solely by their

IQ scores – even when these differ by less than three points – would have yielded a transitive ordering, and yet only pairs of candidates who are anyway very close in terms of their qualifications would have had their orders reversed by the two rules.

[4] The cyclic ordering of Table I's rows in this example is the reverse of the ordering obtained in the previous examples that relied on this table. This is due to the particular definition of pp. Note, however, that the reverse relation, $p(A < B) > 0.5$, is also cyclic, and orders the rows in the reverse cycle.

[5] The rules of this game dictate that Rock overcomes Scissors (by smashing them), Scissors overcome Paper (by shredding it), and Paper overcomes Rock (by enfolding it). Hence, every pair consists of a "winner" and a "loser". Since, however, the players must choose from among all *three* strategies, and do so simultaneously, the options are rendered equivalent, and both players are equally likely to win.

[6] Note that $p(A \leq B) \leq 1/2$ is not pp, but rather its inverse. However, the two relationships enjoy isomorphic properties, and the proposed interpretation calls for the "winning" variable to be the one with the smaller value. Note further that in the example represented by Table I – which is essentially, though not numerically, identical to Blyth's – the three random variables that form the pp cycle are dependent. Blyth additionally constructed an example of *in*dependent random variables with this property, as follows: A assumes the value 3 with certainty; B assumes the value 2 with probability 0.6, and otherwise is 5. C assumes the value 1 with probability 0.4, and otherwise is 4. Hence, A *pp*s B (since $p(A \geq B) = 0.6$); B *pp*s C (since $p(B \geq C) = 0.4 + 0.6 \times 0.4 = 0.64$); and C *pp*s A (since $p(C \geq A) = 0.6$).

[7] In this example, the three random variables are independent. IN addition, Blyth gave an example of *dependent* random variables with this property, as follows: The variables A, B and C assume the values 1, 2 and 3, respectively, with probability 0.25; they assume the values 2, 1 and 3, respectively, with probability 0.35; and they assume the values 2, 3 and 1, respectively with probability 0.40.

[8] By switching the labels A and C in Figure 1, and altering the values i to $7\text{-}i$, for $i = 1, ..., 6$, an example of this state of affairs is obtained.

[9] The actual taste experience in the pie example is analogous to the actual running time of the runner bet upon – a variable that played no independent role in the betting example.

[10] If the gambles were not identically distributed (i.e., if one had a higher expectation, or more risk, or different skew, than the other), we would have needed to make assumptions as to how the comparative aspect trades off with these other considerations, but for the present purposes this is unnecessary.

REFERENCES

Arrow, K. J.: 1951, *Social Choice and Individual Values*, New York: Wiley.

Bell, D. E.: 1982, 'Regret in Decision Making under Uncertainty', *Operations Research*, 961–981.

Bell, D. E., Raiffa, H., and Tversky, A.: 1983, 'Normative, Descriptive and Prescriptive Interactions', Harvard, unpublished manuscript.

Blyth, C. R.: 1972, 'Some Probability Paradoxes in Choice from among Random Alternatives', *Journal of the American Statistical Association* **67**, 366–373.

Dummet, M. C.: 1984, *Voting Procedures*, Oxford: Clarendon Press.

Fishburn, P. C.: 1970, 'The Irrationality of Transitivity in Social Choice', *Behavioral Science* **15**, 119–123.

Fishburn, P. C.: 1982, 'Nontransitive Measurable Utility', *Journal of Mathematical Psychology* **26**, pp. 31–67.

Fishburn, P. C.: 1984, 'SSB Utility Theory: An Economic Perspective', *Mathematical Social Science* **8**, 63–94.

Lindman, H. R. and Lyons, J.: 1978, 'Stimulus Complexity and Choice Inconsistency among Gambles', *Organizational Behavior and Human Performance* **21**, 146–159.

Loomes, G. and Sugden, R.: 1982, 'Regret Theory: An Alternative Theory of Rational Choice under Uncertainty', *The Economic Journal* **92**, 805–824.

Luce, R. D. and Raiffa, H.: 1957, *Games and Decisions*, New York: Wiley.

MacCrimmon, K. R. and Larsson, S.: 1979, 'Utility Theory: Axioms Versus "Paradoxes" ', in M. Allais and O. Hagen (eds.), *Expected Utility Hypotheses and the Allais Paradox*, Dordrecht, Holland: D. Reidel.

May, K. O.: 1954, 'Intransitivity, Utility, and the Aggregation of Preference Patterns', *Econometrica* **22**, 1–13.

Pratt, J. W.: 1972, 'Comment', *Journal of the American Statistical Association* **67**, 378–379.

Sage, A. P. and White, E. B.: 1983, 'Decision and Information Structures in Regret Models of Judgment and Choice', *IEEE Transactions on Systems, Man and Cybernetics* **13** (3), 136–145.

Samuelson, P. A.: 1948, 'Consumption Theory in Terms of Revealed Preference', *Economica*, N.S., 15.

Simon, H. A.: 1957, *Models of Man*, New York: Wiley.

Slovic, P. and Lichtenstein, S.: 1983, 'Preference Reversals: A Broader Perspective', *The American Economic Review* **73** (4), 596–605.

Slovic, P. and MacPhillamy, D.: 1974, 'Dimensional Commensurability and Cue Utilization in Comparative Judgment', *Organizational Behavior and Human Performance* **11**, 172–194.

Suppes, P. C.: 1958, *Introduction to Logic*, Princeton, N. J.: D. Van Nostrand.

Tversky, A.: 1969, 'Intransitivity of Preferences', *Psychological Review* **76**, 31–48.

Tversky, A. and Kahneman, D.: 1986, 'Rational Choice and the Framing of Decisions', *The Journal of Business* **59** (4), 251–278.

Tversky, A., Sattath, S., and Slovic, P.: 1987, 'Contingent Weighting in Judgment and Choice', unpublished manuscript.

Ullman-Margalit, E. and Morgenbesser, S.: 1977, 'Picking and Choosing', *Social Research* **44** (4), 757–785.

Department of Psychology,
The Hebrew University,
Jerusalem 91905,
Israel