The Economic Journal, **114** (*July*), 709–726. © Royal Economic Society 2004. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

TESTING EXPLANATIONS OF PREFERENCE REVERSAL*

Robin P. Cubitt, Alistair Munro and Chris Starmer

We present a new experimental investigation of preference reversal. Although economists and psychologists have suggested a variety of accounts for this phenomenon, the existing data do not adequately discriminate between them. Relative to previous studies, our design offers enhanced control for economic explanations and new tests of psychological hypotheses. We find a pattern of preference reversals that is inconsistent with all of the best-known explanations of the phenomenon proposed by economists, with the fundamental economic assumption of context-free preferences, and with several psychological theories of preference reversal. We explore the explanatory strategies that survive exposure to our data.

In conventional economic theory, preferences are taken as primitive. A fundamental assumption is that preferences are independent of the tasks that an agent faces. We will call this assumption the postulate of *context-free preferences.* It is crucial for many standard methods of economic analysis and implicit both in traditional models, such as expected utility theory, and in generalisations of them that have recently become prominent in the economics literature.[1] However, it is not indisputable. Empirically, one of the most troubling challenges to it is posed by observations of preference reversal.

A preference reversal (henceforth PR) occurs when the ranking of two (or more) items depends systematically on the method used to elicit it. Since its discovery by psychologists more than three decades ago (Lichtenstein and Slovic, 1971; Lindman, 1971), PR has been established as a replicable phenomenon.[2] This is worrying both for policy evaluation techniques that rely on preference elicitation, such as cost-benefit analysis, and for the foundations of economic theory. Grether and Plott (the first to report PR in the economics literature) famously commented: '*Taken at face value* (emphasis added) the data are simply inconsistent with preference theory and have broad implications about research priorities within economics. The inconsistency … suggests that no optimisation principles of any sort lie behind even the simplest of human choices…' (Grether and Plott, 1979, p. 623).

The subsequent debate has been polarised largely along disciplinary lines. Most psychologists have interpreted PR as evidence for the *context-sensitivity* of preferences[3] and so against conventional economic theory. A number of distinct hypotheses have been suggested to explain why such context-sensitivity might occur. In

* We are grateful for financial support from the Economic and Social Research Council (award no. L211252053) and the Leverhulme Trust (award no. F/00204/K). We thank David De Meza, Eric Johnson, Graham Loomes, Robert Sugden, two anonymous referees and participants at the seminars and conferences at which we have presented the paper for their comments on earlier versions.

[1] See Starmer (2000) for a survey.

[2] See Tversky and Thaler (1990), Hausman (1992, ch. 13) and Seidl (2002) for reviews of the evidence.

[3] The idea of context-sensitivity subsumes both preference construction (i.e. preferences formed only in the process of performing tasks) and the possibility that different tasks cue different, but pre-existing, preferences.

contrast, some economists have interpreted PR as evidence of intransitive preferences. Others have suggested that much of the existing evidence should not be taken at face value, by arguing that PR is an artefact of particular experimental designs. Although there are significant differences between them, these economic explanations have in common the feature that each maintains the postulate of context-free preferences against the alternative of context-sensitivity proposed by psychologists. As we explain below, the existing body of evidence does not discriminate between the economic and psychological interpretations.

In this paper, we report an experiment with two key objectives. One is to test for PR using a design with tight controls for economic explanations of the phenomenon. The other is to test three important psychological explanations of it, known as the prominence, task goal and scale compatibility hypotheses respectively. These objectives are achieved by a novel combination of experimental design features which we explain in Sections 2 and 3. The psychological theories that we consider predict different patterns of PR across our experimental treatments; but if preferences are context-free, no systematic pattern of PR should be observed in any of our treatments.

1. Economic Explanations of Preference Reversal

The classic cases of PR have related to decisions involving pairs of simple monetary gambles. In each of these pairs, one bet (the 'P-bet') offers a relatively large chance of a modest prize, while the other (the '$-bet') offers a smaller chance of a larger prize. In a typical PR experiment, a given subject makes a straight choice between the P-bet and the $-bet and also states a monetary valuation for each of them. The classic finding is a puzzling tendency for subjects to choose the P-bet over the $-bet in the choice task but to place a strictly higher monetary value on the $-bet. We will call such a tendency *standard preference reversal*. We will also call a tendency in the opposite direction (i.e. $-bet chosen in the choice task, but P-bet valued more highly) *counter preference reversal*. A crucial feature of the existing evidence of PR among monetary gambles is that it is *systematic*. Standard reversals have been overwhelmingly more commonly observed than counter reversals.

One theory which might explain this is regret theory (Loomes and Sugden, 1983; Loomes *et al.*, 1989, 1991). On this account, standard PR arises because of a systematic intransitivity in preferences.[4] Another important set of explanations maintains more of the structure of conventional preference theory, by retaining the assumption of transitivity. Instead, it implicates the incentive mechanisms used in most of the experiments which have recorded PR using real payoffs.

Such experiments face two central design issues: how to make the valuation tasks incentive-compatible and how to give incentives for earlier tasks without distorting those for later tasks. Previous experiments on PR with real payoffs have typically

[4] Regret theory, as usually formulated, is consistent with particular directions of preference cycle, including cycles of the form $V_m(P) \sim P \succ \$ \sim V_m(\$) \succ V_m(P)$, where P denotes the P-bet, $ the $-bet, $V_m(.)$ the agent's monetary valuation of a particular gamble and $V_m(\$) > V_m(P)$. This cycle would generate standard PR but regret theory would then not permit counter reversals (see Loomes *et al.*, 1991).

used one or both of the following procedures: the Becker-DeGroot-Marschak mechanism and the random lottery incentive system. The Becker-DeGroot-Marschak mechanism[5] is intended to provide incentives for experimental subjects to reveal true certainty equivalents in monetary valuation tasks. The random lottery incentive system is used in experiments, like those on PR, in which subjects must undertake multiple tasks.[6] It is intended to encourage subjects to think carefully about each task, while controlling for income effects which might otherwise arise if subjects were paid for all of them.

The Becker-DeGroot-Marschak mechanism and the random lottery incentive system are valid experimental procedures if subjects have preferences which satisfy expected utility theory. However, a series of papers published in the 1980s (Holt, 1986; Karni and Safra, 1987; Segal, 1988) demonstrated that, if agents' preferences are context-free but violate the independence and/or reduction axioms of expected utility theory, these are biased elicitation procedures which might generate the spurious appearance of standard PR. Since there is considerable evidence against the independence axiom of expected utility theory (see Starmer, 2000) and some evidence against the reduction axiom (Bernasconi and Loomes, 1992; Bernasconi, 1994) there is a coherent reason to discount the interpretation of many observations of PR as evidence either of context-sensitivity or of intransitivity.

However, this line of criticism can be avoided by utilising a device which Tversky *et al.* (1990) called the *ordinal payoff scheme*. This incentive mechanism is a crucial ingredient of our experimental design.

2. The Ordinal Payoff Scheme: Theory and Evidence

To understand the ordinal payoff scheme, consider an experiment in which there are *n parameter sets*, each consisting of a P-bet and a $-bet. The subject is told that she will face three tasks in relation to each parameter set: she will be asked to choose between the P-bet and the $-bet; and she will be asked to put a value on each of them separately. After she has performed these tasks for all *n* parameter sets, a random device will select one of the parameter sets, each with probability $1/n$. Then, another random device will select either *CHOICE* or *VALUATION*, each with probability 1/2. If it selects *CHOICE*, the subject will receive whichever gamble she chose in the choice task for the selected parameter set; if it selects *VALUATION*, she will receive whichever of the two gambles she put a higher value on (with a coin-toss as a tiebreaker, in the event of equal valuations). Whichever gamble she receives is then played out for real, using the subject's earlier response to it.

[5] In an experiment that uses this device to elicit a selling price for some gamble, the subject is first given an entitlement to the gamble and then asked to state the minimum price at which she would be prepared to sell it back to the experimenter. After this, a random device is used to generate a price. If, and only if, this price exceeds the price stated by the subject, the subject is required to sell the entitlement to the gamble back to the experimenter at the randomly generated price.

[6] In an experiment that uses the random lottery incentive system, each subject performs a number of tasks, knowing from the start that at the end of the experiment exactly one task will be selected at random to be for real, i.e. that it will be played out for real using the response that the subject has given. For discussion, see Cubitt *et al.* (1998, 2001).

The ordinal payoff scheme offers considerable improvement in experimental control relative to previous designs because, if any systematic pattern of PR is observed in an experiment which uses it, this cannot be explained by a theory which follows the standard choice theoretic approach of economics, even one that relaxes expected utility theory. To see why, note that any theory which assumes context-free preferences implies the existence of a *choice function* $C(\cdot)$ which governs behaviour in all tasks and which picks out, for any non-empty feasible set F of gambles, a non-empty preferred subset of F. We will refer to this approach to modelling as *generalised economic theory* (*GET*) – 'economic theory' because it retains the postulate of context-free preferences, 'generalised' because it allows these preferences to violate expected utility theory. Thus, our definition of *GET* subsumes all of the models which economists have put forward to explain PR.

Consider the implications of *GET* in an experiment that uses the ordinal payoff scheme. This is most straightforwardly illustrated for the special case[7] in which $n = 1$. Assume that the choice function is applied to the feasible set of compound gambles determined jointly by the possible responses to all tasks in the experiment.[8] Considering all possible responses to choice and valuation tasks, the subject faces the feasible set $F = \{P, \$, L_1, L_2, L_3\}$ where $L_1 = (P, 0.5; \$, 0.5)$, $L_2 = (P, 0.5; (P, 0.5; \$, 0.5), 0.5)$ and $L_3 = (\$, 0.5; (P, 0.5; \$, 0.5), 0.5)$. P (resp. \$) can be obtained by choosing P (resp. \$) in the choice task and giving it a strictly higher value; L_1 can be obtained by choosing one gamble in the choice task and giving the other gamble a strictly higher value; L_2 (resp. L_3) can be obtained by choosing P (resp. \$) in the choice task and giving equal valuations to the two gambles. Responses that display a PR between choice and valuation imply the selection of L_1 from this feasible set.

Within this design, the observation of an *instance* of PR is not inconsistent with *GET*. For example, an agent who is indifferent between P, \$ and lottery L_1 might select L_1; and a non-expected utility maximiser might have a strict preference for L_1 over both P and \$.[9] However, there are two distinct ways in which an agent may select lottery L_1: either by choosing P and valuing \$ more highly (the standard reversal); or by choosing \$ and valuing P more highly (the counter reversal). Since these two ways of responding select exactly the *same* element of F (i.e. L_1), *GET* can provide no reason to expect PR in one direction to be more frequent than PR in the other. So, observation of *systematic* PR in an experiment that utilises the ordinal payoff scheme would be a violation of *GET*.

Given that, unlike earlier designs, it controls for violations of expected utility theory, it is significant that the ordinal payoff scheme has, to our knowledge, never been used in the form set out above. In fact, we know of only two studies of PR which have used any close variant of it - Cox and Epstein (1989) and Tversky *et al.* (1990). Neither of them utilised the scheme in its pure form; and they produced

[7] Cubitt *et al.* (2000) shows that the argument generalises to any finite $n > 0$.

[8] We examine the implications of relaxing this assumption in Section 6.

[9] In order for L_1 to be strictly preferred to both P and \$, the agent would have to violate the property of *betweenness*. This is a weakening of the independence axiom, satisfied by expected utility theory and some, but not all, generalisations of it. In the present context it implies that L_1 can only be strictly preferred to P (resp. \$) if it is strictly dispreferred to \$ (resp. P).

strikingly different results. In Cox and Epstein's variant of the ordinal payoff scheme, payoffs depended only on relative (rather than absolute) valuations, but subjects were paid for each task and gamble payoffs were framed differently between choice and valuation tasks. Cox and Epstein reported an overall level of reversing behaviour in line with earlier studies but did not find the typical asymmetry between the two types of reversals. They concluded that the prior evidence of systematic reversals is best accounted for by *GET* (for example, by a non-expected utility model relaxing independence and reduction). However, their design created the potential for contaminating income effects and endowment effects.[10] Tversky *et al.* (1990) found the usual systematic pattern in reversals and concluded in favour of psychological models of PR. However, they used a design in which, for a subject in one group, all gambles – including the ordinal payoff scheme itself – were hypothetical and, for a subject in the other group, a real gamble would be faced only with a 15% probability.

Given the inconclusive evidence from existing studies that attempt to control for expected utility violations, it is an open question whether the phenomenon of PR undermines the postulate of context-free preferences. The most promising available means of resolving it is through a pure implementation of the ordinal payoff scheme.

3. A New Experiment

The central feature of the experiment reported here is the use of the ordinal payoff scheme, in its pure form, in conjunction with an experimental manipulation of the valuation tasks faced by subjects. The advantages of the ordinal payoff scheme do not depend on the valuation tasks being monetary. (All that is required is that valuations can be compared on some numerical scale.) We introduced a new type of valuation task, known as *probabilistic valuation*, using it in conjunction with more conventional *monetary valuation* tasks.

Subjects were randomly allocated between two experimental groups. The groups faced identical parameter sets and choice tasks but in one group (the MV-group) the corresponding valuations were monetary, whereas in the other (the PV-group) they were probabilistic. Each subject faced a total of eighteen tasks: six choices and twelve valuations. Tasks were presented one at a time on a computer screen, with the order randomised independently for each subject. In order to describe the tasks more precisely, it is first necessary to define the category of gambles used in the experiment.

[10] Cox and Epstein reported the results of tests for income effects, which did not indicate the presence of such effects. The additional possibility of an endowment effect in their design arises because, in valuation tasks, subjects were told that they would be paid a fixed sum, the prizes of the gambles were reduced by this sum and subjects were told that they would play out only the gamble to which they gave a higher value. In effect, they would sell the lower-valued gamble for the fixed sum. Although this preserves the final outcome implications of the gambles and controls for independence/reduction violation, it introduces a second framing difference between choice and valuation tasks, besides that inherent in the difference between choice and valuation, namely the separation of the fixed sum from residual payoffs in valuation tasks. Although Cox and Epstein argued that their findings did not suggest framing effects between choice and valuation, it is possible that these findings resulted from endowment effects offsetting the standard difference between choice and valuation.

A *basic gamble* g_i is a single-stage gamble in which the subject receives the *money-prize* $m_i > 0$ with *win-probability* $p_i \in [0, 1]$, and zero otherwise. Formally, such a gamble can be denoted by the ordered pair (m_i, p_i) of these *attributes*. In our experiment, gambles were described and subsequently played out using a bag containing a set of discs numbered from 1 to 100, one of which would be drawn if the gamble was played for real. In each choice task, the subject was presented with a pairwise choice between basic gambles. Figure 1 gives an example screen-display. To play option A from this example for real, the subject would draw a disc from the bag: a number from 1–31 would win £32, a number above 31 would win nothing.

In all valuation tasks, the subject was presented with a *yardstick*, which was a basic gamble with one of its attributes left unspecified. A different yardstick was used for each group. In the yardstick for the MV-group, the win-probability was fixed at unity and the missing attribute was the money-prize. To give a monetary valuation of some gamble g_i, the subject's task was to set the money-prize m_y of the yardstick such that $(m_y, 1) \sim g_i$. The subject's response was her *monetary valuation* of g_i. In the yardstick for the PV-group, the money-prize was fixed at £10 and the missing attribute was the win-probability. To give a probabilistic valuation of some gamble g_i, the subject's task was to set the win-probability p_y of the yardstick such that (£10, $p_y) \sim g_i$. The subject's response was her *probabilistic valuation* of g_i. Example screen displays for monetary and probabilistic valuation tasks are shown in Figures 2 and 3 respectively.

The six parameter sets used in the experiment are described in Table 1. Each row of the Table specifies a pair of basic gambles with win-probability expressed as a number of chances, out of 100, of winning a positive prize (in pounds sterling). Parameter sets 1 to 5 each contain a $-bet and a P-bet. These parameter sets were

Option A:	You get £32.00 for numbers 1 to 31
Option B:	You get £8.00 for numbers 1 to 97
Choose A or B:	

Fig. 1. *The Display for a Choice Task*

Set the missing amount so that the yardstick
and the option are equally attractive to you

yardstick:	You get £ ??? for numbers 1 to 100
option:	You get £18.00 for numbers 1 to 19
Pounds:	Pence:

Fig. 2. *The Display for a Money Valuation Task*

Set the missing number of discs so that the yardstick and the option are equally attractive to you

yardstick: You get £10.00 for numbers 1 to ???

option: You get £18.00 for numbers 1 to 19

Enter a number in the range 1 to 100:

Fig. 3. *The Display for a Probability Valuation Task*

Table 1

Parameters

	$ - bet		P – bet	
Set	Prize (£)	Winning numbers (out of 100)	Prize (£)	Winning numbers (out of 100)
1	32	1–31	8	1–97
2	18	1–19	4	1–81
3	13	1–50	6	1–94
4	17	1–39	5	1–94
5	10	1–50	4	1–92
		Dominated Option		Dominating Option
6	7	1–36	7	1–41

derived from parameters used by Tversky *et al.* (1990). We replaced US dollars with pounds sterling and doubled the nominal amounts. This preserves the ratios of prizes, whilst substantially increasing their real value and the expected value of the gambles.[11] Parameter set 6 is different: its two gambles offer the same prize but one of them (the dominating option) gives more chances of winning it than the other. The extent to which subjects choose (or value) a dominated option over a dominating option provides an indication of the propensity for 'error'. We will comment on the significance of this parameter set in Section 6.

We utilised the ordinal payoff scheme in both groups: every subject played for real some gamble, selected by that scheme in the light of their task-responses. As there were six parameter sets, we implemented the scheme as follows. Before completing the tasks, subjects were told that, at the end of the experiment, they would throw a die to select one parameter set. They would then throw the die a second time. If it came up 1–3, they would play the gamble which they had chosen in the choice task for the selected parameter set; if it came up 4–6, they would play

[11] Win-probabilities are unchanged from those used by Tversky *et al.* (1990), except for rounding. The expected value of gambles in our experiment ranged from approximately £2.50 to £10 (with £1 ≅ $1.5 at the time of the experiment). The average expected value was £5.20.

the gamble that they had valued most highly in the valuation tasks for the selected parameter set.

Our MV-group provides a check on the robustness of the findings of Tversky *et al.* (1990). Since that paper provides the only previous evidence of systematic PR using a design which attempts to control for violations of expected utility theory, this is a significant feature of our experiment. The experimental manipulation of the valuation task provides a further test of *GET* and tests of three important psychological explanations of PR. We explain this in the next Section.

4. Theoretical Predictions

Our design allows well-controlled tests among various groups of hypotheses. Since we used the ordinal payoff scheme in both groups, observation of systematic PR in *either* group in any parameter set would be incompatible with *GET*. Further, from the perspective of *GET*, each group in our design faces an *identical* overall decision problem. It is to respond to six pairwise choice questions and to assign a number to each of twelve gambles. The absolute magnitudes of the numbers are irrelevant to the determination of which gamble is received; all that matters for this is their ordering within each parameter set. The scale in which valuations are given should make no difference. Thus, *GET* would predict no systematic difference between the behaviour of our two groups, in terms either of the responses to the choice tasks or of the pairwise preferences implied by the valuation tasks.

Our design was also intended to discriminate between particular psychological hypotheses about the source of PR. To see this, note that each of our valuation tasks is an example of what is called in the psychology literature a *matching task*, i.e. a task in which the subject has to set the value of some missing attribute so that two multi-attribute items are equally attractive.[12]

The *prominence hypothesis* (Tversky *et al.*, 1988)[13] is concerned with the relationship between choice and matching. It presupposes the existence of an attribute which, in some sense, is the most important to the agent. The hypothesis is that the prominent attribute weighs more heavily in choice tasks than in matching tasks.[14]

The *task goal hypothesis* (Fischer *et al.*, 1999) also postulates the existence of a prominent attribute. Its novelty is to suggest that the weight accorded to this attribute in a task depends upon the goal of the task, as perceived by the subject.

[12] Matching tasks have been extensively studied in psychology, usually in relation to items other than gambles. A precursor of our use of matching on a probabilistic dimension is MacCrimmon and Smith (1986).

[13] The formulation of the hypothesis by Tversky *et al.* (1988) built upon earlier work by Slovic (1975). For more recent discussions, see Fischer and Hawkins (1993), Slovic (1995) and Fischer *et al.* (1999).

[14] One explanation for this hypothesis is that a choice task is qualitative, whereas a matching task is quantitative. Choices may cue qualitative reasoning strategies designed to find a decisive reason to choose one way or another, so bringing the prominent attribute to the fore. Matching tasks, on the other hand, may cue quantitative reasoning strategies involving tradeoffs between attributes. This explanation of the prominence hypothesis was suggested by Tversky *et al.* (1988); Fischer and Hawkins (1993) articulate it in more detail and call it the strategy compatibility hypothesis. This is distinct from the scale compatibility hypothesis.

The hypothesis is that the prominent attribute is weighted more heavily in tasks whose perceived goal is to differentiate between items than in tasks whose perceived goal is to equate them.

The *scale compatibility hypothesis* (Tversky *et al.*, 1988)[15] focuses on the fact that any task must have a response mode. A particular response mode may be compatible with certain types of attribute; for example, the attribute and the response mode might share the same scale. In general, response modes in different tasks may differ in their compatibility with any given attribute. The hypothesis is that the weight of a given attribute in any task is greater the more compatible that attribute is with the response mode of the task.

Each of these hypotheses can explain the standard pattern of PR between P- and $-bets in previous experiments which used monetary valuations. For example, the evidence can be accounted for by the prominence hypothesis if 'win-probability' is the prominent attribute. According to this view, choices are more affected by differences in win-probability than are monetary valuations. Thus, there is a tendency for the P-bet to be chosen even when a higher monetary valuation is put on the $-bet. However, the same evidence can be explained by the task goal hypothesis, provided that subjects see the goal of choice tasks but not of valuation tasks, as being to differentiate between the P-bet and the $-bet. Finally, standard PR in previous experiments can also be explained by the scale compatibility hypothesis, given the claim that the attribute 'money-prize' is more compatible with a monetary valuation scale than it is with the response mode of a choice task. On this theory, subjects pay more attention to the size of the monetary prizes when putting monetary values on gambles than they do when choosing between them. This induces a tendency to value the $-bet more highly even when the P-bet would be chosen in a straight choice.

However, in our design, these three psychological hypotheses have distinct implications because our MV- and PV- groups faced different types of valuation task and because of our use of the ordinal payoff scheme.

Since both monetary and probabilistic valuation are matching tasks, the prominence hypothesis implies that the prominent attribute should weigh more heavily in choice tasks than in *either* type of valuation task. If prominence is the source of PR, any systematic tendency towards PR in a given parameter set should be in the *same* direction in both groups. To be consistent with the explanation of previous evidence by the prominence hypothesis, this must be the direction of standard reversal. Moreover, in our design, there should be no systematic difference between the two groups, in terms of the binary preferences implied by their valuations.[16]

The ordinal payoff scheme drives a wedge between the form of the valuation tasks and their goal. Each valuation task has the form of a matching task, but its goal – just like that of a choice task – is to determine which of two gambles is received, i.e. to differentiate. Provided subjects perceive this, the task goal

[15] See also Tversky *et al.* (1990), Slovic *et al.* (1990) and Slovic (1995).

[16] The prominence hypothesis gives no reason to suppose that the weight on the prominent attribute in a valuation task will depend on the response scale.

hypothesis would not predict PR in either group. On the milder assumption that the perceived goal of the valuation tasks is independent of the valuation response scale, the task goal hypothesis implies that there should be no systematic difference between groups.

Alone among the hypotheses which we have discussed, the scale compatibility hypothesis asserts that the valuation response scale *itself* will affect behaviour. On the assumption that money-prize (resp. win-probability) is more compatible with monetary (resp. probabilistic) valuation than with the choice response mode, the scale compatibility hypothesis predicts a systematic tendency for PRs to be in the *standard* direction (*P* chosen, $ valued more highly) in the MV-group and to be in the *counter* direction ($ chosen, *P* valued more highly) in the PV-group.

We summarise the implications of the hypotheses we have discussed in Table 2. The Table identifies, separately for the groups undertaking monetary and probabilistic valuations, whether a *systematic* tendency towards PR in a particular direction is allowed by each of the hypotheses.

5. Results and Interpretation

A total of 230 subjects – predominantly undergraduate students from a wide range of disciplines at the University of East Anglia – took part across 22 sessions. Adjacent sessions were paired and a coin was tossed, in advance, to decide which of the pair would be in the MV-group and which in the PV-group. As a result of this randomisation, the MV-group finally contained 114 subjects and the PV-group 116 subjects.

The main results are summarised in Table 3 which presents the data for each parameter set, separately for the MV- and PV-groups. The first row reports the sample size. The next two rows show responses to choice tasks (as percentages of subjects choosing *P* and $). For any given parameter set, an individual subject's behaviour in the three tasks relating to it can be classified into one of four categories: (I) *consistent preference for P* (*P* chosen and valued at least as highly as $);

Table 2

Summary of Predictions of Economic and Psychological Hypotheses

	STANDARD Pref. Rev. Allowed?		*COUNTER* Pref. Rev. Allowed?	
Maintained Hypothesis	Monetary Valuation (MV)	Probabilistic Valuation (PV)	Monetary Valuation (MV)	Probabilistic Valuation (PV)
Generalised Economic Theory	No	No	No	No
Prominence*	Yes	Yes	No	No
Task Goal†	No	No	No	No
Scale compatibility‡	Yes	No	No	Yes

Notes:
*Conditional on the assumption that win-probability is the prominent attribute.
†Conditional on the assumption that the perceived goal of valuation tasks in an ordinal payoff scheme is differentiation.
‡Conditional on the assumptions that money-prize is more compatible with monetary valuation than with choice and that win-probability is more compatible with probabilistic valuation than with choice.

Table 3

Summary of Results

	Set 1		Set 2		Set 3		Set 4		Set 5		Set 6*	
	MV	PV	MV	PV	MV	PV	MV	PV	MV	PV	MV	PV
Sample Size	114	116	114	116	114	116	114	116	114	116	114	116
Choice												
% choosing *P*	51.75	51.72	57.89	54.31	40.35	39.66	31.58	42.24	24.56	25.00	95.61	98.28
% choosing $	48.25	48.28	42.11	45.69	59.65	60.34	68.42	57.76	75.44	75.00	4.39	1.72
Choice & Valuation												
I: Consistent, prefer *P*	9	29	16	44	11	27	4	18	3	16	92	100
II: Consistent, prefer $	53	34	43	29	62	37	74	47	84	60	4	1
III: Standard reversal	50	31	50	19	35	19	32	31	25	13	17	14
IV: Counter reversal	2	22	5	24	6	33	4	20	2	27	1	1

*In Parameter set 6: P6 dominates $6.

(II) *consistent preference for* $ ($ chosen and valued at least as highly as *P*); (III) *standard reversal* (*P* chosen, $ valued more highly); and (IV) *counter reversal* ($ chosen, *P* valued more highly). The last four rows of Table 3 report the frequencies of each category of behaviour, for each parameter set and each group. Note that the *P*/$ terminology does not apply to parameter set 6 and, to begin with, we focus on the data for parameter sets 1–5 only. Figure 4 provides a compact visual summary of the distributions across behaviour categories I to IV, separately for the MV and PV-groups, aggregating across parameters sets 1–5.

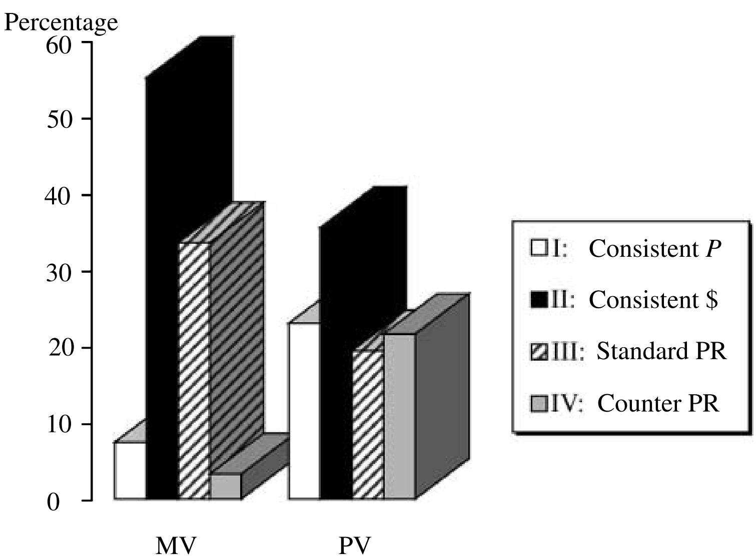

Fig. 4. *Percentages of Different Categories of Behaviour* (aggregate for parameter sets 1–5)

The MV-group corresponds most closely to traditional PR experiments with monetary valuations. The data from this group provide the basis for a simple test of *GET* because *GET* rules out the possibility of any asymmetric pattern of reversals. Figure 4 reveals clear evidence of systematic PR: in the MV-group, standard reversals account for around 35% of all behaviour and, crucially, they occur much more frequently than counter reversals. Table 3 shows that this asymmetric tendency towards standard PR occurs for each of parameter sets 1–5 (as a proportion of total reversals, the standard reversal rate varies between 85% and 96%). We tested the null hypothesis, separately for each parameter set, that the proportion of MV-group subjects who rank $ over *P* is the same under choice and valuation. In each case, using a one tailed t-test and a 1% significance level, we can confidently reject the null in favour of the alternative that the proportion who rank $ over *P* is greater in valuation compared with choice. These results are consistent with the classic findings of the PR literature; they conform closely with the findings of Tversky *et al.* (1990); and they are inconsistent with *GET*.

Comparison of behaviour across the groups provides tests of both economic and psychological hypotheses. The upper part of Table 3 shows that, when faced with identical tasks, the behaviour of the two groups was almost identical, since the only tasks that were common across groups were the choices. However, the behaviour of the two groups in relation to which gamble was valued more highly, within each parameter set, was different. Table 4 summarises tests of the hypothesis that subjects' relative valuations in the PV-group are drawn from the same population as those for the MV-group. Using a Pearson chi-square test, for each of parameter sets 1 to 5, the null hypothesis of no significant difference is rejected at the 1% level.

Figure 4 shows that, although the aggregate level of reversing behaviour (combining standard and counter reversals) in the PV-group is similar to that of the MV-group, the distribution of reversals is markedly different. For the PV-group, at the aggregate level, counter reversals occur more frequently than standard reversals (21.7% of behaviour, as against 19.5%). Inspection of the lower part of Table 3 reveals that, for some parameter sets, counter reversals are more common than standard reversals, but for other parameter sets the pattern is reversed. Unlike the MV-group, there appears to be no consistent pattern to the distribution of reversals in the PV-group that is general across parameter sets. Since the tasks faced by the two groups are formally equivalent, this systematic difference between groups cannot be explained by *GET*.

A simple between-group comparison also eliminates the task goal and prominence hypotheses as contenders for organising the data. Neither theory predicts any

Table 4

Do the PV-group Relative Valuations Come from the Same Population as the MV-group Relative Valuations?

Parameter set	1	2	3	4	5	6
χ^2_3	258.5**	145**	162.2**	122.9**	381.5**	3.48

** = Significant at 1% level.

significant difference between groups, with respect to the pairwise preferences implied by valuations.[17] Among the hypotheses that we have considered, the only theory which predicts a systematic difference between our two experimental groups is the scale compatibility hypothesis. So are these observed differences consistent with scale compatibility? The natural interpretation of the scale compatibility hypothesis is that money-prize is more compatible with monetary valuation than with choice and that win-probability is more compatible with probabilistic valuation than with choice. The fact that standard reversals are a lower proportion of total reversals in the PV-group than in the MV-group is consistent with this. However, the scale compatibility hypothesis would not lead us to expect counter reversals in the MV-group or standard reversals in the PV-group. The high frequency of standard reversals in the PV-group, therefore, cannot be explained by this version of scale compatibility acting alone.

Interestingly, although the aggregate data for the PV-group does not show any preponderance of one type of reversal over the other, the pattern is more systematic at the level of the individual. For example, for each of parameter sets 2–5, we tested the hypothesis that the proportion of individuals in the PV-group who exhibit a counter reversal in that parameter set is independent of whether the individuals concerned showed a counter reversal in parameter set 1. In each case the null-hypothesis is rejected (with p values varying from 0.05 to 0.004) in favour of the alternative that individuals who have a counter reversal in parameter set 1 are more likely to have a counter reversal in the other parameter set. (For instance, out of the 22 individuals who exhibit a counter reversal for parameter set 1, 9 also have a counter reversal for set 2 (i.e. 41%), compared to 15 out of 94 (16%) amongst the subjects who do not have a counter reversal in parameter set 1.) We used a chi-squared test to examine the general hypothesis of independence of the probability of counter reversals across all parameter sets. This is decisively rejected at the 1% level - relative to the pattern of reversals that would be expected given independence, there are too many individuals in the tails of the distribution, either with no counter reversals or with 4 or 5 counter reversals. The corresponding test for standard reversals also leads to rejection of the null hypothesis and reveals a similar pattern of extremes. Finally, to investigate association between counter reversal and standard reversal, we test the null hypothesis that the proportion of individuals who exhibit a standard reversal in each of parameter sets 1,2, 4 and 5 is independent of whether the individuals concerned showed a counter reversal in parameter set 3. Out of the 33 individuals who exhibit a counter reversal for parameter set 3, very few have standard reversals in the other tasks. For parameter set 2, four out of the 33 have a standard reversal (compared to 15 out of 83 in the control group) and this is not significantly different, but for parameter sets 1, 4 and 5, the null hypothesis is rejected at probabilities between 0.01 and 0.04. Here and earlier in the paragraph, the tests reported are simply illustrative.

[17] The comparison between groups is inconsistent with the task goal hypothesis, conditional only on the assumption that the perceived goals of tasks are independent of the valuation response scale. Conditional on the stronger assumption that the perceived goal of all tasks in the ordinal payoff scheme is differentiation, the task goal hypothesis would also be unable to explain the systematic pattern of reversals in the MV-group.

In each case, it makes no qualitative difference to the results which parameter set is selected as the baseline.

To sum up, preference reversals are common in both of our experimental groups. In the MV-group there is a pronounced tendency, in each parameter set, for those subjects who are not consistent to commit reversals in the standard direction only. In the PV-group, both types of reversal are common in the aggregate but, at the individual level, there is a tendency for those subjects who are not consistent to commit reversals in the same direction across parameter sets.

6. Possible Explanations for our Findings

Given that neither *GET* nor any of the psychological hypotheses that we have considered, acting in isolation, can explain our findings, it seems natural to ask how our data can be organised? The most striking findings in need of explanation are the highly systematic pattern of reversals in the MV-group and the marked difference between groups. We can think of three possible explanatory strategies: one assumes decision errors; another is a 'myopic' specification of regret theory; while a third invokes multiple decision-heuristics. They have in common the feature that each invokes an element of bounded rationality. In this Section, we briefly explain them and offer a speculative evaluation of their respective merits, keeping in mind that our experiment was not designed to discriminate between them.

6.1. *Stochastic Preference Models*

There has been considerable interest among economists in stochastic theories of decision making (Loomes and Sugden, 1995). One surprising product of this research has been to demonstrate that patterns of behaviour which violate standard theories of preference, like expected utility theory, may nevertheless be consistent with stochastic versions of those preference models. Might some stochastic variant of *GET* explain our data?

Any explanation of our findings must account for the systematic difference between groups. In order to do this, a stochastic version of *GET* might suppose that subjects have true preferences which can be represented by some non-stochastic *GET* but that they may make errors in tasks, leading them to mis-record their preferences. This could generate systematic differences between groups if subjects in one group are more prone to mistakes than the other. For example, this could happen if, notwithstanding the formal equivalence of the tasks faced by the two groups, their perceived difficulty differed between groups. Without further refinement, such a model could not explain the patterns in the individual level data, but perhaps it could provide an account of the difference between groups at the aggregate level.

To investigate this, consider responses to parameter set 6 tasks recorded in the final columns of Table 3. Notice that very few subjects select the dominated option in parameter set 6 choice tasks (4.4% in the MV-group, 1.7% in the PV-group), while dominance is violated more frequently in valuation tasks (18% in the

MV-group; 13% in the PV-group). To the extent that violations of dominance reveal 'error', we conclude that errors are more likely in valuation, relative to choice, but there is very little difference between the rates of error in the MV- and PV-groups. The difference between groups is not significant at the 5% level in either choice or valuation. Moreover, the differences that are observed look to be in the 'wrong' direction: in the aggregate, most reversals (standard + counter) are observed in the PV-group. It seems plausible to think that if the between group difference in reversing is to be explained by errors relative to *GET*, error rates will need to be higher in the PV-group. That is the opposite of what we observe in parameter set 6. So, we conjecture that a stochastic variant of *GET* is not an obviously promising line of attack for explaining differences between our two experimental groups.

6.2. *Regret Theory and Bounded Rationality*

We have argued that, because our design uses the ordinal payoff scheme, no *GET* can explain systematic PR in our experiment. In line with the tradition of unbounded rationality in economic theory, we assumed that agents behave as if they consider the set of choice and valuation tasks in the experiment as a single choice among compound gambles. However, this may be too taxing for a subject, in terms of the computation required, its demands on memory, and so on. An alternative would be to assume bounded rationality. While we offer no formal model of bounded rationality here, one account of how a boundedly rational agent might approach the tasks would be to suppose that subjects consider each in isolation and do their best to give honest and considered responses. We will call agents who behave in this way *myopic agents*. Can *GET* explain systematic PR in our design if we assume myopic agents in place of unboundedly rational ones? In Cubitt *et al.* (2000), we show that for most *GET* models the answer is 'No'. However, we also demonstrate that it is possible to construct a *GET*-style model which might explain PR in our design, given myopic agents who violate transitivity. Specifically, we show that a variant of regret theory with myopic agents is consistent with the asymmetry between standard and counter PR observed among MV-group subjects; it also provides a reason to expect *some* difference between the behaviour of the two groups. However, it cannot explain why standard reversals are observed in the PV-group or why certain individuals are systematically prone to them. So, the model provides only a partial fit to our data.

6.3. *Conjunctions of Psychological Processes*

Perhaps the simplest approach to explaining the data would be a hybrid theory based on some conjunction of psychological processes. The tendency for the proportion of all reversals which are in the standard direction to fall as one moves from the MV-group to the PV-group is in the direction predicted by scale compatibility. So, a natural question to ask is whether there is some additional hypothesis which, when combined with scale compatibility, would explain those features of the data otherwise inconsistent with it? The feature of the data most in need of additional explanation is the occurrence of standard reversal in the PV-group. An obvious candidate explanation for this is the prominence hypothesis, which predicts standard reversals in both groups. A theory which proposed the simultaneous operation of scale compatibility and prominence effects would predict a relatively high frequency of standard reversal in the MV-group (where both effects operate in the same direction); and, while standard reversals should be less frequent in the PV-group (assuming a non-negligible compatibility effect), they may still occur as a result of the operation of prominence. If the relative intensity of the two effects varies across individuals, such an account could also be compatible with our individual-level data. While there is ongoing dispute concerning the relative importance of scale compatibility effects and prominence effects, some researchers – we believe Slovic *et al.* (1990) were the first – have explicitly suggested the combination of scale compatibility and prominence effects as explanations of PR. Our data are broadly consistent with that conclusion.

7. Conclusion

Previous studies of preference reversal show an apparent inconsistency between agents' behaviour in situations in which they choose between gambles and in which they put values on them. *Prima facie*, this phenomenon is incompatible with a fundamental postulate of economic theory, namely that agents have context-free preferences, i.e. preferences which are independent of the tasks which they face. However, economists have previously provided good theoretical reasons for scepticism about this prima facie reading of the data. For example, many experiments which report preference reversal do not embody satisfactory controls, if viewed as tests of the postulate of context-free preferences. In the designs which economists criticised, the appearance of preference reversal could arise if agents have context-free, but non-expected utility, preferences. In contrast, psychologists have tended to be more sympathetic to the prima facie reading.

One motivation for the study reported here was therefore to look for preference reversal in a design with tight controls for economists' explanations of the phenomenon. The other was to test some of the leading explanations proposed by psychologists. To these ends, we employed a distinctive incentive mechanism, using it in a purer form and with stronger incentives than has (to our knowledge) been done previously. Additionally, our experiment introduced a new type of task for valuing gambles, called probabilistic valuation, which we used in conjunction with conventional monetary valuation tasks. There were two experimental groups. One undertook choice and monetary valuation tasks, the other choices and probabilistic valuations. Our two main findings were (i) strong evidence of systematic preference reversals in one direction in the group whose valuations were monetary; (ii) a marked difference in the pattern of preference reversals between the two groups. Explanations of preference reversal which retain the postulate of context-free preferences can explain neither of these findings; however, nor can any of the leading psychological hypotheses which we have considered, when acting in isolation. Our experiment therefore renews the puzzle of preference reversal.

Given this, it is natural to ask what explanatory strategies are suggested by our data. We have briefly considered three possibilities, each of which involves an important element of bounded rationality. Two of them – one based on stochastic choice, another on a combination of regret and myopia – retain some of the spirit of standard economic theory, whilst departing from it in important respects, especially in the latter case. The third invokes a particular combination of psychological mechanisms acting in conjunction. These accounts were constructed in the light of our findings, so our experiment was not designed to discriminate between them. However, we have shown that our data provide grounds for reservations about the first two of these accounts. Although postulating a combination of psychological mechanisms is an unfamiliar form of explanation in economics, it may offer the most straightforward explanation of our findings.

University of East Anglia
University of Nottingham

Date of receipt of first submission: September 2000
Date of receipt of final typescript: July 2003

Technical Appendix is available for this paper: www.res.org.uk/economic/ta/tahome.asp

References

Bernasconi, M. (1994). 'Nonlinear preferences and 2-stage lotteries: theories and evidence', ECONOMIC JOURNAL, vol. 104, pp. 54–70.

Bernasconi, M. and Loomes, G. (1992). 'Failures of the reduction principle in an Ellsberg-type problem', *Theory and Decision*, vol. 32, pp. 77–100.

Cox, J.C. and Epstein, S. (1989). 'Preference reversals without the independence axiom', *American Economic Review*, vol. 79, pp. 408–26.

Cubitt, R.P., Munro, A. and Starmer, C. (2000). 'Preference reversal: an experimental investigation of economic and psychological hypotheses', CeDEx discussion paper no. 2000-1, University of Nottingham.

Cubitt, R.P., Starmer, C. and Sugden, R. (1998). 'On the validity of the random lottery incentive system', *Experimental Economics*, vol. 1, pp. 115–31.

Cubitt, R.P., Starmer, C. and Sugden, R. (2001). 'Discovered preferences and the independence axiom: testing for common consequence and common ratio effects using single task designs', *Journal of Economic Methodology*, vol. 8, pp. 385–414.

Fischer, G.W. and Hawkins, S.A. (1993). 'Strategy compatibility, scale compatibility, and the prominence effect', *Journal of Experimental Psychology: Human Perception and Performance*, vol. 19, pp. 580–97.

Fischer, G.W., Carmon, Z., Ariely, D. and Zauberman, G. (1999). 'Goal-based construction of preferences: task goals and the prominence effect', *Management Science*, vol. 45, pp. 1057–75.

Grether, D. and Plott, C. (1979). 'Economic theory of choice and the preference reversal phenomenon', *American Economic Review*, vol. 69, pp. 623–38.

Hausman, D. (1992). *The Inexact and Separate Science of Economics*, Cambridge: Cambridge University Press.

Holt, C.A. (1986). 'Preference reversals and the independence axiom', *American Economic Review*, vol. 76, pp. 508–15.

Karni, E. and Safra, Z. (1987). '''Preference reversal'' and the observability of preferences by experimental methods', *Econometrica*, vol. 55, pp. 675–85.

Lichtenstein, S. and Slovic, P. (1971). 'Reversals of preference between bids and choices in gambling decisions', *Journal of Experimental Psychology*, vol. 89, pp. 46–55.

Lindman, H. (1971). 'Inconsistent preferences among gambles', *Journal of Experimental Psychology*, vol. 80, pp. 590–7.

Loomes, G., Starmer, C. and Sugden, R. (1989). 'Preference reversal: information processing effect or rational non-transitive choice?', ECONOMIC JOURNAL, vol. 99, pp. 140–51.

Loomes, G., Starmer, C. and Sugden, R. (1991). 'Observing violations of transitivity by experimental methods', *Econometrica*, vol. 59, pp. 425–40.

Loomes, G. and Sugden, R. (1983). 'A rationale for preference reversal', *American Economic Review*, vol. 73, pp. 428–32.

Loomes, G. and Sugden, R. (1995). 'Incorporating a stochastic element into decision theories', *European Economic Review*, vol. 39, pp. 641–8.

MacCrimmon, K. R. and Smith, M.R. (1986), 'Imprecise equivalences: preference reversals in money and probability' University of Columbia working paper no.1211, November.

Segal, U. (1988). 'Does the preference reversal phenomenon necessarily contradict the independence axiom?', *American Economic Review*, vol. 78, pp. 233–5.

Seidl, C. (2002). 'Preference reversal', *Journal of Economic Surveys*, vol. 6, pp. 621–55.

Slovic, P. (1975). 'Choice between equally valued alternatives', *Journal of Experimental Psychology: Human Perception and Performance*, vol. 1, pp. 280–7.

Slovic, P. (1995). 'The construction of preferences', *American Psychologist*, vol. 50, pp. 364–71.

Slovic, P., Griffin, D. and Tversky, A. (1990). 'Compatibility effects in judgment and choice'. In (R. Hogarth, ed.), *Insights in Decision-Making*, pp. 5–27. Chicago: University of Chicago Press.

Starmer, C. (2000). 'Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk', *Journal of Economic Literature*, vol. 38, pp. 332–82.

Tversky, A., Sattath, S. and Slovic, P. (1988). 'Contingent weighting in judgement and choice', *Psychological Review*, vol. 95, pp. 371–84.

Tversky, A., Slovic, P and Kahneman, D. (1990). 'The causes of preference reversal', *American Economic Review*, vol. 80, pp. 204–17.

Tversky, A. and Thaler, R.H. (1990). 'Preference reversals', *Journal of Economic Perspectives*, vol. 4, pp. 201–11.