The Conditions of Occurrence of the Preference Reversal Phenomenon

by

Robert M. Hamm

January 1984

Institute of Cognitive Science

University of Colorado
Campus Box 345
Boulder, CO 80309-0345

ICS Technical Report #90-11

Short Abstract

The stability of the preference reversal phenomenon, a reversal of preference between a high-risk and a low-risk gamble when they are evaluated by two different evaluation methods, was investigated by engaging each of 109 subjects in a directed discussion designed to influence their evaluation strategies. Two stable subsets of subjects were found: those with no reversals and those with a particular, predicted reversal pattern. The discussions generally increased the number of subjects having no reversals, and decreased the number of subjects having a particular, unpredicted reversal pattern. There was, however, no reduction in the number of subjects having the predicted reversal pattern.

Individual analysis, which had not previously been used to investigate the preference reversal problem, was able to reveal the existence of the two stable groups and the nature of the effects of the advice, and is recommended for preference reversal investigations where there are stable individual differences.

Abstract

The stability of the preference reversal phenomenon in the face of manipulations designed to make subjects consistent was investigated. The preference reversal phenomenon is a reversal of preference between a high-risk and a low-risk gamble when they are evaluated by two different evaluation methods, choosing between the pair of gambles and naming a limit price for each gamble individually. Of the two possible forms of reversal, the form predicted by Lichtenstein and Slovic (1971), which involves choosing the low risk bet but setting a higher limit price for the high risk bet, is more commonly observed.

The stability of the preference reversal phenomenon was tested by engaging each of 109 subjects in a directed discussion designed to influence the strategies they used on the evaluation methods. Subjects were advised to attend more to the values or the probabilities in the gambles, or encouraged to take a more analytic or a more intuitive approach. Two stable subsets of subjects were found: those with no reversals and those with the predicted reversal pattern. Except for the intuitive advice, the influence conditions increased the number of subjects having no reversals, and decreased the number of subjects having the unpredicted reversal pattern. There was, however, no reduction in the number of subjects having the predicted reversal pattern.

There were two major methodological results. First, two methods that have previously been used for measuring the extent to which a group of subjects have preference reversals were shown to be inconsistent: (a) comparing the reversal rates conditional on whether the high-risk or low-risk bet was chosen, and (b) testing whether the results can be explained by different error rates on the two evaluation methods. It is argued that the latter method is better. Second, individual analysis, which had not previously been used to investigate the preference reversal problem, was able to reveal the existence of the two stable groups and the nature of the effects of the advice, while the two aggregate methods could not. The use of individual analysis is recommended for preference reversal investigations where there are stable individual differences.

Finally, it is argued that progress in understanding the preference reversal phenomenon depends on researchers addressing the issue of generalization over (a) subjects (b) different knowledge or attitude states on the part of the subject, (c) evaluation methods, and (d) gambles.

The Conditions of Occurrence of the Preference Reversal Phenomenon

Consistent, rational behavior is an ideal perhaps more distant than generally thought. An inconsistency in the expression of preferences, the preference reversal phenomenon, has recently been demonstrated (Grether & Plott, 1979; Lichtenstein & Slovic, 1971): the preference between two gambles that is revealed by one evaluation method is the opposite of the preference revealed by another. Subjective expected utility theory, or indeed any theory of decision making in which the elements of gambles--the probabilities and values--are evaluated and combined in a consistent way in all situations. We can no longer speak generally of the utility of something to someone; we can speak only of its utility to him or her under particular evaluation conditions. On the basis of recent replications (Pommerehne, Schneider, & Zweifel, 1982; Reilly, 1982) which showed the effect to be smaller and less stable than in the original studies, it has been argued that the preference reversal phenomenon does not present as major a problem as had been suggested. However, in this paper we show that it is stable despite some manipulations designed to make the preferences consistent.

Lichtenstein and Slovic (1971) demonstrated that when people use two methods to evaluate a pair of monetary gambles--choosing between the two gambles, and setting a limit price for each gamble--they often choose the gamble to which they would assign less money. For example, consider these two two-outcome gambles:

Gamble	Value of Winning (V_w)	Probability of Winning (P_w)	Value of Losing (V_l)	Probability of Losing (P_l)	Expected Value (EV)
Hi-P	$4.00	35/36	-$1.00	1/36	$3.86
Hi-V	$16.00	11/36	-$1.50	25/36	$3.85

Typically, when choosing between a high probability (Hi-P) and a high value (Hi-V) bet of nearly equal expected value, people will select the Hi-P bet; yet when naming the lowest or limit price for which they would sell each gamble, they will indicate that the Hi-V bet is worth more money.

To answer initial criticisms that the preference reversals they had found were not important because the bets involved hypothetical or negligibly small prizes, Lichtenstein and Slovic (1973) replicated the study as a game in a gambling casino, demonstrating that the phenomenon still occurs when substantial amounts of money actually change hands.

Grether and Plott (1979), concerned that the inconsistency between choice and pricing behavior contradicted a basic assumption of economic theory, that people seek to optimize, designed a series of experiments to discredit the work as applied to economics. They tested whether alternative explanations might account for the preference reversals. First, subjects might really be indifferent between the Hi-P and Hi-V bets and, when forced to indicate a preference, choose capriciously. Grether and Plott therefore allowed "indifference" as a possible choice response. Second, the reversal could be due to the combined effect of the subjects' strategic responding when setting a selling price (over-pricing) and ceiling effects (because there is more room to overprice the Hi-V bets, compared to their expected value). To prevent strategic responding, Grether and Plott devised a condition that did not refer to market behavior, asking the subject instead to name "the exact dollar amount such that you are indifferent between the bet and the amount of money." Third, since subjects might not be sufficiently motivated to be rational, Grether and Plott varied the extent to which subjects' pay depended on the wisdom of their choices. Though this manipulation might not be sufficient to eliminate reversals, it could at least indicate whether insufficient motivation were a factor. Finally, suspecting that subjects do not take psychology experiments seriously, Grether and Plott ran theirs as "economics experiments".

To their surprise, the preference reversal phenomenon occurred after all these manipulations, indicating that it indeed exists in situations where economic theory is generally applied. The only class of psychological theory consistent with the phenomenon, they reluctantly concluded, holds that preference is influenced by the method with which it is expressed.

Additional work by economists has confirmed Grether and Plott's (1979) findings. Pommerehne, Schneider, and Zweifel (1982) showed that neither repeated experience with such gambles nor access to records of their previous answers eliminated their subjects' reversals. Reilly (1982) used an idea of Lichtenstein, Slovic, and Zink's (1969) and provided subjects with information about the expected values of the gambles, but this did not prevent the reversals. Each of these manipulations seemed, however, to cause moderate reductions in the rate of reversals.

Mowen and Gentry (1980) have shown preference reversals in a managerial decision task in which the uncertain prospects were product development programs rather than monetary gambles. Interestingly, when these decisions were made by groups of 3 or 4 subjects a greater incidence of preference reversals occurred than when the same prospects were evaluated by individuals.

Goldstein (1983) demonstrated that preference reversals can be due to both the response mode and the attribute the subject is evaluating (i.e., how desirable it would be to play the gamble versus how much money the gamble should be sold for). By manipulating both of these factors he discovered conditions which produce the opposite preference reversal pattern from what had been observed before. That is, subjects selected the Hi-V bet yet indicated, with their single gamble evaluations, that the Hi-P bet was better.

Lichtenstein and Slovic's (1971) preference reversal phenomenon is fairly firmly established by this body of work, despite some variation in the extent of the reversal pattern. The majority of people express reversed preferences about particular, carefully selected pairs of gambles, when using choice and some form of pricing as their two evaluation methods, in situations sufficiently similar to "real life" to occasion some concern. It remains to be established why these reversals occur and whether they can be prevented.

Why do preference reversals occur? The most likely general explanation for the preference reversals is that the two evaluation methods evoke different evaluation strategies, each of which seems justifiable but which produce preferences that are inconsistent with each other. Lichtenstein and Slovic's (1971) particular explanation is that the monetary response mode of the pricing task induces a strategy of anchoring on the V_W and adjusting insufficiently to take the P_W into account. Thus the V_W dimension is given most weight. In contrast, the choice task induces a dimension-wise comparison procedure in which the important dimensions V_W and P_W are given about equal weight. Consequently the ratio of the importance of V_W to the importance of P_W is higher with the pricing evaluation method than with the choice method. In the space of gambles varying in V_W and P_W, these two methods would therefore produce different families of hyperbolic indifference curves; judicious or fortuitous selection of pairs of gambles can produce preference reversals between the two methods. Although this particular explanation does not cover Goldstein's (1983) findings, still the general form of Lichtenstein and Slovic's (1971) explanation seems reasonable: the evaluation methods evoke different response strategies.

Can preference reversals be prevented? They would not occur if the subject would faithfully use one internally consistent (transitive) evaluation strategy in all response modes. Further, if subjects' reversals were called to their attention with some form of outcome feedback, they would probably adjust, if only on an ad hoc basis, in order to become consistent. The existence of the preference reversals indicates, then, that subjects are neither accurately following one gamble evaluation strategy under all response modes nor adequately monitoring themselves for inconsistency. These behaviors may be beyond subjects' information processing abilities, or simply may be something they would not want to do. It would be trivial to instruct subjects to use response strategies guaranteed not to reverse (such as determining their pricing and choice responses from calculation of the expected values of the gambles), or to alert them to the possibility of inconsistency and give them sufficient information to discover and avoid it. Either approach would probably eliminate preference reversals (cf. Berg, Dickhaut, & O'Brien, 1983; Reilly, 1982), but this kind of specific intervention may not generalize to real world situations in which the subject is not controlled or assisted by the experimenter.

However, there may be some way of sensitizing people to the features of the gambles or to their own evaluation processes so that they could have consistent preferences without depending on the continued presence of the experimenter. It may be possible (a) to affect subjects' choice of strategy; (b) to induce them to apply the same relative weights to the P_W and V_W in all evaluation modes; or (c) simply to develop their understanding of the situation. These possibilities were tested in the experiment described below.

Several methodological issues are also explored in this study:

Aggregate versus individual analysis. All previous studies of the preference reversal phenomenon have aggregated data across subjects. To facilitate comparison with earlier studies, we too present an aggregated analysis. But since aggregation is known to obscure individual differences and even to produce misleading conclusions (Hammond, McClelland, & Mumpower, 1980, Chapter 10), we further analyze our data in terms of individuals.

Measuring preference reversals. Preference reversals in groups of people have previously been measured in two ways, neither of which is wholly satisfactory. Grether and Plott (1979), Mowen and Gentry (1980), Pommerehne, Schneider, and Zweifel (1982), and Reilly (1982) index the extent of preference reversal by comparing the rate of reversal when the Hi-P bet has been chosen to the rate when the Hi-V bet has been chosen. If the reversal rate for the former is higher, the general pattern of choosing the Hi-P bet but paying more for the Hi-V bet is confirmed. A problem with this method is that it is arbitrary whether to define the preference reversal as conditional on the choice or on the prices, yet the results produced by these alternatives can be quite different. Furthermore, the index can be influenced by differential error rates for the two evaluation methods. In approximately half of the conditions in the studies by Grether and Plott (1979), Mowen and Gentry (1980), Pommerehne, Schneider, and Zweifel (1982), and Berg, Dickhaut, and O'Brien (1983), and in all the conditions in Reilly (1982), different error rates on the two evaluation methods, rather than true preference reversals, might account for the observed pattern of reversals. The Reversal Rate Comparison index does not distinguish between these two possibilities.

A second method of measuring preference reversal, used only by Lichtenstein and Slovic (1971), involves attempting to fit a null model to the subjects' responses. The null model assumes that there is a single preference rate for Hi-P over Hi-V gambles and that reversals are due solely to error. There can be different error rates on the two evaluation methods. Rejection of this null model implies that the subjects have different preferences on the two evaluation methods. The disadvantage of this method is lack of a general statistical test for rejecting the null model (judgment, rules of thumb, or statistical tests applicable only in special circumstances are used instead). However, the method is descriptive and it addresses the possibility of different error rates.

In the present study, we will demonstrate the relative advantages of these two methods for measuring the existence and extent of the preference reversal phenomenon.

Buying versus selling methods of obtaining limit prices. Previous studies of the preference reversal phenomenon, except for Lichtenstein and Slovic's Experiment II (1971), have used selling limit prices. The predominance of the Selling Limit Price method is partly due to the erroneous belief that with the Selling Limit Price method the price expresses the expected utility of the gamble no matter what the subject's utility function, whereas in the Buying Limit Price method the price equals the gamble's expected utility only if the subject's utilities are linear or exponential. In fact both methods have the same limitation.

The erroneous argument seems to originate with Krantz and Tversky (1965, cited in Becker & McClintock, 1967), is clearly present in Raiffa (1968, pp. 89-90), and is repeated by Lichtenstein and Slovic (1971, footnote 5, p. 54) and Grether and Plott (1979, footnote 5, p. 630). The error results from the failure to recognize that when one sells a gamble one loses its prizes (Hamm, 1979, Appendix 1-f). In terms of Farquhar's (1982) typology, the argument confuses the expected utilities in the Sale and Gift situations.

As the buying situation is more familiar for college students, and hence more likely to produce accurate preference reports, we used the Buying Limit Price method to get limit prices in the present study.

Method

Subjects evaluated three sets of 6 gamble pairs, using both choice and pricing methods (see Figure 1). During the second set of gamble pairs, one of six influence conditions was administered. In these influence conditions the subjects were not told to use a particular strategy, nor alerted to the possibility of reversals. Rather, they were directed to focus on general issues concerning the evaluation of gambles. Each influence was applied equally to the choice and pricing response modes. Comparison between the first and third sets of gamble pairs allowed a test of whether the influence condition affected the rate of preference reversals.

Insert Figure 1 about here

The Six Influence Conditions

Value versus Probability Advice. In the first two conditions the subjects were advised to pay more attention to either the probabilities or the values in the gambles. These conditions tested whether preference reversals were due to the evaluation methods' induction of strategies which involve different relative weights on these dimensions. Explicit advice with respect to those relative weights was expected to counter the tendency of the response modes to induce different relative weights.

Analytic versus Intuitive Advice. In the third and fourth conditions the subjects were advised to take either a more analytical or a more intuitive approach to the task.[1] Analytic advice was expected to cause fewer reversals, by one of two mechanisms. It might make the subject use an analytic strategy, such as computing the expected value of the gambles, on one or both evaluation tasks. It also might attune the subject to the consistency between the methods. Intuitive advice was expected to have the opposite effect, through encouraging subjects to make uncritical use of the immediately available strategies, which are assumed to be the ones that induce preference reversals.

Control conditions: Discussion and Practice. The fifth condition involved undirected discussion, in which the subject was asked to explain his or her evaluations. This discussion was expected to stimulate the development of the subjects' evaluation strategies and awareness of possible inconsistencies, and hence cause fewer reversals. This condition also served as a control for the four advice conditions, which involved such discussion in addition to the advice.

In the sixth, Practice, condition there were no influence manipulations. The subject simply did the three sets without discussion. If, as Lindman (1971) and Pommerehne, Schneider, and Zweifel (1982) have suggested, experience evaluating these gambles causes fewer preference reversals, we would expect to find less indication of the preference reversal pattern on the third set than the first, whether or not the subject had received any advice. Thus this practice condition served as a control for the other conditions.

Previous workers have demonstrated that the preference reversal phenomenon occurs with both hypothetical and real gambles; when equivalent values are determined via buying and selling limit prices as well as via non-market evaluation situations; with economics students and adult gamblers as well as with psychology subject pool students; and with economist as well as psychologist experimenters (Grether & Plott, 1979; Lichtenstein & Slovic, 1971; 1973; Pommerehne, Schneider, & Zweifel, 1982; Reilly, 1982). We therefore felt that it was sufficient to conduct a psychology experiment in which college students stated their limit buying prices for hypothetical gambles.

One hundred and nine subjects (54 M and 55 F) between 15 and 55 years of age (mean 22) participated in a single individual 1- to 5-hour session, with 1 of 6 experimenters.[2] A computer was programmed to present gambles, record answers and response times, count reversals, and present the advice to the subjects.[3]

The independent variables in the experiment were crossed in a 2 x 2 x 2 x 2 x 6 factorial design: Sex of Subject by Task Order (Choice or Pricing first) by Gamble Set Order by Experimenter Type by Influence Condition. Half of the 96 cells in this design were empty. In this paper we report only the effects of Influence Condition. Further details are available in Hamm (1979).

Procedure

Most subjects were college students (a few were summer school students or older people) without a particular interest in gambling. They were run individually with the experimenter constantly present, in a separate room with a CRT terminal connected to the computer. The computer terminal and terminology were unfamiliar to the subjects, as were the gamble display and the pricing task. Although subjects were told there was no time pressure, they may occasionally have been aware of tension between the experimenter and other users over conflicts in the computer schedule.

The First Stimulus Set

First, the gamble display was explained. The subjects were told to consider each decision independently and not to worry about accumulation of wins or losses from these hypothetical gambles. They were then given instructions for their first task.

Instructions for the choice task showed a sample choice and asked subjects to indicate their choice or state that they were indifferent between the gambles.[4] Subjects were asked to practice on the sample shown on the sheet. All showed that they understood the basic concept by choosing the dominating bet.

The Pricing task instructions introduced a "buying" version of the Becker, DeGroot, and Marschak (1964) bidding procedure (see Appendix A). This procedure encourages the subject to pay exactly what he or she thinks the gamble would be worth, without trying to get a bargain by offering lower prices. The subject was asked to name the highest price he or she would be willing to pay for a ticket for the hypothetical gamble, and was told to imagine that the actual price would be determined by a lottery. If the price drawn were lower than the subject's limit price then he or she would play the gamble for that price; if it were higher, the subject would not play. The universe of possible prices was the range between the value of the win and the value of the loss. Negative prices were considered payments to the subject. The subject could name a negative number as the highest price, indicating that he or she would have to be paid at least that amount in order to play the gamble. (Reilly, 1982, and Goldberg, 1983, similarly extended the Selling Limit Price method into the negative range.) For practice, the subject was asked to evaluate the sample gamble shown on the instruction sheet. After having completed all the evaluations in the first task, the subject was introduced to the other task and completed the set.

During this first stimulus set, the experimenter maintained a formal, distant attitude toward the subject. The experimenter would pay attention to what the subject did, sometimes taking notes, but conversation was discouraged through statements such as "I'll be right here if there are any questions" (implying, "if there aren't, let's not talk") and "Whenever the computer says "HIT RETURN WHEN READY", stop and I'll tell you what to do next." There were two purposes for this approach. First, conversation would increase the computer's response time measures. Second, since subjects might learn from conversation, it was saved for use as a manipulation during the second stimulus set.

After the first set, in every condition except Practice, the Experimenter directed the computer to print the subject's results and the advice corresponding to the condition.

The Second Stimulus Set

In the second stimulus set, all subjects did six Pricings, three Choices, six Pricings, and then three Choices (see Figure 1). In the Practice condition, the subjects were allowed to continue as they had done on the first stimulus set, with little conversation. In each of the other five conditions, the subjects were told, "We are going to do this set differently. I'd like you to think aloud as you look at these gambles, and tell how you come to your answer. Just say everything that comes to mind." The experimenter often prompted the subject by asking questions and calling the subject's attention to different aspects of the stimulus and of the subject's behavior. Subjects in the Discussion condition continued in this manner through this entire stimulus set.

Subjects in the advice conditions were given the computer's printed advice at the halfway point, when no gamble was displayed on the screen (see Appendix B for a sample). They were told that the computer had analyzed the pattern of their responses to the first stimulus set and had selected some statements of advice for them on the basis of this analysis. It was further explained that the analysis program was at an exploratory stage, and their help in evaluating its advice was needed. They were asked to read each statement and rate, on a 1 to 7 scale, whether this advice applied to them in particular.

This approach served a number of purposes. The experimenters found it difficult to give the subjects false or irrelevant advice and defend it during a quarter-hour conversation, but easier to discuss advice that had come "from the computer" and to convince subjects that the advice was particularly addressed to them. This was essential, for in pilot work subjects had paid little attention to paragraphs that obviously did not apply to them personally.

In each condition the same basic advice was written in each of the four modes of psychiatric discourse: empathic, objective, psychoanalytic, and counterprojective (Havens, 1978). The task of evaluating these statements forced subjects to pay close attention to them. The experimenter had the subject read each statement aloud, and they discussed its meaning and its appropriateness for the subject. The experimenter would take the position that the advice was probably correct, and during the second half of the second stimulus set he or she would keep the discussion focused on the relevance of the advice.

The advantage of allowing the subjects to accept or reject the advice is that being able to reject advice as "mistaken" prevented them from losing self esteem or faith in the experiment. If the advice told subjects to do something they were already doing, they could reject it but still be motivated to follow the advice just to prove that they did not need it.

The Third Stimulus Set

The last stimulus set, like the first, was run with little conversation. In all conditions except Practice, the experimenter announced at the beginning, "I will not be taking notes on this set, so you don't have to talk. Just do it whatever way seems most natural to you." Though this statement may seem to undermine the effect of the advice, it was used not only to interrupt the momentum of conversation but also to interrupt any tendency the subject may have had to make judgments in a certain way in order to please the experimenter. We were less interested in our power to elicit such behavior than in what lasting effects the advice and discussion might have.

After the last gamble of the third set, the computer displayed, for each gamble pair where there was a preference reversal, the stimuli and the subject's responses. In order to investigate whether the reversal was due to motivational or cognitive causes, subjects were asked whether they felt this was a problem requiring correction, and indicated which responses they would change in order to become consistent.

Finally, subjects were informed that the advice was false, i.e., given to them independently of their performance. Most had not suspected this. Data from those who had been suspicious was included in the analysis because even though they did not believe that the advice applied to them they had still spent 15 minutes discussing the gambles and their strategies in the desired way.

Results

Confrontation

All but five of the 109 subjects committed at least one reversal. Of the 81 subjects who committed a reversal on the "After" set of gamble pairs and were confronted with their inconsistencies, 80 said that the inconsistencies represented a problem and changed their answers so as to be consistent. Most of them changed their limit prices rather than their choices, thus indicating that they had more confidence in the choice method of evaluating gambles.

Preference Reversal Pattern

Table 1 shows the four possible preference patterns a subject could have, using the two evaluation methods on any pair of gambles. If the Hi-P bet were chosen, a higher Buying limit price could have been set for it (Cell A) or for the Hi-V bet (Cell B); if on the other hand the Hi-V bet were chosen, more money could have been offered for it (Cell D) or for the Hi-P bet (Cell C). Cells B and C involve inconsistencies between the two methods of indicating preference. Previous researchers (except Goldstein (1983) and Lichtenstein and Slovic (1973) with the negative expected value gamble pairs) have observed Cell B to be more common than Cell C; Lichtenstein and Slovic's (1971) theory (that setting a limit price calls attention to the value of winning, while choosing causes more weight to be put on the probability of winning) predicts the reversal pattern of Cell B but not of Cell C.

Insert Table 1 about here

In addressing whether the preference reversal phenomenon occurred in this study (i.e., whether the reversals were in a regular pattern or due only to random error) and whether the influence conditions caused it to disappear, we shall first present analyses of data aggregated across subjects, then an analysis of individual subjects' data.

Analyses Involving Data Aggregated across Subjects

Table 2 shows the data aggregated across subjects for each of the six conditions, and for all conditions combined, before and after the manipulations. The entries for each of the cells (as defined in Table 1) are the proportion of the responses from N subjects, six pairs each, that displayed the indicated pattern. The entries do not sum to 1.00 because some subjects were indifferent on the choice task and/or offered the same price for both bets. We shall next compare the Reversal Rate Comparison method and the Two-Error-Rate Model method (described above) as measures of the existence and extent of preference reversals. Each of them is applied to the data in Table 2.

Insert Table 2 about here

First Aggregate Method: Comparison of Predicted and Unpredicted Reversal Rates

The preference reversal phenomenon can be seen by comparing the proportion of times that the prices reflect the opposite preference from the choice, for those gamble pairs where the Hi-P bet as opposed to the Hi-V bet was chosen. In terms of Table 1, this means subtracting C / (C + D) from B / (A + B).[5]

The proportion of reversals of each type and the difference between the proportions, both before and after the influence, are shown in Table 3. Additionally, the change due to the influence is shown. In every case except Analytic (Before), there is a higher reversal rate when the Hi-P bet was chosen than when the Hi-V bet was chosen, consistent with the pattern observed by previous researchers. As subjects in all the Before sets received identical treatment, their results may be combined. For this All Conditions (Before) group, Hi-P choices have a reversal rate of .48 (the predicted reversal) and the Hi-V choices a rate of .28 (the unpredicted reversal). This is comparable to the rates in Lichtenstein and Slovic's Experiment II (1971), .51 predicted reversals and .27 unpredicted reversals. As in the present study, they used the Buying Bid method for obtaining prices. Studies using the Selling Bid method (e.g., Grether & Plott, 1979; Lichtenstein & Slovic, 1971, Experiments I and III) have generally produced larger differences (see discussion). The effect of the advice conditions is shown in the rightmost 3 columns of Table 3. The Analytic Advice, Intuitive Advice, and Discussion conditions involved increases in the extent of preference reversals, the Value Advice condition involved no change, and the Probability Advice and Practice conditions involved decreases.

Insert Table 3 about here

In summary, the method of comparing the predicted and unpredicted reversal rates reveals a general pattern of preference reversal. Three influence conditions increased the extent of this preference reversal, two decreased it slightly, and one produced no change. Thus, using the difference between predicted and unpredicted reversal rates as an index of the preference reversal phenomenon, preference reversals occurred in this study and were not generally reduced by the influence conditions.

Second Aggregate Method: The Two-Error-Rate Model

We said above that even if subjects' underlying preferences for Hi-P versus Hi-V bets did not vary with evaluation method, differences in predicted and unpredicted reversal rates could be caused by a difference in error rates between the two evaluation methods. Lichtenstein and Slovic (1971) showed how to test whether such a Two Error Rate model fits data in the form of Tables 1 and 2, or whether a Two Preference Rate (i.e., preference reversal) model is necessary.

Assume that: (a) subjects have a probabilistic preference between Hi-P and Hi-V bets, such that they prefer Hi-P bets a proportion p of the time; (b) they choose incorrectly a proportion r of the time; and (c) their limit prices fail to reflect their true preferences a proportion s of the time. (In fitting this model to the data of a number of individuals, we need not assume they all have the same parameters p, r, and s, only that their individual parameters are normally distributed about the group's mean parameters.) Thus, proportion $p(1 - r)s + (1 - p)r(1 - s)$ of the answers would fall in Cell B of Table 1--those where the Hi-P bet was really preferred and the subject indicated this correctly with the choice yet erred with the pricing response, and those where the Hi-V bet was really preferred and the subject indicated this correctly when naming prices yet erred with the choice.

Lichtenstein and Slovic (1971) show how the model parameters may be calculated from data expressed as in Table 1 (when A + B + C + D are normalized to sum to 1):[6]

$$p(1 - p) = \frac{AD - BC}{(A + D) - (B + C)}$$

$$p(1 - r) + (1 - p)r = A + B$$

$$p(1 - s) + (1 - p)s = A + C$$

Applying the quadratic equation to the first expression produces two solutions for p; the one where the expressions of preference are incorrect less than half the time (r and s < .5) is used.

Estimated parameters of the Two Error Rate model for a particular data set must be evaluated to determine whether the model accounts sensibly for the data. If it does not, a model which allows different 'true preferences' on each method, in addition to different error rates, would be needed. Such a model would predict true preference reversals for some possible gamble pairs.

Evaluating the fit of the model requires decisions about three questions (Figure 2):

1. Are the parameters p, r, and s technically acceptable? For some data sets, p(1 - p) is estimated at more than .25, so that there is no real solution for p; for others the probabilities p, r or s are outside the 0 to 1 range. Such results do not necessarily mean the Two Error Rate model is wrong, for sample variation can cause odd fits.

2. If the parameters are technically acceptable, are they plausible? One case of implausible fit occurs when one of the error rates is greater than .5 (Category C). It is also implausible if a higher error rate is found for choices (r) than for limit prices (s), because as we reported above, most subjects had more confidence in the Choice Method than in the Limit Price method.

3. The sampling characteristics of the Two Error Rate model are not known. Even if the parameters of the Two Error Rate model are beyond meaningful bounds, the cause may simply be sampling variation rather than different preferences on the two methods. If p, r, or s are less than 0 or greater than 1, there is no way to test how strong the

preference reversal pattern is. However, Lichtenstein and Slovic (1971) showed that when $p(1 - p) > .25$, the model can be tested statistically even though there is no real solution to the quadratic equation. The assumption is made that the population p equals .5; therefore, the numbers in the B and C cells should be equal. If the difference between them is significant, according to McNemar's test for correlated proportions (Hays, 1973, p. 740), then one must reject a Two Error Rate model that sets parameter p to .5.

In Figure 2 the possible modes of success or failure of the Two Error Rate model are labeled. A is an unequivocal success, with no hint of preference reversals; B and C implausible; D a technical failure not significantly different from a Two-Error-Rate model; E one that is; and F, G, and H are untested technical failures. Thus, only in Category E is there surely a true preference reversal pattern, and only in Category A is there surely none; all other categories require the researcher to judge how strong the preference reversal is. We adopt the liberal rule of thumb that only categories A and D should be considered not to have a preference reversal pattern.

Insert Figure 2 about here

The fits of the Two-Error-Rate model to the Before and After data of each of the six Influence groups, and of all conditions combined, are shown in Table 4. There was evidence for a preference reversal pattern in all but three of the conditions--Intuitive Advice (Before), Discussion (Before), and Probability Advice (After). However, the preference reversal pattern was unequivocal only in the Practice (Before) and Discussion (After) conditions. Just as with the Reversal Rate Comparison method, the Two-Error-Rate Model method indicates that the preference reversal occurs in most of the groups, both before and after the influence attempts. As for the influences, Probability Advice and Practice decrease the preference reversal pattern, Value and Analytic Advice cause no change in degree of preference reversal, and Intuitive Advice and especially Discussion cause a stronger preference reversal pattern.

Insert Table 4 about here

However, the two methods of gauging the extent of the preference reversal phenomenon do not indicate the same order of strengths of reversal patterns among conditions. For example, by the Reversal Rate Comparison method, the subjects in the Value Advice condition had the strongest preference reversals; but by the Two-Error-Rate Model method, there is only weak evidence for preference reversal in the Value advice condition--it hinges on the implausibility that pricing should be more accurate than choice when most subjects had more confidence in choice. In such cases of conflicting implications, the Two-Error-Rate Model method is to be given more credence because it treats error separately from preference. Further, its descriptive features enable it to explain why the Reversal Rate Comparison method should give such a high score to the Value Advice group: the p parameter is very low, indicating that these subjects are risk seekers with a high preference for the Hi-V bets on both methods. Whether or not their general preference changes between evaluation methods, they would tend to have a high rate of reversals on the rare occasions when they go against their general preference and choose the Hi-P bet.

Influence Conditions

Did the influence conditions succeed in eliminating the preference reversals or lessening their frequency? Let us consider each pair of manipulations in turn.

Discussion and Practice conditions. These conditions were designed as controls for the advice conditions. What effects would simple repetition or undirected discussion have on the subjects' preference reversals? Table 4 shows that the subjects in the Practice condition had a strong preference reversal on their first set of gamble pairs, while the subjects in the Discussion condition had none. On the second set, both groups exhibited the preference reversal pattern. These results indicate that the preference reversal pattern does not tend to disappear spontaneously. Therefore any disappearance following an advice condition will have to be taken seriously. Additionally, comparing Discussion to Practice indicates that subjects who explained their strategies tended to exhibit more of the preference reversal pattern than subjects who did not.

Value and Probability Advice. Before asking whether advice to attend more to the values of the prizes or to the probabilities in the gambles eliminated the preference reversal pattern, let us see whether these advice conditions affected subjects' relative attention to the values and probabilities in the gambles. Advice to attend to the value dimensions of the gambles would be expected to increase the subjects' preference for the Hi-V bet; Probability Advice would make the Hi-P bet more preferred. Table 2 shows that for each of these advice conditions, the Hi-P preference proportion moved in the expected direction. After individuals' proportion data had been rescaled by the arcsine transformation (Cohen & Cohen, 1975, p. 255), ANOVA contrasts showed that this move in opposite directions was significant, $F(1, 61) = 4.52$. Similarly, the p parameter of the Two-Error-Rate model, which indicates the overall preference for the Hi-P bet, decreased from .25 to .15 with the Value Advice and increased from .42 to .50 with the Probability Advice, as would be expected (Table 4).

Although Value and Probability Advice changed the subjects' evaluations sufficiently to affect the preference rates on each of the evaluation methods, there is little evidence for a change in preference reversal pattern. There were small decreases in the rates of predicted and unpredicted reversals alike (Table 3). The Two-Error-Rate model showed only weak evidence for the preference reversal pattern before the influence, and even weaker evidence afterwards (Table 4).

There is thus no evidence that Value or Probability Advice lessens the occurrence of the preference reversal phenomenon, despite the effect each has on the subjects' evaluations. This test is inconclusive, however, because these particular groups of subjects did not exhibit the preference reversal very strongly at the outset.

Analytic and Intuitive Advice. Before asking whether the Analytic Advice eliminated the preference reversal pattern and whether the Intuitive Advice strengthened it, we must ask whether these influences affected the subjects' strategies. One index of an effect is a change in response time: Analytic Advice should make subjects spend more time on their answers; Intuitive Advice, less. There were, however, sex differences in response time changes. On the choice task, both sexes slowed down after Analytic Advice, as expected (males from 16 to 31 seconds per choice, females from 21 to 22 seconds); but the females speeded up after Intuitive Advice (from 19 to 14 seconds) whereas males slowed down (from 33 to 42 seconds). On the limit price setting task, though both sexes speeded up after advice to evaluate intuitively (males from 43 to 35 seconds; females from 27 to 18 seconds), only males took longer after advice to evaluate analytically (from 25 to 33 seconds); females sped up drastically (from 45 to 23 seconds.) It is as if many men resisted the advice to think intuitively while many women resisted the advice to think analytically. Thus, it is not clear whether the Analytic and Intuitive Advice conditions had the expected effect on all subjects' evaluation strategies.

With respect to their effects on preference reversals, there was weak evidence for the expected result that Analytic advice would decrease preference reversals while Intuitive Advice increased them. By the evidence of the change in predicted and unpredicted reversal rates (Table 3), there was an increase in extent of preference reversal in each condition. By the evidence of the Two-Error-Rate model method, subjects in the Analytic Advice condition exhibited preference reversal patterns that are untestable both before (G) and after (H) the advice. The pattern for subjects with Intuitive Advice moved from no preference reversals (D) to reversals for which no test of statistical significance (H) can be made. Thus it seems that Intuitive Advice produced a stronger preference reversal pattern, as predicted, but Analytic Advice did not produce a weaker one.

In summary, only 3 of the 12 conditions exhibited no preference reversal pattern, 2 before the advice and 1 after. With respect to the seriousness of the preference reversal pattern, Probability Advice made it disappear, Practice produced a decrease (though not a disappearance), Value and Analytic Advice produced no change, and Intuitive Advice and Discussion produced increases.

Analysis in Terms of Individuals

The use of data aggregated across subjects has been criticized for obscuring the real effects at the level of the individual (e.g., Hammond, 1976).

Although the aggregate analyses of our data showed that few influences decreased the extent of preference reversal for the group of subjects, analysis at the level of individuals allows us to see whether the manipulations worked for any individuals. If individuals are stably different from each other, we can ask whether the influence conditions caused significant numbers of them to have different preference reversal patterns. Did the Value, Probability, or Analytic Advice make subjects become consistent? Did the Intuitive Advice make them inconsistent?

Because each subject evaluated only six gamble pairs in each set, the Reversal Rate Comparison method and Two-Error-Rate Model method would provide very unstable measures of an individual's preference reversal pattern. Therefore a simpler measure was used. Each individual's preference pattern in a set of gambles was classed as one of four categories:

0 Subject had no inconsistent responses. On none of the six pairs of Hi-P bet and Hi-V bet gambles was the subject's preference as indicated by choice different from his or her preference as indicated by buying limit prices.

= Subject had an equal number of predicted reversals (chose Hi-P bet but offered more for Hi-V bet) and unpredicted reversals (chose Hi-V bet but offered more for Hi-P bet).

> Subject had more predicted than unpredicted reversals (the classic Preference Reversal Phenomenon pattern).

< Subject had more unpredicted than predicted reversals.

Table 5 shows how many of the 109 subjects fell into each of these categories before and after the influence manipulations, over all conditions and Table 6 shows the results for each condition separately. The marginals on the right of Table 5 indicate that before the influence, 45 subjects had more predicted than unpredicted reversals, while 30 subjects had more unpredicted reversals. The bottom marginals show that after the influence, 48 subjects had more predicted reversals and only 19 had more unpredicted. Thus the individual data support the conclusion that the classic preference reversal phenomenon persisted despite manipulations designed (except for the Intuitive Advice) to get rid of it.

Insert Tables 5 and 6 about here

The cells in Table 5 indicate how many subjects had particular patterns before and after the influence. Thus, 9 subjects changed from having mainly the unpredicted reversal pattern before the influence to exhibiting mainly the predicted reversal pattern after the influence; 24 subjects exhibited the predicted reversal pattern both before and after the influence.

Is there sufficient evidence for stable individual differences to justify this individual analysis? We can check whether subjects tended to have the same pattern both times they were observed, that is, before and after the influence. This is a conservative test of stability, because the influences were expected to change the subject's reversal pattern. Weak evidence for stability, under these conditions, would justify individual analysis. Each row in Table 5 is tested (binomial test) to see whether significantly more subjects than chance had the same pattern on the second set of gamble pairs as on the first. The method of obtaining the expected frequencies for this test is explained in Appendix C. To apply the binomial test, let r = the number of stable subjects observed (o); n = the number of subjects in the row (N); and p = the proportion of subjects expected to be stable if they were responding randomly (e/N). The probability that the number of stable subjects could be at least as different from e as o, and in the same direction, is given by the cumulative binomial distribution

$$\sum_{x = r}^{n \text{ or } 0} \binom{n}{x} p^x (1 - p)^{n - x}$$

(Mosteller, Rourke, & Thomas, 1970). The summation goes from r to n or to 0, depending on whether o is greater or less than e. The logic of this test is two-tailed but the cumulative binomial distribution gives the probability for just the one tail of the binomial distribution on which the observation appeared, so the conventional 95% confidence limit is at $p = .025$.

Table 7 shows that the subjects who had no reversals (0 category) were stable to a highly significant degree ($p < .0001$), the subjects with the predicted reversal pattern (> category) were marginally stable ($p = .033$), and those with the other two patterns were not stable. Overall, the stability ($p = .050$) did not reach statistical significance. However, since the influence conditions would make people unstable, there is sufficient evidence for stable individual differences to justify individual analysis.

Insert Table 7 about here

Did the influence conditions cause people to shift among categories? In particular, did they cause significant increases or decreases in the number of subjects having reversals, and did they cause changes in the numbers of subjects having predominantly predicted or unpredicted reversals? The cells of Tables 5 and 6 that are involved in each of these patterns of transition are:

To Fewer Reversals: > to 0, < to 0, = to 0.

To More Reversals: 0 to >, 0 to <, 0 to =.

To Predicted Reversals: < to =, = to >, < to >, 0 to >.

To Unpredicted Reversals: > to =, = to <, > to <, 0 to <.

(Note that the transitions 0 to > and 0 to < are included twice; hence the sum of o's in a row of Table 8 may exceed N.) Did more, or fewer, subjects make these transitions than would be expected by chance? The number of subjects that would be expected to have these transitions by chance is derived by the method explained in Appendix C. Table 8 shows the results of binomial tests for each pattern, for each influence condition. For most of the influence conditions, and for the total set of subjects, a significant number of subjects came to have no reversals. There was no change in the number of subjects having the predicted reversal pattern, but significantly fewer subjects than would be expected shifted to the unpredicted reversal pattern.

Insert Table 8 about here

There was little differentiation among the Advice conditions and the control conditions. In each condition except Intuitive Advice a significant number of subjects came to have fewer reversals. That the Intuitive Advice did not produce such a switch is not surprising, since it was not expected to reduce preference reversals. In each condition except Intuitive and Probability Advice, significantly fewer subjects than would be expected by chance switched to having predominantly unpredicted reversals. We do not know why the Discussion and Practice control conditions had similar effects to the Value and Analytic Advice conditions.

In summary, subjects who started off with no reversals or with the classical preference reversal pattern tended to stay in the same pattern that they started with. Further, with the exception of the Intuitive Advice condition, significant numbers of subjects shifted to the pattern of having no reversals. Despite this growing subset of consistent subjects, the influence attempts did not make a significant number of subjects stop having the classical preference reversal pattern, though four of the six conditions prevented them from having the unpredicted reversals pattern.

In conclusion, analyzing data by counting how many individual subjects exhibit particular patterns gives a more complicated picture than aggregated data provides. While it supports the conclusion we drew from the aggregate data, that the preference reversal phenomenon occurs and that the influence conditions did not generally succeed in eliminating it or lessening its overall degree of occurrence, the analysis of individuals also reveals that for a subset of individuals, there is no preference reversal problem.

Discussion

The two important questions researchers can ask of the preference reversal phenomenon are how general it is and why it occurs. Its occurrence in previous studies and in the plurality of subjects in our study, despite influence conditions designed to make them become consistent, shows that it is general. The stable individual differences we observed suggest that in order to understand why the preference reversals occur, it is necessary to understand the cognitive processes used by each individual when doing each evaluation method. In this section, we will discuss the problems of measuring the preference reversal pattern and of exploring and explaining its generality. We will show what our study has contributed, and sketch the research needed for delineating its generality and explaining it.

Measuring the Existence and Extent of the Preference Reversal Phenomenon

Our replication of the preference reversal phenomenon and our finding that advice and discussion do not eliminate it are complicated by the fact that there are several methods for measuring the preference reversal pattern. Previous researchers have used methods that aggregate across subjects; of these the Reversal Rate Comparison method has been used more often than the Two-Error-Rate Model method. The present study was the first to use analysis on the individual level. For most purposes, individual analysis is preferable. We will compare the three approaches by showing what each enables us to say about the data of the group of subjects in the Value Advice condition (see Table 2).

The two aggregate methods agree in finding no major change in the extent of preference reversal pattern due to the Value advice. The Reversal Rate Comparison method (Table 3), which involves subtracting unpredicted from predicted reversal rates, indicated that this group of twenty-five subjects had a strong preference reversal pattern (a difference of .44 between the proportions of predicted and unpredicted reversals) both before and after the advice.

The Two-Error-Rate Model method (Table 4) showed that this group of subjects had a strong basic preference for the Hi-V bet (preferring it 75% of the time before the advice), that the Value advice increased their preference for it (to 85%), and that they erred more often in expressing their preferences when choosing than when pricing (eight times as often before, and twice as often after). With this method, evidence for the existence of the preference reversal pattern is weak because no statistical rejection of the Two-Error-Rate Model is involved. The existence of a preference reversal pattern is asserted solely because it is implausible that subjects should express their preferences more accurately when setting limit prices than when choosing, when they have greater confidence in their choices. However, it is possible for people to feel more confident in a procedure where they are less accurate (e.g., Hammond, Hamm, Grassia, & Pearson, 1983). Therefore, the Two-Error-Rate Model method does not indicate as strongly as the Reveral Rate Comparison method does that the Value Advice group had preference reversals.

Additionally, by the Two-Error-Rate Model method's measure the subjects in other conditions, especially Analytic Advice and Practice, had stronger preference reversal patterns than the subjects in the Value Advice condition, while by the Reversal Rate Comparison method's measure the preference reversal pattern was strongest with the subjects in the Value Advice condition.

The description emerging from the individual analysis (Tables 6, 7, and 8) is fundamentally different from what the aggregate analyses showed. Fourteen of the 25 subjects had more predicted than unpredicted reversals before the Value advice, and 10 did after the advice, although 8 subjects (increased from 3) now had no reversals. Where both aggregate methods indicated that there was no change in the group's preference reversal pattern due to the Value Advice, analysis at the level of individuals showed a significant shift in the numbers of subjects with various patterns (Table 8), a shift toward having no reversals and away from having more unpredicted reversals. Where the aggregate methods disagreed in the extent they claimed the group as a whole had a preference reversal pattern, individual analysis shows that it is not possible to speak of the group's pattern; rather, some individuals have predominantly one mode of preference reversal, some the other, and some don't have it at all.

Differences between the findings of the aggregate methods are trivial in comparison with what both these methods obscure with their general descriptions: the diversity of preference among individuals. Only one previous study, Lichtenstein and Slovic (1971), has reported the number of individuals with various patterns; subsequent researchers have reported only aggregate analyses. Because analysis on the level of individuals cuts closer to the truth than the aggregate analyses do, we recommend it for future investigations of the preference reversal phenomenon. Additionally, past results should be reanalyzed, for their comparisons between various conditions are based on aggregate measures and are likely to be misleading.

The present study is the first attempt to change the preference reversal phenomenon that has looked at individuals. Finding that the treatment influences some subjects but not others raises a new class of questions for further research: What made some subjects respond to the advice? Why were the others not reached? How could they be?

In doing individual analyses, it would be desirable to use a more precise measure of the occurrence and extent of the preference reversal pattern than was used in this study. If more evaluations of gamble pairs were made in each condition by each individual, it would be possible to use the Two-Error-Rate Model method. This would be chosen over the Reversal Rate Comparison method for the same reasons that it was preferable in the aggregate analysis (see above).

Of course, analysis of individuals' preference reversal patterns is not necessary or appropriate in all cases. It would not be necessary if the researcher knew that all the individual subjects were homogeneous, i.e., that there were no (or negligable) stable individual differences. In our study we found evidence for such stable differences; no other researchers have looked.

Individual analysis would not be appropriate if the research were concerned with reversals that occur when a group uses one method as opposed to the other (see Mowen & Gentry, 1980). For example, consider the project of developing measures of the spectrum of health/disease states in order to guide policy makers in allocating resources for the prevention or treatment of the various diseases. To develop such a 'health status index', large samples of people are questioned about their preferences for an exhaustive categorization of health states. The preference reversal issue rises because, as it happens, one of the leading health status indices has been developed by researchers who asked their informants to use magnitude estimation procedures to evaluate these health states (Kaplan, Bush, & Berry, 1979)--similar to the Pricing evaluation mode--while another leading group has asked its informants to make choices between gambles over health states (Torrance, Boyle, & Horwood, 1982)--similar to the Choice evaluation mode. If the population's evaluations of two health states should exhibit a preference reversal between these two evaluation methods, then the two leading health status indices could lead to conflicting policy recommendations. Researchers concerned with discovering whether such a preference reversal occurs would analyze the groups' preferences in the aggregate; individual analysis would not be appropriate.

However, for most of the research questions about the preference reversal phenomenon (e.g., its challenge to the assumption in classical economics that the individual is rational), analysis at the level of the individual is appropriate and, our results suggest, necessary. Therefore, further work on the preference reversal phenomenon should either measure the preference reversal pattern on the individual level or verify that this is not necessary by showing that the subjects are a homogenous group.

The Generalization Problem

Although preference reversal has unquestionably been demonstrated between the two methods, choosing between two gambles and evaluating single gambles by naming limit prices, it is important to know how often such reversals will occur in the real world. This problem can be analyzed in terms of four distinct generalization questions: generalization over subjects, over different knowledge or attitude states on the part of the same subject, over evaluation methods, and over gambles.

Generalization over subjects. Preference reversals have been demonstrated mainly with groups of college students, but also with gamblers at a Las Vegas casino (Lichtenstein & Slovic, 1973). Because even the experienced gamblers showed reversals, it has been assumed that the phenomenon will occur with any group of people who can understand the gambles and make the prescribed evaluations.

Our finding of large and somewhat stable individual differences in whether reversals occurred and in whether they were of the predicted or unpredicted type suggests that one can not state that college students generally exhibit the preference reversal phenomenon. New research must determine why some individuals have preference reversals while others are consistent.

Generalization within the individual over different belief, attitude, or skill states. Our study is the first to directly investigate the generality of preference reversal within the individual--that is, whether an individual whose preferences were inconsistent in one situation would also exhibit the inconsistency in different situations. Because we were interested in discovering conditions that would make subjects consistent, our approach to the question was not to sample broadly over various states of the subject's everyday life (cf. Brunswik, 1944). Instead, we looked at the effects of influencing the subject, by written advice and then during conversation with the experimenter, to adopt a different way of evaluating the gambles. This approach offers indirect evidence on the generalization question: If preference reversals were still to appear under conditions in which the subject was influenced to avoid them, they would surely appear in situations lacking any pressure toward consistency.

Within our narrower scope, we nevertheless wanted our influence attempt to be like the influences available for people in real life. People are not often confronted about inconsistencies between their preferences as expressed with choice and with single-item evaluations, but they are often questioned about particular features of their judgment policies. We designed our influence attempts to be similar to three of four common types of judgment-criticizing conversation:

1. Objection to the judge's attention to details of the situation, suggesting that the judge is, in effect, not weighting cues properly. In our Value and Probability Advice conditions, we told subjects to adjust the relative attention they were paying to V_W and P_W.

2. Objection to the way the judge is combining information. We did not explicitly address this in our study.

3. Objection to the judge's general approach. In the Analytic and Intuitive Advice conditions, we suggested that the subject's entire approach was too intuitive, or too analytical, and urged a general reorientation.

4. Close attention to the judge's own explanation and exploration of his or her strategies. Our Discussion condition facilitated development of the subject's understanding without giving direction.

For three reasons, we chose not to give subjects any information about their own preference reversals, or even about the possibility of reversals. First, we did not want to induce subjects to avoid preference reversals just because we had implied that they were bad.

Second, since the preference reversal phenomenon is not likely to be recognized when it occurs in real life, people would not ordinarily receive explicit feedback about it. Rather, if they changed at all it would be through partial analysis of their practice, induced by influences that were not triggered by the existence of preference reversals. Hence we designed influences that gave general rather than specific advice, which had been selected independently of the subject's actual pattern of preference reversal. This advice was given with equal vigor to the subject in both evaluation modes. All this was done without revealing the possibility or the undesirability of a preference reversal.

Third, investigation of such direct communication seemed unnecessary. We assumed that if the subjects were told about the reversals, they would both want to, and be able to, avoid being inconsistent in the future. The assumption that subjects would be motivated to avoid having preference reversals seems to be supported. When, at the end of the experiment, our subjects were shown their reversals and given the opportunity to change responses, almost all of them were disturbed enough, in retrospect, to make their responses consistent. However, other researchers have disproved the assumption that subjects would be able to prevent reversals prospectively once they knew about them. Berg, Dickhaut, and

O'Brien (1983) gave subjects the experience of losing money because of their inconsistent preferences, yet most subjects were not able to avoid making reversals when they evaluated another set of gambles. This is consistent with work on multiple cue probability learning, which has shown that 'outcome feedback', i.e., information about the consequences or correctness of one's judgments, is not very effective in inducing learning (Einhorn, 1980; Hammond, Summers, & Deane, 1973). We should not have expected that subjects who knew that they had made inconsistent judgments and who had suffered for them would then be able to avoid making them in the future. To take a further lesson from Hammond, Summers, and Deane's (1973) work, influence attempts would have been more successful had the subjects been given 'cognitive feedback', that is, had they been shown what they were doing, why it was wrong, and how to do differently.

Even though we overestimated how helpful the information we excluded from our constrained influence attempts would have been in making a subject consistent, we found that the various advice conditions were partially effective. The number of people having no reversals increased, although reversals did not disappear (in fact, there were more subjects who had the predicted kind of reversal after the advice than before). Only the Intuitive Advice decreased the number of subjects who had no reversals, as would have been expected. The lack of large differences among the advice and control conditions, however, means that it is not clear which features of our influence attempts were effective--perhaps just discussion.

In conclusion, our exploration of the generality of the preference reversal phenomenon over different states of the same individual involved changing the subjects' attention to the dimensions in the gambles (Value and Probability Advice), or their general strategy (Analytic and Intuitive Advice), or letting their understanding change through trying to justify their strategy (Discussion). Ironically, after electing not to use the most effective influences available because it was so obvious that they would work, we found that the effect of the influences we actually used was small: although some subjects became consistent, many still had preference reversals.

Had we simply sampled belief, attitude, and skill states within subjects randomly, the preference reversal phenomenon would likely have appeared very general. More effective influences would be needed to take most individuals beyond the preference reversal phenomenon's limits of generalization. To this end, future researchers should test the effect on the preference reversal phenomenon of (a) giving the subject advice specific to his or her strategy on each evaluation method, (b) teaching the subject to recognize and avoid preference reversals, and (c) supplying outcome and cognitive feedback--that is, telling subjects when they have produced inconsistent evaluations and how they have done so.

Generalization over evaluation methods. The preference reversal phenomenon may be generally defined as an inconsistency between the two most basic types of behavior for indicating preferences--choosing among a set of objects, and assigning a number to one object (see review in Hamm, 1979, Appendix 1). Indeed, several specific versions of each type of evaluation method have been used in previous demonstrations of the preference reversal phenomenon. For assigning a number to one gamble, Lichtenstein and Slovic (1971) used both Selling and Buying Limit Price methods, and Grether and Plott (1979) used Selling Limit Prices and a method in which the subject named an amount of money such that he or she was indifferent between receiving the money or the gamble. Goldstein (1983), in addition to the usual Selling Limit Price method, asked subjects to rate the gamble according to how willing they were to play it. He also added a version of the choice method in which subjects indicated for which gamble they would set a higher selling limit price.

Although the others found the same basic preference reversal pattern with their various methods, Goldstein (1983) did not find it with his new methods. One might conclude, then, that the limits of generalization of the preference reversal phenomenon have been established. However, there are five reasons why such a conclusion would be premature. First, aggregating data across subjects obscures individual differences in the generalizability of the preference reversal phenomenon across methods. Second, no preference reversal investigations except Goldstein's (1983) have used two methods of choice or of single gamble evaluation with the same subjects, which is necessary for careful exploration of generalization gradients. Third, preference reversals have usually been measured with the reversal rate comparison method, which is flawed (as we argued above). Fourth, although Goldstein (1983) did not find the typical "predicted" preference reversal pattern with the new methods, he still found preference reversals. Thus he has not located a true limit of generalization; rather, he has crossed the boundary to a distinct preference reversal phenomenon. Fifth, although preference reversals with a particular type of gamble pair may disappear under a change in evaluation method, there may still be reversals if a different kind of gamble pair is used (the next section discusses this issue). These five problems imply that the question of the limits of generalization of the preference reversal phenomenon over evaluation methods has not been fully addressed.

On a more general level, simply to identify whether the preference reversal phenomenon generalizes to new variants of the evaluation methods is to neglect the scientifically interesting aspect of generalization--the explanation of <u>why</u> the phenomenon does or does not generalize to new situations. This question has not been effectively pursued because (a) most researchers since Lichtenstein and Slovic (1971; see Slovic & Lichtenstein, 1983) have been interested mainly in whether the preference reversal phenomenon "really exists", and so they have focussed on demonstrating rather than exploring the preference reversals; (b) measurement of the extent of the preference reversal pattern has been so imprecise that discussion of generalization gradients has not been possible.

To fill this gap, we will describe the differences between the sizes of the preference reversal pattern we found using the Buying Limit Price method for evaluating single gambles, and the pattern others have found using the Selling Limit Price method.

The Becker, DeGroot, and Marschak (1964) procedure used in this study and in previous demonstrations of the preference reversal phenomenon establishes a subject's evaluation of an object through (a) eliciting a limit price, (b) determining the actual price through a lottery, and (c) exchanging if the actual price is not less favorable for the subject than the limit price. There are Buying and Selling versions of this procedure. In the Buying version, the subject names the upper limit on possible prices he or she would pay to play the gamble. In the Selling version the subject owns the right to play the gamble and names the lowest price at which he or she would sell it.

Our study used the Buying Limit Price method for evaluating single gambles. This was not done for purposes of generalization; on the contrary, we assumed the methods were equivalent and chose the Buying Limit Price method because our student subjects would typically be more familiar with buying than selling. However, in the aggregate our subjects had lower preference reversal rates (measured, for comparability, by the Reversal Rate Comparison method) than subjects in other studies. The difference between the predicted reversal rate and the unpredicted reversal rate was .20 before the influence and .30 after (averaged across conditions). Similarly, the difference for Lichtenstein and Slovic's (1971) Experiment II, the only other study using the Buying Limit Price method, was .24. This may be contrasted with the larger preference reversal scores obtained when the Selling Limit Price method was used: Lichtenstein and Slovic (1971, Experiments I and III)--.77, .45; Lichtenstein and Slovic (1973)--.72; Grether and Plott (1979; four groups using Selling Limit Prices; and two using the method of stating a price equivalent to the gamble)--.48, .60, .57, and .40; .73 and .49; Pommerehne, Schneider and Zweiful (1981; three groups using Selling Limit Prices)--.35, .40, .33; Reilly (1981; two groups with the Selling Limit Prices)--.33, .52. Reilly went on to get differential reversal rates as low as ours only after applying consistency-increasing manipulations (.24, .47, .03, .22). It can be seen that the preference reversal pattern obtained when the Buying Limit Price method is used to evaluate the single gambles is generally less than when the Selling Limit Price method is used.

How can this difference be explained? Although these two versions of the method have equal theoretical claim to express the subject's values truthfully (Farquhar, 1982; Hamm, 1979, Appendix 1-f), they have been shown in studies that were not concerned with the preference reversal phenomenon to produce systematically different numbers (Coombs, Bezembinder, & Goode, 1967; Slovic & Lichtenstein, 1968). The Selling Limit Prices are higher (cf. Crozier, 1978; Thaler, 1979). Further, Coombs, Bezembinder and Goode (1967) found that the correlation between Buying Limit Prices and the gambles' expected values (.36) is less than that between Selling Limit Prices and expected value (.70). This information suggests two reasons that these methods might produce different preference reversal patterns, one reason having to do with the gambles' ranges and the other having to do with noise.

First, the prize ranges of the gambles being evaluated interact with whether the method uses Buying or Selling Limit Prices. As Lichtenstein and Slovic (1971) note, Selling Limit Prices generally fall between the gamble's expected value (EV) and the amount that could be won (V_w), and Buying Limit Prices between EV and the amount to lose (V_l). In preference reversal studies, the stimuli have always been pairs of gambles with equal expected value, where one has a low probability of winning (P_w) and a high V_w, while the other has a high P_w and a low V_w. V_w's are above EV, of course, and V_l's are below. Consequently, in a pair of Hi-P and Hi-V gambles, the Hi-P bet with its lower V_w has a shorter range of possible Selling Limit Prices than the Hi-V bet. We need not make strong assumptions about the subjects' ability to accurately perceive the EV's of the two gambles, nor about the shape of their distributions of selling bids, to see that the average Selling Limit Price would be higher for the Hi-V bets. There would be no such range effect with the Buying Limit Price method, because the natural bottom of the range (either 0 or V_l), which would have the determining effect on Buying Limit Prices, was intentionally constructed not to be systematically related to whether the gamble is a Hi-V bet. Therefore, the prize range hypothesis predicts that with the Selling Limit Price method the Hi-V bet would consistently be given the higher Limit Price, while with the Buying Limit Price method, neither gamble would consistently be evaluated higher.

Second, as Coombs, Bezembinder, and Goode (1967) noted, the Buying Limit Price method is noisier than the Selling Limit Price method, having a lower correlation with the gambles' expected values. No matter how the extent of preference reversal is measured, added noise would make it more difficult to observe preference reversal patterns.

In summary, we have argued that there is a generalization gradient for the preference reversal phenomenon across evaluation methods, because the aggregate reversal pattern decreases when the Buying instead of the Selling Limit Price method is used. We have identified two possible causes for the decrease. Further work using within-subject comparison across evaluation methods, single-subject analysis, and the Two-Error Rate Model measure of the preference reversal pattern, with the full range of possible evaluation methods, is necessary to both establish and explain the true limits of generalization of the preference reversal phenomenon.

Generalization to other gamble pairs. Since Lichtenstein and Slovic (1971), little attention has been paid to the nature of the gamble pairs involved in the preference reversal phenomenon. Subsequent preference reversal researchers have used the six gamble pairs of their (1971) Experiment 3. Thus the preference reversal phenomenon has been repeatedly demonstrated with only a very small set of gamble pairs.

The gamble pairs Lichtenstein and Slovic (1971) introduced in their Experiment 3 were selected to have nearly equal expected values, and a very large difference in value of winning (V_w) between the Hi-V bet and the Hi-P bet. The constraint that the gamble pairs have nearly equal expected values has two important features. First, there is no normative reason to prefer one bet to another (cf. Pauker, 1981). Second, gamble pairs with equal expected values are a very small subset of the universe of possible gamble pairs. If we consider only the P_w and V_w dimensions, gambles with expected values equal to a given gamble must fall on the same hyperbolic curve. In the space of gambles defined by the three dimensions V_w, P_w, and V_l, the locus of equal gambles is a surface, again a small subset of the possibilities. Most gamble pairs people face in real life (whatever their particular content, cf. Mowen & Gentry, 1980) would not have equal expected values.

There has been no systematic exploration of the limits of generalization for gamble pairs with unequal expected value, either for groups or for individuals. In general, if people's evaluations are at all related to the norms, then when gamble pairs have unequal expected values there will be fewer preference reversals. On the other hand, if a particular subject's judgment policies for choosing and setting limit prices are known, it should be possible to select particular unequal expected value gamble pairs that have as high a chance of reversing as equal pairs do.

Lichtenstein and Slovic (1971) have explored generalization of the preference reversal phenomenon over gamble pairs with equal expected values. In their Experiment 2, 74 subjects evaluated 49 such gamble pairs. Across the pairs, the difference in V_w between the Hi-V bet and the Hi-P bet was related to the number of subjects who set a higher buying limit price for the Hi-V bet. At the same time, the difference in V_l between the gambles was related to the number of subjects who chose the Hi-V bet. Thus, the relative size of the winning prize controlled the subjects' limit prices but the relative size of the losing penalty controlled their choices.

In selecting gamble pairs for their Experiment 3, Lichtenstein and Slovic (1971) focussed on the effects of V_w (and P_w which complements it) rather than V_l, and so simply counterbalanced V_l so that it favored choice of the Hi-V bet in three of the six gamble pairs. While ignoring the V_l dimension made it easy to present the preference reversal phenomenon, it poses problems for researchers interested in exploring the limits of generalization of the phenomenon.

It is necessary for researchers exploring the influence of other factors on the preference reversal phenomenon to be aware of the role of V_1 in the gamble pairs. At least two researchers have failed to do this. Reilly (1982), in Stage 2 of his study, used 6 gamble pairs of which 4 had more attractive V_1's in the Hi-V bet. Knowing how V_1 can influence the choices, we would predict fewer than usual preference reversals here. This result is evident in Reilly's data when it is analyzed with the Two-Error-Rate Model method, but it is not so obvious with the Reversal Rate Comparison method, which he used. Reilly (1982) attributed the changes in reversal rates to his subjects' increased motivation, but the simpler explanation is that the changes are due to the changes in relative value of V_1 in his new gamble pairs.

Similarly, in the present study 6 new gamble pairs were created to complement the 6 from Lichtenstein and Slovic's (1971) Experiment III. However, V_1 was not counterbalanced among the Hi-V and Hi-P bets. Instead, 5 of the 6 new pairs had a more attractive V_1 in the Hi-V bet. Thus the Hi-V bet would be favored in the choice as well as the single gamble evaluations, and there would be fewer preference reversals. Indeed, across all conditions in our study, only 47% of the non-indifferent choices favored the Hi-V bet with Lichtenstein and Slovic's (1971) gamble pairs, while 62% did with our gamble pairs. Further, the fit of the Two-Error-Rate Model (see Figure 2) to the judgments of the Lichtenstein and Slovic gambles yielded evidence for the preference reversal phenomenon both before (Category E, McNemar Test $X^2 = 9.13$, $p < .01$) and after (Category E, $X^2 = 26.60$, $p < .001$) the influence, while the judgments of the new set of gambles had only a weak preference reversal pattern (Category B both before and after the influence). (Since set of gamble pairs was counterbalanced before and after the influence, within each influence condition, using the alternative set of gambles did not compromise the evaluation of our manipulations.) Surprisingly, when Berg, Dickhaut, and O'Brien (1983) subsequently used this new gamble set, the V_1 effect was not evident in their data.

Researchers should incorporate the V_1 dimension into their descriptions and explanations of the preference reversal phenomenon, rather than merely controlling for it. When the variance due to V_1 is "accounted for", i.e., incorporated in the explanatory model, rather than just being "uncorrelated with the important factors", i.e., counterbalanced but excluded from the model, the explanatory model is more powerful.

Further, researchers should acknowledge that the V_1 dimension is intrinsic to the preference reversal phenomenon. Explanation of the preference reversal phenomenon depends on description of individuals' strategies in each evaluation mode. Since the V_1 dimension has been present in all of the gambles with which the preference reversal phenomenon has thus far been demonstrated, every individual might have taken the V_1 into account when choosing and even when setting limit prices. There is evidence for this in our study: in the choice evaluation mode, a preliminary analysis of 80 subjects showed that for 17 of the 18 gamble pairs used in our experiment (including 6 fillers not otherwise analyzed), the majority of the subjects chose the gamble with the more attractive V_1. For the evaluation of single gambles, at least one of our subjects was influenced by the V_1 dimension. She frequently did not want to play a gamble, and so set the lowest possible buying limit price. Since the possible prices ranged from V_1 to V_w for each gamble, her single gamble evaluation was completely determined by V_1 whenever she used this strategy.

Recognition of the V_1 dimension's role makes it sensible to consider an expanded set of possible choice strategies, such as selecting the gamble whose worst possible outcome (V_1) is more attractive (maximin), or selecting the gamble that is better on two out of the three dimensions. It is plausible that subjects would use such strategies; this is bound to frustrate the researcher who attempts to explain the preference reversal phenomenon only in terms of different V_w/P_w weight ratios between modes.

To summarize, the limits of generalizability of the preference reversal phenomenon over all possible gamble pairs are not known, because it has been investigated only with gamble pairs that have equal expected values. Unequal gamble pairs would probably show fewer preference reversals, unless they had been selected to take advantage of the differences in the individual subject's judgment policies on the two modes. Among gamble pairs with equal expected values, the V_1 dimension is an important factor in controlling the choices and hence in whether there is reversal between choices and single gamble evaluations.

Summary of generalization discussion. Considering the extent to which the preference reversal phenomenon generalizes over these aspects will help researchers keep its importance in perspective. Additionally it provides a systematic framework within which researchers may seek explanations for the preference reversal phenomenon. To summarize:

1. Generalization over subjects. Examination of individual subjects' reversal patterns shows that the preference reversal pattern is not as general over subjects as had previously been thought. Some people do not have reversals, and others have predominantly the unpredicted reversal pattern.

2. Generalization within the individual over different belief, attitude, or skill states. We sought to move subjects beyond the limits of generalization by inducing states in which we assumed subjects would no longer have preference reversals. Since the phenomenon was more

general over states than we had expected, although some people came to have no reversals, the preference reversal phenomenon may be very general over unselected changes in the subject's state.

Generalization over different evaluation methods. Fewer reversals occur with the buying than the selling version of the limit price method. Two possible causes for this limitation are that the range of the gambles influences buying prices differently than selling prices, and that the buying price method is noisier than the selling price method.

4. Generalization over gambles. The gamble pairs with which the preference reversal phenomenon has been demonstrated have equal expected values. Generalization to gambles with large EV differences is likely to be limited. The V_1 dimension also influences the reversals, thus offering limits to generalization that have not been systematically explored.

Explaining the Preference Reversal Phenomenon

Researchers have not been able to produce a convincing general explanation of preference reversals. What people do in each evaluation mode to produce the reversals is not known. Goldstein's work (1983) makes clear the full extent of this failure. By manipulating response mode (choice versus single gamble evaluation) and attribute (selling limit price versus willingness to play) independently, he showed that both the usual reversal pattern (predicted on the basis of Lichtenstein & Slovic's, 1971, explanation) and the opposite, unpredicted reversal pattern could be found between choice and single gamble evaluations. These results defy the available explanations.

The lack of adequate explanation may be attributed to the interests motivating the preference reversal phenomenon research of the last decade and to the methods that have been used. Many researchers have been interested in disproving the existence or belittling the importance of the reversals, a tendency decried by Slovic and Lichtenstein (1983). They have sought manipulations that would make the preference reversals disappear, studied groups instead of individuals, and used only gross measures of the preference reversal pattern. When the reversal pattern does not completely disappear, they can not gauge how far along the generalization gradient they have moved it. They have also failed to recognize the usefulness of exploring generalization gradients in explaining the preference reversal phenomenon. When a factor that controls generalization is identified, theories must account for its effect. The biggest problem, as our study has made clear, is that explaining preference reversal has been impossible because researchers have aggregated over their subjects, and thereby obscured large individual differences, not only in preference reversal pattern but also in evaluation strategy on each mode.

In the light of these criticisms, we propose the following method for future research:

1. Decide which subjects, which subject states of belief, attitude or skill, which evaluation methods, and which kind of gambles are of interest.

2. Describe each subject's evaluation strategy for each evaluation mode. Use techniques such as regression analysis for single gamble evaluation (see Slovic & Lichtenstein, 1968) and ordered metric analysis for choice (see McClelland & Coombs, 1975). These techniques need to be applied to evaluations of a broad sample from the space of possible gambles, not solely to gambles with equal expected value.

3. Predict from these descriptions which gamble pairs will reverse, and test these predictions on the subject.

4. If predictions are confirmed, accept the pair of strategy descriptions as an explanation of this individual's preference reversals between these two evaluation modes in this psychological state.

This method would provide an explanation for a particular individual's preference reversals. To explore the types of generalization, it could be used (a) with a variety of subjects (b) over changes in the subject's state effected, for example, by outcome feedback or cognitive feedback, (c) with various pairs of evaluation modes, and (d) with different classes of gambles. Such an approach, which couches explanations in terms of individuals' strategies in each evaluation mode, could easily solve the problems of explaining the preference reversal phenomenon and establishing its limits of generalization.

References

Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single response sequential method. Behavioral Science, 9, 226-232.

Berg, J. E., Dickhaut, J. W., & O'Brien, J. R. (1983). Preference reversal and arbitrage (Working Paper 1983-2). Minneapolis: University of Minnesota, Graduate School of Management.

Brunswik, E. (1944). Distal focussing of perception: Size-constance in a representative sample of situations. Psychological Monographs, No. 254.

Brunswik, E. (1956). Perception and the representative design of psychological experiments. Berkeley: University of California Press.

Camilli, G., & Hopkins, K. D. (1978). Applicability of chi-square to 2 X 2 contingency tables with small expected cell frequencies. Psychological Bulletin, 85, 163-167.

Camilli, G., & Hopkins, K. D. (1979). Testing for association in 2 X 2 contingency tables with very small sample sizes. Psychological Bulletin, 86, 1011-1014.

Cohen, J., & Cohen, P. (1975). Applied multiple regression/correlation analysis for the behavioral sciences. Hillsdale, NJ: Lawrence Erlbaum Associates.

Coombs, C. H., Bezembinder, T. G. G., & Goode, F. M. (1967). Testing expectation theories of decision making without measuring utility or subjective probability. Journal of Mathematical Psychology, 4, 72-103.

Crozier, W. R. (1978). Evaluating the worth of gambles. _British Journal of Psychology_, _69_, 179-183.

Einhorn, H. (1980). Overconfidence in judgment. _New Directions for Methodology of Social and Behavioral Science_, _4_, 1-16.

Farquhar, P. H. (1982). _Utility assessment methods_ (Working Paper 81-5). Davis: University of California, Graduate School of Administration.

Goldstein, W. M. (1983, March). _Judgment versus choice: The preference reversal phenomenon_. Paper presented at the meeting of the Public Choice Society, Savannah, Georgia.

Grether, D. M., & Plott, C. R. (1979). Economic theory of choice and the preference reversal phenomenon. _The American Economic Review_, _69_, 623-638.

Hamm, R. M. (1979). _The conditions of occurrence of the preference reversal phenomenon_. Unpublished doctoral dissertation, Harvard University. (University Microfilms No. 80-12,829).

Hammond, K. R. (1978). _Psychology's scientific revolution: Is it in danger?_ (Tech. Rep. No. 211). Boulder: University of Colorado, Center for Research on Judgment and Policy.

Hammond, K. R., McClelland, G. H., & Mumpower, J. (1980). _Human Judgment and Decision Making: Theories, Methods, and Procedures_. New York: Praeger.

Hammond, K. R., Summers, D. A., & Deane, D. H. (1973). Negative effects of outcome-feedback in multiple-cue probability learning. _Organizational Behavior and Human Performance_, _9_, 30-34.

Havens, L. L. (1976). Participant observation. New York: Jason Aronson, Inc.

Hays, W. L. (1973). Statistics for the social sciences (2nd ed.). New York: Holt, Rinehart, and Winston, Inc.

Kaplan, R. M., Bush, J. W., & Berry, C. C. (1979). Health status index: Category rating versus magnitude estimation for measuring levels of wellbeing. Medical Care, 17, 501-525.

Krantz, D. H., & Tversky, A. (1965). A critique of the applicability of cardinal utility theory (Tech. Rep. MMPP 65-4). Ann Arbor: University of Michigan, Michigan Mathematical Psychological Program. (Cited in Becker, G. M., & McClintock, C. G. (1967). Value: Behavioral decision theory. Annual Review of Psychology, 18, 239-286.)

Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. Journal of Experimental Psychology, 89, 46-55.

Lichtenstein, S., & Slovic, P. (1973). Response-induced reversals of preference in gambling: An extended replication in Las Vegas. Journal of Experimental Psychology, 101, 16-20.

Lichtenstein, S., Slovic, P., & Zink, D. (1969). Effect of instruction in expected value on the optimality of gambling decisions. Journal of Experimental Psychology, 79, 236-240.

Lindman, H. (1971). Inconsistent preferences among gambles. Journal of Experimental Psychology, 89, 390-397.

McClelland, G. H., & Coombs, C. H. (1975). ORDMET: A general algorithm for constructing all numerical solutions to ordered metric structures. Psychometrika, 40, 269-290.

Mosteller, F., Rourke, R. E. K., & Thomas, G. B., Jr. (1970). Probability with statistical applications (2nd ed.). Reading, MA: Addison-Wesley.

Mowen, J. C., & Gentry, J. W. (1980). Investigation of the preference-reversal phenomenon in a new product introduction task. Journal of Applied Psychology, 65, 715-722.

Pauker, S. G. (1981). The toss-up. New England Journal of Medicine, 305, 1467-1469.

Pommerehne, W. W., Schneider, F., & Zweifel, P. (1982). Economic theory of choice and the preference reversal phenomenon: A reexamination. The American Economic Review, 72, 569-575.

Raiffa, H. (1968). Decision analysis: Introductory lectures on choices under uncertainty. Reading, MA: Addison-Wesley.

Reilly, R. J. (1982). Preference reversal: Further evidence and some suggested modifications in experimental design. The American Economic Review, 72, 576-584.

Slovic, P., & Lichtenstein, S. (1968). Relative importance of probabilities and payoffs in risk-taking. Journal of Experimental Psychology Monographs, 78(Pt. 2).

Slovic, P., & Lichtenstein, S. (1983). Preference reversals: A broader perspective. The American Economic Review, 73, 596-605.

Thaler, R. (1980). Toward a positive theory of consumer choice. Journal of Economic Behavior and Organization, 1, 39-60.

Torrance, G. W., Boyle, M. H., & Horwood, S. P. (1982). Application of multi-attribute utility theory to measure social preferences for health states. Operations Research, 30, 1043-1069.

Author Notes

This paper is based on the author's Ph.D. dissertation (Hamm, 1979). Help at all stages from Harold Bursztajn is gratefully acknowledged. Thanks are due to dissertation advisors Duncan Luce, Reid Hastie, Robert Rosenthal, and Richard Herrnstein, to Kenneth Hammond for emphasizing individual analysis, and to Janet Grassia for good editing. Data collection was supported by NSF Dissertation Assistance Award # 78-01977.

Address reprint requests to Robert M. Hamm, Center for Research on Judgment and Policy, Institute of Cognitive Science, Campus Box 344, University of Colorado, Boulder CO 80309.

Footnotes

[1]These conditions were suggested by Harold Bursztajn.

[2]In addition to the author, David Alwin, Harold Bursztajn, Marlene Kliman, Michael Law, and Janet Wilson served ably as experimenters.

[3]The timesharing PDP-9 computer of the Computer Based Laboratory, Department of Psychology and Social Relations, Harvard University, was used. hanks are due to Scott Bradner, Director of CBL, for advice and cooperation on this project.

[4]Experience showed that the indifference concept required further clarification, for some subjects used it when they did not like either bet. After this confusion was recognized, experimenters would supplement the written instructions with the verbal comment: "If you don't like either gamble, remember that you have to take one of them. Pick the one you would find less unattractive."

[5]We do not include in the denominators of these terms the cases where equal prices were offered, as earlier authors did (e.g., Grether & Plott, 1979; Mowen & Gentry, 1980; Reilly, 1982), because it dilutes the contrast of interest.

[6]The formula for $p(1 - p)$ given by Lichtenstein and Slovic (1971) was inadvertently wrong by a factor of -1.

Table 1

Basic Preference Reversal Data Pattern

Pattern for presenting data concerning preferences for pairs of gambles (Hi-P and Hi-V) revealed by two evaluation methods (Choice and Limit Price Setting). Parameters for the Two-Error-Rate Model.

		GAMBLE FOR WHICH A HIGHER LIMIT PRICE WAS SET		
		Hi-P bet	Hi-V bet	Row Sums
GAMBLE CHOSEN	Hi-P bet	A $pr's' + p'rs$	B $pr's + p'rs'$	A + B $pr' + p'r$
	Hi-V bet	C $prs' + p'r's$	D $prs + p'r's'$	C + D $pr + p'r'$
	Column Sums	A + C $ps' + p's$	B + D $ps + p's'$	A + B + C + D 1

p = rate of preference for Hi-P bet

r = choice error rate

s = limit price error rate

$p' = (1 - p)$

$r' = (1 - r)$

$s' = (1 - s)$

$A + B + C + D = 1$

Table 2

Preference Reversal Data: Aggregate Analysis

Distributions of answers in the Before and After stimulus sets, for each of the six influence conditions. "P>V" indicates that a higher limit price was set for the Hi-P bet than for the Hi-V bet.

		Before Influence			After Influence		
		Pricing Preference			Pricing Preference		
Condition	Choice	P>V	V>P	Total	P>V	V>P	Total
Value	Hi-P	.14	.25	.39	.11	.17	.28
Advice	Hi-V	.10	.39	.49	.10	.49	.59
N = 25	Total	.24	.64	.88	.21	.66	.87
Probability	Hi-P	.24	.19	.43	.31	.18	.49
Advice	Hi-V	.14	.33	.47	.11	.28	.39
N = 24	Total	.38	.52	.90	.42	.46	.88
Analytic	Hi-P	.26	.14	.40	.18	.21	.39
Advice	Hi-V	.17	.31	.48	.06	.36	.42
N = 12	Total	.43	.45	.88	.24	.57	.81
Intuitive	Hi-P	.29	.17	.46	.24	.17	.41
Advice	Hi-V	.13	.31	.44	.08	.39	.47
N = 12	Total	.42	.48	.90	.32	.56	.88
Discussion	Hi-P	.19	.16	.35	.20	.22	.42
	Hi-V	.19	.31	.50	.08	.33	.41
N = 24	Total	.38	.47	.85	.28	.55	.83
Practice	Hi-P	.24	.25	.49	.24	.13	.37
	Hi-V	.06	.39	.45	.01	.44	.45
N = 12	Total	.30	.64	.94	.25	.57	.82
All	Hi-P	.22	.20	.42	.21	.18	.39
Conditions	Hi-V	.13	.34	.47	.08	.38	.46
N = 109	Total	.35	.54	.89	.29	.56	.85

Table 3

Preference Reversal Data: Aggregate Analysis

Extent of preference reversals in each condition, before and after manipulations, measured with the conditional reversal method. Proportion of Hi-P Reversals = B / (A + B). Proportion of Hi-V Reversals = C / (C + D). Each subject evaluated 6 pairs of gambles.

INFLUENCE CONDITIONS	Before Influence			After Influence			Change		
	Hi-P Reversals	Hi-V Reversals	Dif	Hi-P Reversals	Hi-V Reversals	Dif	Hi-P Reversals	Hi-V Reversals	Dif
VALUE N = 25	.64	.20	.44	.61	.17	.44	-.03	-.03	0
PROBABILITY N = 24	.44	.30	.14	.37	.28	.09	-.07	-.02	-.05
ANALYTIC N = 12	.35	.35	0	.54	.14	.40	.19	-.21	.40
INTUITIVE N = 12	.37	.30	.07	.42	.17	.25	.05	-.13	.18
DISCUSSION N = 24	.46	.38	.08	.52	.20	.32	.06	-.18	.24
PRACTICE N = 12	.51	.13	.38	.35	.02	.33	-.16	-.11	-.05
ALL CONDITIONS N = 109	.48	.28	.20	.48	.18	.30	-.02	-.11	.09

Table 4

Two Error Rate Model Method's Measures of Preference Reversal Pattern

Parameters of the Two-Error-Rate Model's fit to the Before and After Data, and its category, for each of the six Advice Conditions and for the total.

		BEFORE					AFTER				
Condition	N	p(1-p)	p	r	s	Cat.	p(1-p)	p	r	s	Cat.
Value	25	.187	.249	.387	.048	B	.128	.151	.244	.129	B
Probability	24	.243	.419	.362	.017	B	>.25	~.5	~	~	D
Analytic	12	.248	.458	-.044	.364	G	.239	.394	.413	-.457	H
Intuitive	12	>.25	~.5	~	~	D	.239	.396	.336	-.157	H
Discussion	24	.223	.337	.229	.337	A	>.25	~.5	~	~	E
Practice	12	>.25	~.5	~	~	E	.236	.380	.297	-.312	H
All Subjects	109	.238	.392	.370	.005	B	.233	.370	.341	-.112	H

Table 5

Preference Reversal Data: Individual Analysis, All Conditions

Changes in individual consistency patterns. Number of subjects who had particular patterns of inconsistencies between choices and bids, before and after treatment, for all conditions.

			AFTER				
			No Reversals 0	Equal Numbers of Reversals =	More Predicted Reversals >	More Unpredicted Reversals <	Total Before
BEFORE	No Reversals	0	5	2	6	1	14
	Equal Numbers	=	6	3	9	2	20
	More Predicted	>	12	3	24	6	45
	More Unpredicted	<	5	6	9	10	30
	Total After		28	14	48	19	109

Table 6

Preference Reversal Data: Individual Analysis, Each Condition

Number of subjects with each inconsistency pattern for each of the six influence conditions.

Value Advice

BEFORE	AFTER 0	=	>	<	
0	1		2		3
=	1	1	1	2	5
>	6		6	2	14
<		1	1	1	3
	8	2	10	5	25

Probability Advice

BEFORE	AFTER 0	=	>	<	
0	1		2		3
=	2	2	3		7
>		2	3	3	8
<		1	2	3	6
	3	5	10	6	24

Analytic Advice

BEFORE	AFTER 0	=	>	<	
0					0
=	1		1		2
>	3		2		5
<	1	2	2		5
	5	2	5	0	12

Intuitive Advice

BEFORE	AFTER 0	=	>	<	
0		2		1	3
=			1		1
>			3		3
<	1	2	1	1	5
	1	4	5	2	12

Discussion

BEFORE	AFTER 0	=	>	<	
0	1		1		2
=	1		2		3
>	1	1	6	1	9
<	2		3	5	10
	5	1	12	6	24

Practice

BEFORE	AFTER 0	=	>	<	
0	2		1		3
=	1		1		2
>	2		4		6
<	1				1
	6	0	6	0	12

Table 7

Individual Analysis: Tests of Stability of Preference Reversal Pattern

Binomial tests of the number of subjects showing stability of each category and of all categories, across all conditions. Expected (e) and observed (o) numbers of stable subjects are given; e is calculated by the procedure described in Appendix C. The probabilities represent the chance that the number of stable subjects would be at least as different from e as o is.

		Category (Row in Table 5)				
		0	=	>	<	All Categories
Total Number of Subjects, N		14	20	45	30	109
Number of Stable Subjects	Expected, e	0.22	4.20	17.42	11.62	33.46
	Observed, o	5	3	24	10	42
Probability of o being this extreme or more		<.0001	.738	.033	.684	.050

Table 8

Individual Analysis: Measures of Change in Preference Reversal Pattern

Binomial tests of the number of subjects showing various transition patterns for each condition and across all conditions. Expected (e) and observed (o) numbers of subjects are given for the specified transitions; e is calculated by the procedure of Appendix C. The probabilities represent the chance that the number of subjects showing each pattern would be at least as different from e as o is.

		Transition Pattern				
Condition		To Fewer Reversals	To More Reversals	To Predicted Reversals	To Unpredicted Reversals	Stable
Value	e	0.34	2.87	4.89	11.46	7.68
N = 25	o	7	2	5	4	9
Probability		<.0001***	.441	.558	.002***	.352
Probability	e	0.33	2.87	7.46	8.65	6.94
N = 24	o	2	2	8	5	9
Probability		.043	.440	.482	.087	.236
Analytic	e	0.19	0	3.76	3.76	4.29
N = 12	o	4	0	5	0	2
Probability		<.0001***	+	.312	.011*	.139
Intuitive	e	0.14	2.87	4.53	3.34	3.35
N = 12	o	0	3	4	1	4
Probability		.869	.575	.503	.112	.443
Discussion	e	0.34	1.92	7.91	7.31	8.02
N = 24	o	4	1	6	2	12
Probability		<.0001***	.417	.276	.011*	.069
Practice	e	0.14	2.87	2.53	5.52	3.18
N = 12	o	4	1	2	0	6
Probability		<.0001***	.180	.520	<.0001***	.070
All	e	1.48	13.41	31.08	40.04	33.46
Conditions	o	23	9	30	12	42
N = 109						
Probability		<.0001***	.124	.457	<.0001***	.050

+ The test is not applicable because no subjects met the precondition for the transitions.
* exceeds 95% confidence limits (i.e., < .025)
** exceeds 99% confidence limits (i.e., < .005)
*** exceeds 99.9% confidence limits (i.e., < .0005)

Figure Captions

Figure 1. Design of the experiment. Each subject did all three sets of gambles.

Figure 2. Flow Chart determining possible conclusions about the fit of the Two-Error-Rate Model to the data. Categories A, B, and C involve technically acceptable Two-Error-Rate Models. However, categories B and C are increasingly implausible. Categories D through H do not have technically acceptable parameters. Category D is not significantly different from the Two-Error-Rate model, while Category E is. No test is available for evaluating the seriousness of the unacceptability of the parameter fit for Categories F, G, or H.

Figure 1

First Set of Gambles	Second Set of Gambles			Third Set of Gambles
12 pricings, 6 choices or 6 choices, 12 pricings	6 pricings, 3 choices S's in all conditions but Practice start discussion	Computer advice given to S's in the four advice conditions	6 pricings, 3 choices	12 pricings, 6 choices of 6 choices, 12 pricings

Figure 2

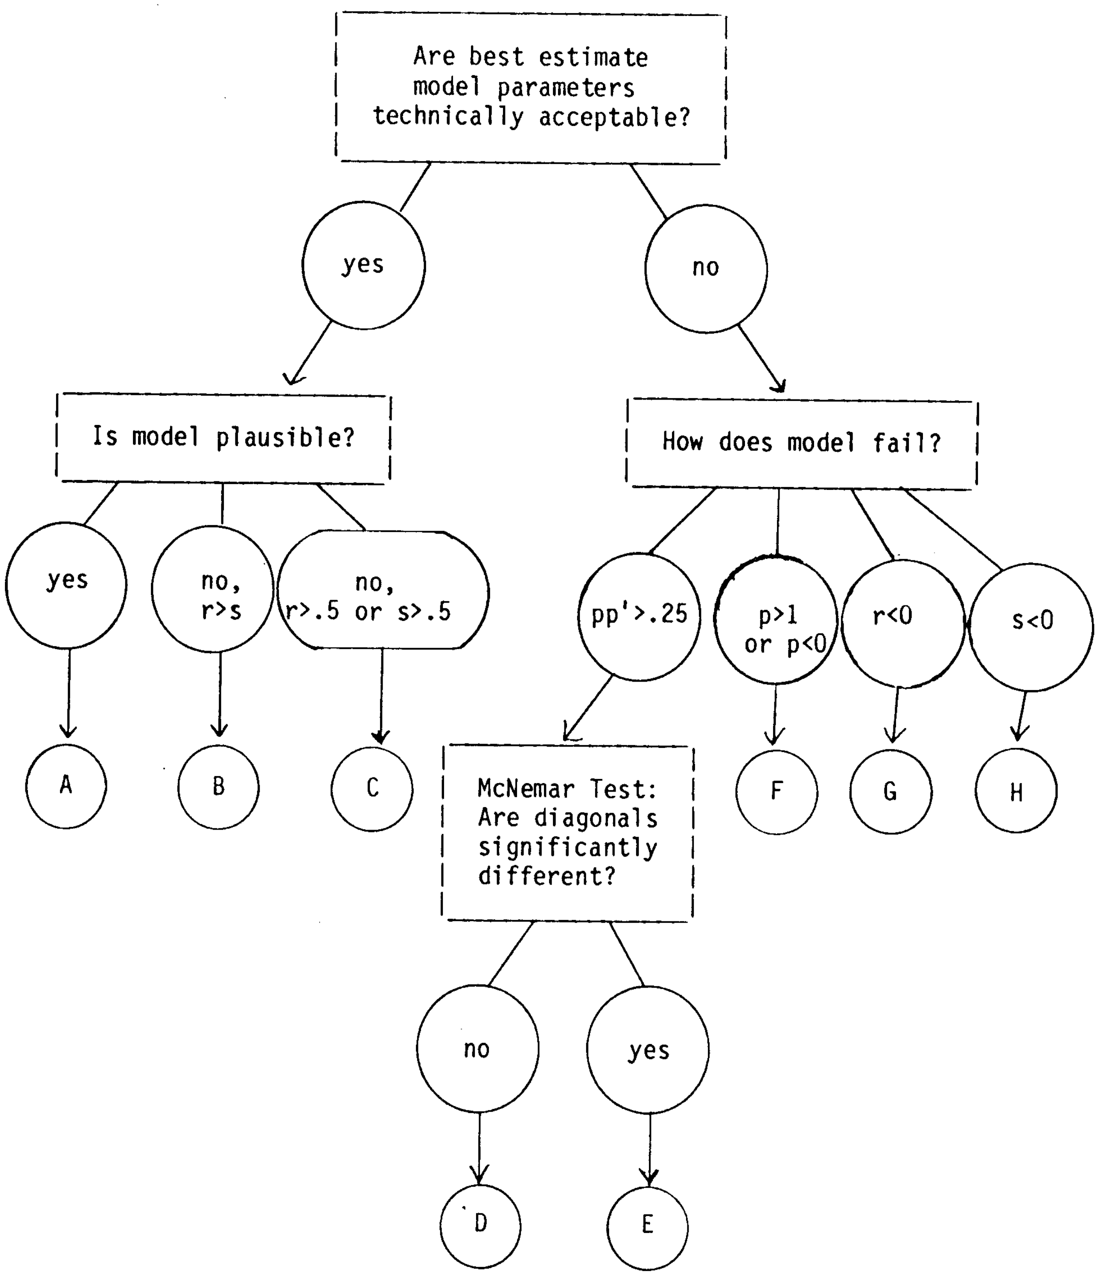
Are best estimate model parameters technically acceptable?
yes
no
Is model plausible?
How does model fail?
yes
no, r>s
no, r>.5 or s>.5
pp'>.25
p>1 or p<0
r<0
s<0
A
B
C
McNemar Test: Are diagonals significantly different?
F
G
H
no
yes
D
E

Appendix A

Instructions for Buying Limit Price Evaluation Mode

In this section, you will decide the highest amount of money you would be willing to pay in order to play each of the gambles. Consider this example:

$10.00 XXXXXXXXXXXXXXXXXXXXXXXX XXXXXXXXXXXX -$1.00

If you want to play this gamble, you have to buy a ticket. You won't get the ticket price back, no matter how the gamble turns out. Then, if you win the gamble, you receive ten dollars. If you lose it, you pay out one dollar. "You pays your money and you takes your chances."

We would like you to say the very highest price that you would be willing to pay for a ticket. This means that a price any higher than the one you named would be just too much.

To help yourself think of this price limit, imagine that every possible price between the amount you would win (i.e., $10.00) and the amount you would lose (i.e., -$1.00) is written onto a slip of paper and put into a hat. Each price is only written once, so there would be one each of "$5.21", "$9.99", "0.00", "-0.79", etc. After you name the highest price you would pay for a ticket for the gamble, a slip is drawn from the hat. If the price on the slip is higher than your price limit, then you would not play. But if it is lower, you get to buy the ticket for the price named on the slip and play the gamble.

Sometimes, you might find the gamble so unattractive that you would not play it, even if you could for free. If so, think of the smallest amount of money that someone would have to give you before you would be willing to play the gamble. Write a minus sign ("-") in front of this answer, to show that you would have to be paid to take the ticket. Remember, though, that the lower you set your limit, the smaller are your chances of being able to play the gamble, because more of the slips will be higher than your limit.

You might be tempted to treat this as a bargaining situation and try to get to play for as cheaply as possible by offering less for the gamble than you are really willing to pay. However, in this situation it is in your best interest to be accurate; that is, the best thing you can do is be honest. If the price limit you state is too high or too low, then you are passing up opportunities that you prefer. For example, suppose that you would be willing to pay $4.21 but you state that $2.82 is your limit. If the ticket price drawn from the hat is between the two (e.g., $3.16) you would have to forego playing the bet, even though you would have been willing to pay this price. Similarly, if you say $6.05 when $4.21 is really the highest you would be happy paying, then if the slip from the hat says $5.00 you would have to pay more than you wanted to.

Please read these instructions again and, for practice, decide what your highest price limit would be for the above gamble. Write that here _______. Take as long as you need.

Appendix B

Advice Given to Subjects in Analytic Advice Condition

On the basis of your answers on the first set of gambles, the computer has selected these statements of advice for you. Please indicate after each one whether the advice is relevant for you.

23. You are using your heart rather than your mind. In this world it doesn't pay to neglect your analytic skills.

    Applies 7 6 5 4 3 2 1 Doesn't Apply

29. It is not easy for you to resist the feelings you have about these gambles. Feelings may be important, but you also have to think.

    Applies 7 6 5 4 3 2 1 Doesn't Apply

37. You should consider the possibility that it is easy to have an intuition about a gamble, but that does not convey much information about whether a gamble is really good.

    Applies 7 6 5 4 3 2 1 Doesn't Apply

77. Some people would scoff at you for taking a rational approach on a task like this, but you shouldn't let that get in the way of your doing your best.

    Applies 7 6 5 4 3 2 1 Doesn't Apply

Appendix C

Individual Analysis: Determining Expected Number of Subjects Having Each Pattern

This appendix describes the method used to determine how many subjects would be expected to have each of the patterns defined above: 0, =, >, and <. Each pattern represents the subject's performance on 6 gamble pairs. On each pair the subject could be consistent or have a predicted or unpredicted reversal. The numbers of predicted (i) and unpredicted (j) reversals a subject could have when evaluating 6 gamble pairs, and the resulting pattern, are shown in Table C-1.

Insert Table C-1 about here

The probability of a subject having each pattern is determined by the sum of the probabilities of the cells in Table C-1 that represent the pattern. The probability of each cell is determined by the number of possible ways that the subject could have i predicted and j unpredicted reversals out of 6 gamble pairs. There is 1/2 chance of having no reversals on a particular gamble pair, and 1/4 chance of having each kind of reversal. Assuming the pairs are independent, the probability of having no reversals (the upper left cell in Table C-1) is $(1/2)^6$ or .0156. There are six possible patterns with only one predicted reversal: it could be the first, the second, ..., or the sixth pair that reversed. The chance of each of these six is $(1/4) * (1/2)^5$ or .0078, so the probability of one predicted reversal and no unpredicted reversals is .0469.

More generally, the probability of the (j, i) cell in Table C-1, having i predicted reversals and j unpredicted reversals out of N gamble pairs, is:

$$\binom{N}{i} * (1/4)^i * (3/4)^{(N-i)} * \binom{N-i}{j} * (1/3)^j * (2/3)^{(N-i-j)}$$

where the binomial coefficient, $\binom{N}{i}$, is the number of ways there could be i predicted reversals out of N gamble pairs; $(1/4)^i$ is the probability of a predicted reversal; $(3/4)^{(N-i)}$ is the probability of it not being a predicted reversal; the binomial coefficient, $\binom{(N-i)}{j}$, is the number of ways there could be j unpredicted reversals out of the N-i gamble pairs that remain after the predicted reversals have been accounted for; $(1/3)^j$ is the probability of an unpredicted reversal (from the universe of unpredicted reversals and consistent gamble pairs); and $(2/3)^{(N-i-j)}$ is the probability of the gamble pair being evaluated consistently.

Summing over the cells in Table C-1 that are 0, =, > or <, the a priori probabilities of these patterns occurring by chance are:

0	.0156
=	.2100
>	.3872
<	.3872

To calculate the number of subjects expected by chance to have a particular transition pattern (as represented by a cell in Table 5 or 6), one multiplies the number of subjects who had the starting pattern by the probability of having the second pattern by chance. For the transition < to 0, the expected number of subjects is n(<)p(0) or 30 * .0156 or .4680. The expected number of subjects for each of the transition patterns defined in the text is the sum of the expected number for each of the individual transitions. Thus, the e for the "To Predicted Reversals" transition pattern is

$n(<)p(=) + n(=)p(>) + n(<)p(>) + n(0)p(>)$

or

$30 * .2100 + 20 * .3872 + 30 * .3872 + 14 * .3872$

or

31.0808.

Table C-1

Possible Reversal Patterns for Six Gambles

The possible combinations of predicted and unpredicted reversals in six gamble pairs. Pairs that do not reverse are consistent. Cell labels: No reversals (0); more (>), equal (=), or fewer (<) predicted reversals than unpredicted reversals.

		Number of Gamble Pairs with Predicted Reversal Pattern, i						
		0	1	2	3	4	5	6
Number	0	0	>	>	>	>	>	>
of Gamble	1	<	=	>	>	>	>	
Pairs with	2	<	<	=	>	>		
Unpredicted	3	<	<	<	=			
Reversal	4	<	<	<				
Pattern, j	5	<	<					
	6	<						