SCHOOL OF ECONOMICS
DISCUSSION PAPER 98/17

Non-transitive Choice: Event-splitting effects or framing effects?

Steven J. Humphrey*
University of Nottingham, UK

A significant body of experimental evidence documents systematic violations of the transitivity axiom of expected utility theory in a manner consistent with regret theory. In a recent paper Starmer and Sugden (in press) test a variety of alternative explanations of this evidence, including event-splitting effects, and conclude that anticipated regret provides the best explanation, but is sensitive to minor changes in problem representation. Humphrey (1998) shows that the apparent sensitivity of regret-cycling to problem representation in Starmer and Sugden's data is best explained by the event-splitting argument they rule-out. This paper reports an experiment which discriminates between regret-driven and event-splitting-driven cycles in a manner which controls for the sources of ambiguity in Starmer and Sugden's tests. There is some evidence that event-splitting effects are important, but significant non-transitive choices are observed which cannot be explained by event-splitting effects, and the data are best organised by regret theory plus framing effects.

Keywords: Regret theory, event-splitting effects, non-transitive choice, experimental economics

JEL classification: D81, C91

---

* I am grateful for financial assistance from the University of Nottingham, award no. 97NLRG865. I would also like to acknowledge the programming assistance of Matt Ayres and the assistance of Chris Jackson in running the experiment. Correspondence to: Steven J. Humphrey, School of Economics, University of Nottingham, Nottingham NG7 2RD, UK.

1. Introduction

The transitivity axiom of expected utility theory [EUT, see e.g. von Neumann and Morgenstern (1944)] is a fundamental property of rational choice in economic theory. It is often justified as statement of logic or, perhaps more controversially, by arguing that an agent who does not adhere to transitivity can, in principle, be 'money pumped' into bankruptcy[1]. In a series of experiments Loomes, Starmer and Sugden (1989, 1991) observe significant violations of transitivity in a manner consistent with regret theory (Loomes and Sugden, 1987). More recently, Starmer and Sugden (in press) draw attention to two possible alternative explanations of this evidence – event-splitting effects (ESEs, Starmer and Sugden, 1993; Humphrey, 1995) and framing effects (Harless, 1992) – and re-run the original tests to discriminate between these explanations and, in the case of ESEs, the original hypothesis of regret-aversion. A similar incidence of non-transitive choice is observed between the original problem representation and one which controls for ESEs, but a lower incidence of regret-cycling is observed in slightly different (but also ESE-controlled) displays which Harless (1992) claims to reduce the impact of regret-aversion. Consequently, Starmer and Sugden (in press) rule-out ESEs as driving observed non-transitivities and conclude, in a similar manner to Harless, that regret-aversion is more likely to motivate decision-making under some problem representations than it is under others.

Humphrey (1998) shows that Starmer and Sugden's conclusions may stem, in part, from peculiarities in observed behaviour not explained by any theory, and that their data do not rule-out ESEs being responsible for observed differences in the incidences of non-transitive choices between subjects. This paper presents new experimental tests of whether individuals systematically violate transitivity in

[1] The controversy here stems from the argument that violations of transitivity are necessary but not sufficient for an agent to be money-pumped. Although an agent might violate transitivity (and according to regret theory, do so rationally) in choice they would recognise attempts to be money-pumped and avoid them.

decision-making under risk, and reports a design which controls for the sources of ambiguity in Starmer and Sugden's data.

2. Regret theory and non-transitive choice

The form of regret theory employed in the prediction of systematic non-transitive choices is that presented in Loomes and Sugden (1987) and axiomatized by Sugden (1993). Assume that an individual is faced with a choice between two acts A_i and A_j which are defined over a finite number of states of the world. Let x_{ik} and x_{jk} be the respective consequences of A_i and A_j in state of the world S_k, and p_k be the probability of state of the world S_k occurring. Regret theory assumes the existence of a real-valued function $F(x_{ik}, x_{jk})$ which is skew-symmetric ($F(x, y)$=-$F(y, x)$ for all x, y), unique up to multiplication by a positive constant, and increasing in its first argument. The function can be interpreted as the anticipated net advantage of choosing A_i rather than A_j if state of the world S_k occurs such that[2]:

$$A_i \;\begin{matrix}\succ\\ \sim\\ \prec\end{matrix}\; A_j \;\Leftrightarrow\; \textstyle\sum_k p_k F(x_{ik}, x_{jk}) \begin{matrix}>\\ =\\ <\end{matrix} 0 \qquad (1)$$

Note that expression (1) shows regret theory to embody EUT as a special case where $F(x_{ik}, x_{jk})=u(x_{ik})-u(x_{jk})$ and $u(.)$ is a von Neumann-Morgenstern utility function. The novel predictions of regret theory arise from the assumption of *regret-aversion*, which states that for any three consequences x, y, z where $x>y>z$, $F(x, z) > F(x, y) + F(y, z)$. It is this assumption which engenders the prediction of systematic non-transitive choices over a set of pairwise decision problems (problem 1 {A or B}, problem 2 {B or C}, problem 3 {A or C}) between each of the three acts (A, B and C) contained in a standard triple. Table 1 shows an example of a standard triple.

---

[2] Where '$\sim$' and '$\succ$' represent relations of indifference and strict preference, respectively.

TABLE 1: Standard Triple

Prospect/Act	High Consequence	(Prob)	Low Consequence	(Prob)
A	a	p_a	d	$1\text{-}p_a$
B	b	p_b	e	$1\text{-}p_b$
C	c	p_c	f	$1\text{-}p_c$

In the standard triples with which this paper is concerned, $a > b > c > d = e = f = 0$ and $1 \geq p_c > p_b > p_a > 0$. The only way in which transitivity can be violated over a standard triple is if either of the two patterns of cyclical choice, ABC or BCA, are observed[3]. Assuming, for simplicity, that p_c=1 an application of the decision rule in expression (1) to the standard triple yields expression (2). Expression (2) is required to hold if one is to observe a BCA violation of transitivity:

$$p_a[\Phi(a, c)\text{-}\Phi(a, b)\text{-}\Phi(b, c)]+(p_b\text{-}p_a)[\Phi(b, 0)\text{-}\Phi(b, c)\text{-}\Phi(c, 0)] > 0 \qquad (2)$$

Since, by the assumption of regret-aversion, both terms in square brackets are positive, regret theory allows BCA violations of transitivity. Starmer and Sugden (in press) term this violation the *clockwise* cycle. The *anticlockwise* cycle, ABC, would imply the reverse inequality in expression (2), and is thereby inconsistent with regret-aversion[4].

---

[3] Where 'ABC', for example, denotes A being chosen in the choice between A and B, B in the choice between B and C, and C in the choice between A and C.

[4] All six other possible patterns of choice over the standard triple do not violate transitivity and are, therefore, consistent with EUT. Since regret theory accommodates EUT as a special case they are also consistent with regret theory.

FIGURE 1: Problem representations

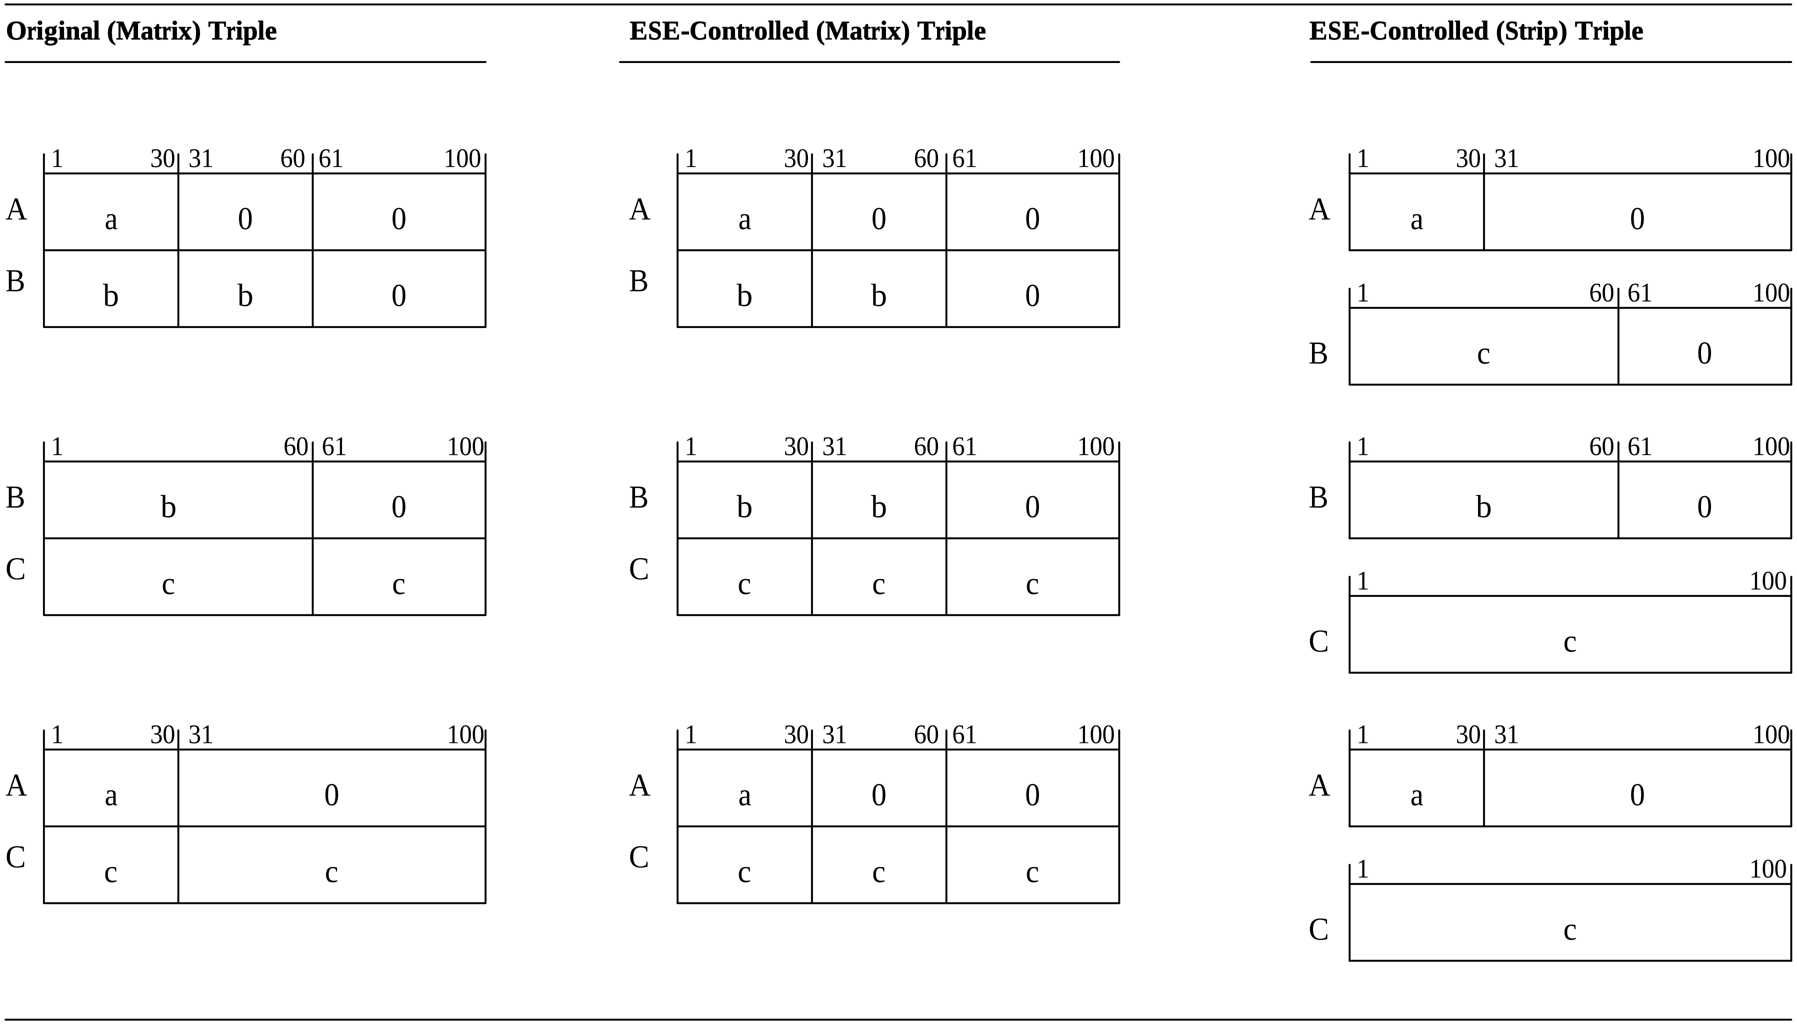

3. Previous evidence, event-splitting effects and Harless-framing

Starmer and Sugden (in press) report tests of cycling asymmetry (BCA>ABC) of the above type using three types of display. Examples of these displays are illustrated in figure 1. The rationale for comparing the incidence of cycling asymmetry between the displays stems from the speculation that the original displays (Loomes, Starmer and Sugden, 1989, 1991) may have been susceptible to ESEs (Starmer and Sugden,1993; Humphrey, 1995) or framing effects (Harless, 1992). Taking ESEs first, Starmer and Sugden (in press) draw attention to the fact that in the original matrix triples, outcome b is split between two events in the first decision problem, but represented as a single event in the second. This difference in the representation of the event leading to consequence b may contribute to cycling asymmetry by rendering option B more attractive in the first problem in the triple than in the second.

The prevailing explanation of ESEs is in terms of a modified version of prospect theory (Kahneman and Tversky, 1979). Assume that an individual is faced with a choice between two prospects, P_i (i=1,2), where each prospect is represented by a vector of probabilities p_{ij} (j=1,..,n) with p_{ij} being the probability that prospect P_i results in consequence x_j. According to prospect theory individuals will then attempt to maximise the value of

$$\Sigma_j \pi(p_{ij}) v(x_j) \tag{3}$$

where $\pi(p_{ij})$ is the subjective probability assessment of the objective probability p_{ij} and v(.) is the utility function assigned to increments or decrements of wealth relative to an individual's reference asset position, being unique up to multiplication by a positive constant with $v(0)=0$ at the point of reference. If prospect theory is to explain ESEs it is necessary to make an assumption concerning the manner in which subjective decision weights are assigned to outcomes which occur in more than one state of the world. Starmer (1992) calls this assumption *post-combination*, which states that subjects apply the weighting function π(.) to the probabilities associated with each disjoint event separately, irrespective of whether there are two (or more) disjoint events in which identical outcomes occur. If post-combination and expression (3) are applied to the original (matrix) triple in figure 1 expression (4) is derived, which is required to hold to observe a clockwise violation of transitivity:

$$v(\mathrm{b})[\pi(0.3)+\pi(0.3)-\pi(0.6)] > 0 \qquad (4)$$

Expression (4) will hold if the weighting function $\pi(.)$ is *subadditive* in the relevant region and modified prospect theory explains the original observation of cycling asymmetry[5].

It should be pointed out that in prospect theory it is normally assumed that the utility of zero is zero ($v(0)=0$), and this causes zero consequences to drop out of value functions. It is possible, however, that ESEs are driven by an aversion to the more frequent representation of the zero consequence in prospect A in the first problem in the triple than in prospect A in the third problem. If ESEs are driving observed cycling asymmetry it should be recognised that the split of zero outcomes in prospect A between the first and third problems, and not the split of the non-zero consequence b, could be exerting the major impact. In this sense, if zero consequences are coded as losses, ESEs may be related to loss-aversion. It is, of course, also possible that a combination of splitting *both* b and zero is driving ESEs and potentially contributing to cycling asymmetry[6].

The ESE-controlled matrix and strip displays control for ESEs by ensuring that each consequence is listed with identical frequency in every decision problem in which it is present throughout the standard triple. The ESE-controlled strip display is also designed to control for framing effects. Harless (1992) reports tests of regret theory which compare matrix displays with strip displays, and discovers that evidence of

[5] The modification is to replace the act of combination from the editing stage of prospect theory with the assumption of post-combination. By postulating that the probabilities of the two b consequences in the first problem would be added prior to the evaluation of the prospect (thus becoming identical to prospect B in the second problem), combination precludes ESEs.

regret observed under the former disappears under the latter. Starmer and Sugden (in press) speculate that matrix displays provide a representation of events which is more visually explicit than the strip display and are therefore more likely to induce ‘up and down’ comparisons within events and between acts. This might serve to enhance the salience of anticipated regret and one might therefore expect cycling asymmetry, if driven by regret-aversion, to diminish in the strip display in relation to the matrix displays.

In light of the possible contribution of ESEs and Harless-framing to the original observation of cycling asymmetry, Starmer and Sugden’s (in press) strategy is, first, to compare responses to the original matrix problems and the ESE-controlled matrix problems. If asymmetric cycling is due to ESEs it should be observed under the former and not under the latter. Second, they compare the ESE-controlled matrix problems to the ESE-controlled strip problems. If asymmetric cycling is due to regret-aversion and influenced by Harless-framing, it should be observed under the former and not under the latter. They in fact observe that asymmetric cycling occurs in the original matrix displays, persists at a similar rate in the ESE-controlled matrix display, persists for some parameters but not for others in the ESE-controlled strip display, and where it persists in strip displays clockwise (regret) cycles occur at a diminished rate. The conclusion they propose is that asymmetric (regret) cycling is not attributable to ESEs, but is sensitive to Harless-framing such that it is more likely to be observed in matrix displays which render cross-act, within-event comparisons more salient than other (i.e. strip) displays.

Humphrey (1998) investigates Starmer and Sugden’s (in press) conclusion by disaggregating their data to the level of the individual decision problem and conducting a between-subject comparison of responses to individual decision problems (i.e. the first, second and third problems in each triple) between the

[6] See Humphrey (in press) for a loss-aversion interpretation of ESEs.

displays[7]. This analysis places a different slant on their conclusions. First, the similar incidence of cycling asymmetry between the uncontrolled and controlled matrix displays may stem from differences between the sample groups and not differences between the decision problems. For example, with probabilities as illustrated in figure 1 and consequences of a=£18, b=£8 and c=£4, Starmer and Sugden's data reveal a significant reduction in the choice of option B in the first problem in the controlled display in relation to the uncontrolled display – despite the problems being identical in all respects[8]. This suggests a higher degree of risk-aversion in the group which faced the controlled problems than in the group which faced the uncontrolled problems. A similar comparison for the second decision problems in the triples reveals identical proportions of subjects choosing each option, despite the version of prospect theory which explains ESEs predicting an increase in the proportion of B choices in controlled matrix display relative to the uncontrolled matrix display[9]. This, observation, however, is consistent with event-splitting causing a switch towards B choices in the controlled display but the greater risk-aversion in the group which faced this display, as intimated above, off-setting the switch through a greater tendency towards C choices. In the third decision problem in the triples event-splitting would yield more C choices in the controlled display than in the uncontrolled display[10]. Starmer and Sugden's data do in fact indicate a significant increase in the proportion of option C chosen (49/90 vs. 61/90,

---

[7] Humphrey (1998) is concerned with the decision problems which constitute Starmer and Sugden's (in press) cases 1, 2, 3, 4, 5, 6, 8, 9 and 16. See Starmer and Sugden for the specific details of these triples, and for the other cases they report which are concerned with hypotheses not directly relevant to the discussion here.

[8] The results Starmer and Sugden (in press) report for this triple are 12.2% BCA vs. 2.2% ABC cycling in the uncontrolled matrix display and 11.1% BCA vs. 3.3% ABC cycling in the controlled matrix display. Both of these asymmetries are significant at the 5% level.

[9] Splitting the event which contains consequence b in the second problem in the controlled matrix display also entails splitting an event which contains consequence c. In the original version of prospect theory the decision weight attached to a certainty is $p(1)=1$, and so removing the act of *combination* (which would combine the probabilities associated with b in option B prior to evaluation of the prospect) yields the straightforward prediction of option B being more attractive in the controlled display than in the uncontrolled display, since by subadditivity $2p(0.3)>p(0.6)$. Humphrey (1998) shows that a similar prediction stems from prospect theory minus the editing stage even if it is not required that $p(1)=1$.

[10] This assumes that combination is not performed and it is not required that $p(1)=1$.

with a test of difference in sample proportions based on the normal distribution yielding z=-1.835), but this could, of course, also be the result of greater risk-aversion in the controlled matrix group.

Starmer and Sugden (in press) recognise that since they cannot be sure their sample groups were drawn from the same population, no completely controlled tests can be conducted. Nevertheless, their tests do not clarify what is causing the persistence of asymmetric cycling between the uncontrolled and controlled matrix triples. Greater risk-aversion in the controlled matrix group, being off-set in the second problem and (possibly) bolstered in the third problem by event-splitting could be responsible for the persistence, at a similar rate, of cycling asymmetry in this display[11]. There is similar ambiguity about what is causing Starmer and Sugden's general observation of reduced cycling asymmetry and, in particular, the reduced incidence of clockwise cycles between the matrix and strip displays. The conclusion they draw is that this constitutes evidence of Harless-framing. Humphrey's (1998) analysis of each individual decision problem between the displays also casts this conclusion under some doubt. Of the 30 comparisons between individual matrix problems and their strip display counterparts that Humphrey (1998) conducts, 16 show a significant difference between the proportions of each option chosen. All of these 16 cases could be explained by increased risk-seeking by the groups which faced the strip displays. Of these 16 cases, however, 14 are consistent with between-subject event-splitting. So, although event-splitting cannot be responsible for the within-subject cycling asymmetry observed in some of the strip display triples, it might explain differences between the incidence of cycling behaviour between the displays.

The rest of this paper discusses the design, implementation and results of an experiment which seeks to clarify some of the ambiguities in Starmer and Sugden's (in press) data. The design is not specifically geared towards discriminating between

[11] Note that ESEs, and the modified version of prospect theory which explains them, cannot be responsible for within (controlled matrix) group cycling asymmetry.

ESEs and Harless-framing. Rather, the strategy is to concentrate on discovering whether the ESEs hypothesis can be rejected. If it can, and the data indicate similar patterns to Starmer and Sugden's (in press), then one implication would be that the Harless-framing argument is worthy of further investigation. This said, however, the design does allow a degree of discrimination between ESEs and Harless-framing, and this will be discussed in the results section. The importance of accurately narrowing-down the causes of systematic non-transitive choices lies in conformity with the transitivity axiom being generally considered an essential component of consistent rational decision-making. In this sense the experiment reported here is not primarily concerned with adding to the volumes of documented systematic violations of one or more axioms of EUT. Rather, if such violations of EUT potentially imply significant welfare losses for decision-makers – and violations of transitivity in the form of being money-pumped can imply perhaps the ultimate welfare loss of bankruptcy – then understanding the possible antecedents of such behaviours is of fundamental importance to isolating contexts in which they might occur and assessing any associated welfare implications. In terms of decision theory, discriminating between anticipated regret (and whether it is influenced by Harless-framing) and the modified version of prospect theory which explains ESEs is important in determining how individuals make choices in decision problems of this type. Since the former entails cross-act, within-event evaluations of decision problems and the latter involves within-prospect evaluations, the conceptions of the decision-making process are very different.

4. Experimental design

The experiment is designed around two types of standard triple of the general form illustrated in table 1, one with two sets of parameters and the other with one. The parameter sets are outlined in table 2. Recall that each triple involves three choices; {A or B}, {B or C} and {A or C}. In terms of the frequency with which positive outcomes are represented in the decision problems in Starmer and Sugden's design, triple 1 corresponds to the ESE-uncontrolled triple, triple 2 to their ESE-controlled triples and triple 3 to their 'strip' triples.

TABLE 2: Parameters

		Prospect/ Act	**High-Outcome**	**Freq**	p_1	p_2	**Low-Outcome**	**Freq**	***Prob***
SET 1	***Triple 1***	A	18	1	0.3	-	0	1	0.7
		B_1	8	2	0.3	0.3	0	1	0.4
		B_2	8	1	0.6	-	0	1	0.4
		C_2	4	2	0.6	0.4	-	-	-
		C_3	4	2	0.3	0.7	-	-	-
	Triple 2	B	8	2	0.3	0.3	0	1	0.4
	Triple 3	B	8	1	0.6	-	0	1	0.4
		C	4	1	1.0	-	-	-	-
SET 2		*Same as above, but with high consequences* a=31.5, b=14, c=7							
SET 3	***Triple 1***	A	32	1	0.3	-	0	1	0.7
		B_1	15	2	0.3	0.25	0	1	0.45
		B_2	15	1	0.55	-	0	1	0.45
		C	12.5	2	0.3	0.4	0	1	0.3
	Triple 2	B	15	2	0.3	0.25	0	1	0.45
	Triple 3	B	15	1	0.55	-	0	1	0.45
		C	12.5	1	0.7	-	0	1	0.3

All decision problems are displayed in the strip format illustrated in figure 1, but the frequency with which each consequence is listed in decision problems in which it occurs is varied. The subscripts in the prospect column denote the decision problem in which that prospect occurs. For example, under set 1 and triple 1, prospect B_1 occurs in the first decision problem and offers £8 with probability 0.6, being split between two events each occurring with probability 0.3. Prospect B_2, on the other hand, occurs in the second decision problem in the triple and similarly offers £8 with probability 0.6, but here the event is not split. The 'Freq' column refers to the number of disjoint events in which a particular outcome occurs in a prospect. Where a prospect is not denoted with a subscript it appears as stated in every decision problem in the triple in which it is represented. Table 2 takes triple 1 as the base triple for each parameter set, and only illustrates how each subsequent triple in the set differs from this. For example, in set 1, the only difference between triple 1 and triple 2 is that prospect B occurs in both the first and second decision problems in triple 2 as opposed to B_1 and B_2, respectively, in triple 1.

All three parameter sets involve triples 1, 2 and 3. The main difference between set 3 and sets 1 and 2 is that set 3 involves a different probability distribution and, in particular, does not involve certainties. It is possible that clockwise cycles are driven by some decision-making heuristic which involves the certainty effect (Kahneman and Tversky, 1979), and set 3 is designed to shed light on this possibility. For example, if in triples such as sets 1 and 2 subjects are approximately indifferent between A and B in the first problem, generally attracted to the certainties in the second and third decision problems, but are sometimes tempted away from the certainty in the third problem by the highest payoff in the triple (a) being offered in option A (and not drawn away from the certainty when it is paired option B, which offers only an intermediate outcome), then one might observe significantly more BCA choices than ABC choices. Consequently, if apparent clockwise cycles are due to some decision-making heuristic which involves the certainty effect, then they should perhaps occur with a greater incidence in sets 1 and 2 than in set 3.

Within the context of the above design the general strategy is similar to that employed by Starmer and Sugden (in press); that is, to compare the incidences of cycling asymmetry between the different triples. The design is a *complete within-subject design*, and this precludes any ambiguities in the results such as those which apparently exist between-groups in Starmer and Sugden's data. For example, table 2 illustrates that the only difference between triples 1 and 2, under parameter set 1, is the second decision problem. Problems 1 and 3 in each triple are identical, will only be faced once by each subject, and therefore cannot entail choices in identical decision problems differing significantly between groups as is evident in some of Starmer and Sugden's cases. Also note from table 2 that in non-certain options the low outcome is zero in every triple for every parameter set, and that these are represented with unitary frequency throughout. This controls for the 'loss-aversion' interpretation of ESEs discussed in section 3, and in more detail in Humphrey (in press), and provides the purest test of whether the original observation of cycling asymmetry is due to splitting the event containing consequence b in the first decision problem and not in the second. More generally, the complete within-subject design, with subjects being drawn from the same population and all responding to the complete set of identical decision problems, allows a test of whether any differences in the incidence of cycling asymmetry stem from within-subject ESEs and, specifically, ESEs in particular decision problems between triples. Table 3 illustrates the predictions of the modified version of prospect theory which explains ESEs for each individual problem between each triple.

TABLE 3: ESEs between triples

Parameter Set	Problem	Triple Pair		
		1-2	**1-3**	**2-3**
Sets 1 and 2	1	none	BA	BA
	2	CB	CB $\pi(1)\neq 1$	BC $\pi(1)=1$
			none $\pi(1)=1$	BC/CB $\pi(1)\neq 1$
	3	none	CA $\pi(1)\neq 1$	CA $\pi(1)\neq 1$
			none $\pi(1)=1$	none $\pi(1)=1$
Set 3	1	none	BA	BA
	2	CB	CB	BC/CB
	3	none	CA	CA

As far as EUT is concerned each decision problem (1, 2 or 3) is identical between the triples. Violations of EUT involve switching from the safer option in the first triple problem to the riskier option in the second triple problem in the comparison, or vice-versa. These predictions of modified prospect theory may differ depending on whether $\pi(1)=1$ (as in standard prospect theory) or $\pi(1)\neq 1$. So, for example, 'BA' denotes switching from choosing B in the first problem in the comparison to A in the second, 'none' denotes comparisons where modified prospect theory does not make a prediction (because the problems are identical).

5. Experimental implementation

The experiment was conducted between Christmas and Easter 1998 with a total of 80 subjects, 74% of which were male. Participants were mainly social science undergraduates recruited by e-mail shot which invited them to book a place in one of a number of pre-scheduled sessions on the campus of the University of Nottingham. Each subject participated in the experiment only once and faced a total of 60 decision problems which were presented in sequence on the screen of a computer terminal[12]. Of these 60 problems, 21 belonged to the triples outlined in table 2, 6 were repeat problems from sets 1 and 2 to test the consistency of responses over decision problems which are faced more than once and identical in all respects. The remaining problems tested hypotheses not of direct concern to the present discussion. Each session lasted for approximately 45 minutes and the average payout (paid on the spot) was £8.72.

The experiment employed some standard controls in designs of this type. First, the order in which each subject faced each individual decision problem was randomised. If any systematic violations of EUT are observed they cannot, therefore, be due to the order in which individuals faced decision problems. Second, the display of individual decision problems was randomised such that the riskier and safer options could appear either above or below each other in the display and, although the juxtapositions of the consequences were maintained throughout the experiment, the positions of the consequences in the displays were randomised between left and right. In terms of the first problem in the strip triple in figure 1, for example, this problem would be equally likely to appear exactly as it is in that figure or as a horizontal and/or vertical mirror-image of itself. Third, the experiment employed a random-lottery incentive system whereby after all decision problems had been answered one was randomly selected and played-out for real money. Subjects were informed of this prior to responding to the problems. Holt (1986) argues that if subjects do not adhere to the independence axiom of EUT, random lottery designs may not elicit subjects' true preferences. Starmer and Sugden (in press) point out that tests of cycling asymmetry of the type reported here are not affected by Holt's argument and draw the reader's attention to Loomes, Starmer and Sugden (1991) for a full discussion of why this is the case[13].

---

[12] The experimental software was programmed in Toolbook version 3.0. The computerised displays captured the essential features of the strip displays illustrated in figure 1, and presented them in a windows format.

[13] Table 3 outlines the tests for ESEs and these tests are similarly unaffected by Holt's criticism of random lottery designs. As Starmer and Sugden (1993, p.245) point out, the potential bias Holt

6. Results

6.1. Cycling asymmetry

Table 4 presents the results of the tests of cycling asymmetry for each triple in each parameter set. The first point to note from table 4 is that in the 9 cases reported 8 indicate cycling asymmetry with a significance greater than 5% and of the order

TABLE 4: Results: Cycling asymmetry

Set	Triple	**ABC**	ABA	ACC	ACA	BBC	BBA	BCC	**BCA**	C-WISE %	AC-WISE %	*p*
1	1	**2**	15	3	4	7	19	19	**11**	13.8	2.5	0.0112*
1	2	**0**	9	1	5	9	8	35	**13**	16.3	0.0	0.0001*
1	3	**1**	9	5	2	11	9	36	**7**	8.8	1.3	0.0352*
2	1	**0**	8	4	3	9	8	34	**14**	17.5	0.0	0.0001*
2	2	**0**	12	1	2	14	13	30	**8**	10.0	0.0	0.0039*
2	3	**0**	18	3	3	10	18	19	**9**	11.3	0.0	0.0020*
3	1	**2**	15	3	3	6	15	22	**14**	17.5	2.5	0.0021*
3	2	**0**	4	0	1	10	25	34	**6**	7.5	0.0	0.0156*
3	3	**2**	7	5	4	8	10	36	**8**	10.0	2.5	0.0547

Column groups: Patterns of Choice over the triple (ABC … BCA); CYCLES (C-WISE %, AC-WISE %).

The cycles column shows the percentage of non-transitive choices in the clockwise (BCA) regret-consistent direction and the opposite (ABC) regret-inconsistent direction. The ***p*** column reports a test statistic based on the binomial distribution of the hypothesis that clockwise cycles occur with a greater frequency than anti-clockwise cycles. The null hypothesis is that they occur with equal frequency. An asterisk denotes rejection of the null hypothesis in favour of the alternative with 5% significance.

previously observed (Loomes, Starmer and Sugden, 1989, 1991; Starmer and Sugden, in press). Cycling asymmetry appears to be a robust phenomenon; it appears in all parameter sets, and in this sense is apparently not restricted to triples involving certainties. As discussed in sections 3 and 4, the zero consequence in triple 1 is represented once in every decision problem in which it is present. This, unlike the Starmer and Sugden (in press) design, controls for the loss-aversion interpretation of ESEs (Humphrey, in press). Consequently, the observation of significant cycling asymmetry over triple 1 in sets 1, 2 and 3 suggests that, if ESEs

---

describes cannot exert any influence over pairs of decision problems which involve identical probability distributions over identical consequences. See also Starmer and Sugden (1991).

are responsible for clockwise cycles, it is not coding of zero outcomes as losses and an aversion to greater frequencies of losses which is driving choices.

Recall that Starmer and Sugden's (in press) data revealed a general pattern of similar incidences of clockwise cycles (around 12%) in their ESE-controlled (triple 1) and ESE-uncontrolled matrix displays (triple 2) and a lower incidence in their strip displays (around 9%)[14]. Table 4 shows the incidences of clockwise cycles observed in this experiment to be slightly different. In parameter sets 2 and 3 clockwise cycles occur with similar frequency in triples 2 and 3 and are less frequent than in triple 1. Since triples 2 and 3 both control for ESEs this observation, if taken at face value, would suggest that clockwise cycles are partially driven by ESEs and not Harless-framing. If correct, this conclusion would be contrary to that reached by Starmer and Sugden (in press). Parameter set 1, on the other hand, which employs identical parameters to one of Starmer and Sugden's cases, would tend to support their findings. The similar incidence of clockwise cycles between triples 1 and 2 and lower incidence in triple 3 lending support to the view that ESEs are not as important as Harless-framing in generating cycling asymmetries. These prima facie conclusions will be investigated in more depth in the next section.

The observation of cycling asymmetry in all parameter sets is interesting in the context of the debate over whether violations of EUT are more significant the lower are decision-makers' financial incentives. Hey (1991, p.5), for example, argues that violations will diminish or disappear if decisions involve high enough stakes. Thaler (1987), on the other hand, discusses a set of 'parallel' experiments in which violations of EUT tend to be more significant in the presence of incentives than in their absence. Parameter set 1 in this experiment replicates one of Starmer and Sugden's (in press) cases, but sets 2 and 3 involve substantially higher incentives. In

[14] These figure relate to the triple which involve the same parameters as set 1 here. For Starmer and Sugden's (in press) triples 3, 4, 6, 9 and 16, which are similar but involve prospects where the smallest consequence is positive, the incidence of clockwise cycles in the matrix displays is around 19.5% and in the strip displays around 9%.

particular, parameter set 2 employs identical probability distributions, but the consequence values render the expected values of each prospect 75% greater than the corresponding set 1 prospects. Yet there is little evidence of diminished violation of transitivity, with a total of 13.75% violation across set 1 being compared to 11.88% across set 2. Set 3 entails the highest expected value, but total violations persist at 13.33%. Consequently, there is little overall evidence that increasing incentives by the order of around 75% of expected value from the levels at which cycling asymmetry has been previously observed reduces the incidence of violations of transitivity. It is of course possible that further increases in incentives would reduce the incidence of violations[15].

6.2. Harless-framing or ESEs?

Having replicated cycling asymmetry over a variety of displays and parameter sets the question turns to a more detailed investigation of the above prima facie evidence regarding whether the incidence of clockwise (BCA) cycles is influenced by either ESEs or Harless-framing. The first point to note in this respect is that on the basis of a (one-tailed) test of differences in sample proportions using the normal distribution, table 4 shows that the only significant difference (at 5%) between the incidence of clockwise cycles is between triples 1 (14/80) and 2 (6/80) in set 3 (z=1.9124). Consequently, in a similar manner to Starmer and Sugden (in press), it should be understood that the discussion is concerned with establishing whether ESEs or Harless-framing *contribute* to significant cycling asymmetry by influencing the frequency of clockwise cycles without necessarily rendering differences in the incidences of clockwise cycles between triples *significantly* different in themselves. Since the experiment was designed specifically to shed light on whether ESEs

[15] Should this be true, it is unclear what conclusion should be reached. If decisions involving consequences of up to £32.00 (some decision problems in the experiment which are not discussed here involved consequences of up to £37.00) and expected values of around £9.00 are to be classed as decisions involving stakes too low for EUT to be applicable, then proponents of EUT who use this argument are, by negating the importance of many day-to-day decisions which involve comparable stakes, effectively negating the importance of EUT itself.

contribute to different incidences of clockwise cycles, this will form the initial focus of attention.

TABLE 5: Results: Event-splitting effects

Triple Pair	Patterns of Choice over the problem pair														
	Question 1					**Question 2**					**Question 3**				
	AA	**BB**	**AB**	**BA**	***p***	**BB**	**CC**	**BC**	**CB**	***p***	**AA**	**CC**	**AC**	**CA**	***p***
1-2															
SET 1	-					35	26	8	***11***	0.3238	-				
SET 2	-					21	38	3	***18***	0.0007*	-				
SET 3	-					2	64	6	***8***	0.3953	-				
1-3															
SET 1	14	47	10	***9***	0.6762	34	30	10	***6***	0.8949†	42	25	8	***5***	0.8666†
SET 2	3	55	7	***10***	0.3145	15	44	9	***12***	0.3318†	23	41	10	***6***	0.8949†
SET 3	10	49	10	***11***	0.5000	3	64	5	***8***	0.2905	8	62	3	***7***	0.1719
2-3															
SET 1	14	47	10	***9***	0.6762	31	26	***15***	8	0.1050†	42	25	8	***5***	0.8666†
SET 2	3	55	7	***10***	0.3145	22	36	***17***	5	0.0085†*	23	41	10	***6***	0.8949†
SET 3	10	49	10	***11***	0.5000	2	61	8	***9***	0.5000†	8	62	3	***7***	0.1719

The three blocks of columns in table 5 show the patterns of choice between pairs of different triples in each parameter set for the first second and third decisions problems in the triple-pairs, respectively. Conjunctions of choices involving the same option being chosen in each problem (e.g. AA or BB) are consistent with EUT. The *p* column shows a test statistic based on the binomial distribution that violations of EUT consistent with ESEs (predicted by modified prospect theory) occur with a greater frequency than the opposite violation (the null hypothesis is that they occur with equal frequency). An asterisk indicates rejection of the null and a significant ESE at the 5% level. The bold italicised choice frequencies are those consistent with ESEs and are used as the basis for the test. Table 3 illustrates the predictions of modified prospect theory to be indeterminate or contingent on whether $\pi(1)=1$ or $\pi(1)\neq 1$, and these cases are denoted '†' in table 5.

Table 5 provides the results of a test based on the binomial distribution of whether there exists any significant ESEs between individual problems in a triple between each type of triple in a parameter set. The null hypothesis is that there are no ESEs and the alternative hypothesis is ESEs occur as described in table 3. All tests are within-subject. The choice patterns italicised in table 5 are those consistent with ESEs. Table 5 should be read in conjunction with table 4.

Consider parameter set 1 first. Under parameter set 1 the incidence of clockwise cycles falls from 13.8% in triple 1 to 8.8% in triple 3. This suggests that either ESEs, Harless-framing or a combination of both are influencing choices. Table 5 shows that there are no significant ESEs between triples 1 and 3 under set 1. Note, however, that an inspection of table 3 suggests that if the decision weight $p(1)=1$, then no ESEs are predicted. The suggestion that ESEs are not responsible for reduced clockwise cycles between triples 1 and 3 under set 1 is confirmed by comparing triples 1 and 2. Here the incidence of clockwise cycles remains of a broadly similar order and, if clockwise cycles are driven by ESEs, one would expect fewer cycles under triple 2 than under triple 1. Table 5 reveals no significant ESE between the second problems in these triples. Similarly, a comparison of the two ESE-controlled triples, 2 and 3, reveals a reduction in clockwise cycles from 16.3% to 8.8% which cannot, according to table 5, be explained by significant ESEs. Thus, in parameter set 1, which employs identical parameters to of one of Starmer and Sugden's (in press) cases and closely replicates their findings, is seems that ESEs can be ruled out as explaining different incidences of clockwise cycles.

Under parameter set 2 the incidence of clockwise cycles is greater under triple 1 (17.5%) than under sets 2 (10.0%) and 3 (11.3%), which are of similar order. This general pattern of choices is consistent with ESEs contributing to clockwise cycles under triple 1, but not under either of the ESE-controlled triples. Table 5, however, reveals no significant ESEs between triples 1 and 3, which would be expected if the ESE argument is correct, and creates a suspicion that a degree of Harless-framing may be present. The lack of ESEs between triples 1 and 3 could be explained in the second and third decision problems if subjects recognise certainties such that $p(1)=1$, but this does not explain the fall in clockwise cycles between triples 1 and 3. The influence of ESEs is, however, evident between triples 1 and 2. Table 5 shows a significant ESE in the second decision problem between the triples which operates in the direction consistent with a reduced tendency for clockwise cycling. This is contrary to Starmer and Sugden's (in press) findings. Finally under set 2, if the suspicion of Harless-framing is to be substantiated one would expect fewer clockwise cycles in triple 3 than in triple 2, and this is not the case. Table 5, however, reveals a significant ESE in the second decision problem between triples 2 and 3 which operates in the direction consistent with increasing clockwise cycles. It cannot be ruled-out, therefore, that Harless-framing generates a tendency towards fewer clockwise cycles in triple 3 than in triple 2, but that this is being offset by ESEs in the second decision problem such that a similar order of clockwise cycling is actually observed.

Choices made over parameter set 3 decision problems exhibit a similar pattern to those under set 2. There is a greater incidence of clockwise cycles under triple 1 (17.5%) than under triples 2 (7.5%) and 3 (10.0%), which again suggests that ESEs are an important determinant of the extent of this kind of non-transitive choice behaviour. However, it is again the case that table 5 indicates no significant ESEs between triples 1 and 3. The same is true between triples 1 and 2 and between triples 2 and 3. Overall, then, there is only partial evidence supporting the conjecture that ESEs drive different incidences of clockwise cycles between the different triples. ESEs are ruled-out as explaining this aspect of the data in parameter set 1, can explain some of the observations under parameter set 2 (but without ruling-out the influence of Harless-framing), and are not present at all in parameter set 3. Consequently, it is necessary to consider in more detail whether Harless-framing explains the aspects of the data not appropriately organised by the ESE hypothesis.

First, given that this experiment employed strip displays throughout, there appears to be little support for the suggestion that matrix displays *per se* are more likely to yield the kind of cycling asymmetry observed over standard triples of this type. The point about strip displays, however, is not so much that they separate each option vertically, but rather that they 'un-split' or collapse disjoint events containing identical consequences to a single disjoint event. Thus, strip displays remove the disjoint event 'columns' which are visually explicit in matrix displays and, according to Harless (1992), render anticipated regret less salient. It is in this respect that the important facet of Harless's (1992) argument is the ease with which 'up and down' comparisons within-events and between-acts can be performed. Given that this is partially determined by the number of disjoint events used to display the options which comprise a decision problem and that this can also create the potential for ESEs, the question becomes one of how to distinguish between the influence of Harless-framing and ESEs.

The strategy adopted here is to consider whether there is any evidence that the incidence of clockwise cycles is influenced by the ease with which 'up and down' comparisons, within-events and between-acts, can be performed. It is difficult to ascertain exactly how difficult such a comparison is, but one possible proxy can be illustrated by comparing the first decision problems in the matrix triples and the strip triple in figure 1. The matrix triples entail explicit disjoint event columns, but the strip triple does not. Figure 1 shows that the first problem in the strip triple has ***two*** vertical lines (one in each option) which separate events which yield different outcomes. Neither of these lines have a directly vertically located counterpart in the other option. Thus, if this makes 'up and down' comparisons more difficult, the first problem in the strip triple can be given a 'difficulty rating' of 2. In a similar manner the first problems in the matrix displays would receive a zero difficulty rating. If this line of reasoning is extended, an ordinal difficulty index can be obtained for all decision problems in all triples employed in the experiment. This is illustrated in table 6.

TABLE 6: Difficulty of 'up and down' comparisons

Problem in Triple	**Triple Type**					
	Parameter Sets 1 and 2			Parameter Set 3		
	1	2	3	1	2	3
1	0	0	2	1	1	2
2	0	0	1	2	1	2
3	0	0	1	0	0	0
Total	**0**	**0**	**4**	**3**	**2**	**4**

On the basis of the difficulty proxy, table 6 shows that for parameter sets 1 and 2, triples 1 and 2 entail the least difficulty in performing ‘up and down’ comparisons (with zero difficulty rating) and triple 3 entails the most (rating 4). An inspection of table 4 shows that the (ascending) order of incidence of clockwise cycles for parameter set 1 is (triple) 3, 1, 2. This corresponds to the difficulty ranking of 1 or 2, 3. For parameter set 2 the ascending clockwise cycle ranking is triple 2, 3, 1, which corresponds less well with the difficulty ranking. The ascending clockwise cycle ranking of triples 2, 3, 1 in parameter set 3 does not correspond particularly well with the difficulty ranking of 3, 1, 2. Thus, on the basis of the suggested proxy for the difficulty of performing ‘up and down’ comparisons, there appears to be only partial support for Harless-framing. The best support is provided by parameter set 1, where ESEs are ruled-out and the difficulty ranking corresponds exactly to the frequency of clockwise cycle ranking.

If one interprets the difficulty index in a looser sense and accepts that it does not represent a complete transitive ranking, but does *indicate* how frequencies of clockwise cycles between *pairs* of triples within a parameter set are influenced by Harless-framing, then more support can perhaps be found. For example, in addition to the correlation mentioned for set 1, the frequency differences of clockwise cycles then corresponds to relative comparison difficulties for triple *pairs* 1 and 3 in both parameter sets 2 and 3. This serves to clarify the suspicion of Harless-framing discussed above in relation to these triple pairs and parameter sets. An interesting feature of the data, however, is that the only significant difference (z=1.9124, one-tailed) in clockwise cycles occurs between triples 1 (14/80) and 2 (6/80) in set 3, and this is in the *opposite* direction to that which table 6 suggests would be consistent with the Harless-framing argument. Yet although the split of clockwise cycles is consistent with ESEs, ESEs neither seem to organise the data.

As far as Harless-framing is concerned, it seems that support is somewhat patchy. Interestingly, the clearest results emerge from parameter set 1 which is identical to one of the parameter sets employed by Starmer and Sugden (in press). Here it does indeed seem that choices are governed by some kind of decision heuristic which is captured by regret theory, and that the extent to which the heuristic generates clockwise violations of transitivity is influenced by the ease with which within-event and between-act comparisons can be performed. More generally, if Harless-framing exists, and there is some evidence that it might, it is either a weak phenomenon, or the differences in the difficulty of performing ‘up and down’ comparisons within events and between acts were not sufficiently different between the triples employed in this experiment to allow it to emerge with any general significance. It is, of course, possible that an alternative proxy for difficulty might be more revealing.

6.3. Repeat problems

It is sometimes argued that, by their nature, one-shot experiments provide little opportunity for learning and that violations of EUT might in fact be mistakes arising out of a lack of this learning opportunity[16]. Chu and Chu (1990), for example, show that preference reversals - which Loomes, Starmer and Sugden (1989) relate to clockwise cycles – disappear after several iterations if subjects experience the consequences of their choices after each repetition. This experiment involved 60 decision problems which, although concerned with different hypotheses, were of a generally similar type. It is, therefore, arguable that sufficient opportunity to learn and develop response strategies was afforded in this experiment. Moreover, given the randomisation of problem order, the experiment included appropriate control such that systematic violations of EUT of particular triples or problem pairs cannot reflect the general nature of any learning which occurred during the course of the responding to the set of problems – unless, of course, 60 decision problems is an insufficient number to facilitate the development of response strategies. To investigate this latter possibility two triples were repeated during the course of the 60 decision problems. Given the randomised problem order, in responding to a repeated triple subjects may have responded to any decision problem in the triple

[16] See Thaler (1987) for a more detailed discussion of this point.

twice before they responded to either of the other two decision problems for the first time, or they may equally have responded to all problems in the triple once before they face any of the problems for the second time, or any permutation in between. Table 7 presents the results of a test for cycling asymmetry in the repeated triples, and should be read in conjunction with the corresponding results in table 4.

TABLE 7: Repeat cycling asymmetry

Set	Triple	Patterns of Choice over the triple								CYCLES		*p*
		ABC	ABA	ACC	ACA	BBC	BBA	BCC	**BCA**	C-WISE %	AC-WISE %	
1	1	**2**	12	3	4	6	18	24	**11**	13.8	2.5	0.0112 *
2	1	**1**	9	1	6	6	12	34	**11**	13.8	1.3	0.0032*

Table 7 shows that although the patterns of choice over repeated triple 1 in parameter set 1 differ slightly from those in table 4, the incidences of both clockwise and anti-clockwise cycling are exactly the same. Similarly, under repeated triple 1 in parameter set 2 significant cycling asymmetry persists. Consequently, despite the slightly lower incidence of clockwise cycles following repetition under parameter set 2, and the possibility that additional repetitions may cause further subsidence, there is no convincing evidence that cycling asymmetry is not a genuine feature of individual decision-making behaviour over the standard triples used in this, and previous, experiments.

7. Conclusions

A number of interesting findings have emerged from the experiment reported here. First, significant cycling asymmetries of the type consistent with regret theory have emerged over a variety of parameter sets which both replicate previous findings and add to the existing evidence with parameter sets involving higher incentives. The volume of evidence illustrating this kind of decision-making behaviour reveals it to be extremely robust, and the evidence provided here suggests little reason to expect such behaviour to dissipate with repetition or experience. Second, by employing a complete within-subject design, the experiment controls for differences in group behaviour which seems to have muddied the water in Starmer and Sugden's (in press) recent investigation. Third, and perhaps most importantly, the experimental design facilitated a in-depth investigation of the possible antecedents of the extent of cycling asymmetry observed.

Overall, and somewhat surprisingly, the importance of ESEs seems to be limited. The data under parameter sets 1 and 3, for example, seem best organised by regret theory plus Harless-framing, and whist parameter set 2 does suggest the influence of ESEs, it also suggests the influence of Harless-framing. In fact, one of the more striking features of the data generated by this experiment in the lack of ESEs emerging between individual problems in different triples (table 5 shows only 2 significant ESEs out of 21 cases). In light of the apparent robustness of ESEs in previous studies, this finding is somewhat of a mystery. One explanation, however, would be the control for the loss-aversion interpretation of ESEs (Humphrey, in press) which was present in the experiment reported here, but not in other studies. In this respect it is possible that whist ESEs do not seem to be generally responsible for the incidence of clockwise cycles observed in this experiment, they could be responsible for behaviour observed in the original experiments (Loomes, Starmer and Sugden, 1989, 1991) and recent replications Starmer and Sugden (in press). What is clear is that although parameter set 2 precludes ruling-out the influence of ESEs altogether, ESEs are not necessary to observe cycling asymmetry and differences in the incidence of clockwise cycles between different types of displays.

Finally, as Starmer and Sugden (in press) conclude, it does indeed seem that cycling asymmetry stems from the use of a decision-making heuristic which is influenced by the ease with which individuals are able to within-event and between-act evaluations of decision problems. Although regret theory coupled with Harless-framing does provide a model of this kind of behaviour, and it seems likely that action-based considerations are as important as prospect-based considerations, it is less certain whether anticipated regret is actually driving choices. For example, in many of the decision problems in this experiment subjects would never know what consequence they would have experienced had they chosen the counterfactual (unless the counterfactual was a certainty) since the risk in counterfactual choice was not resolved. In this sense there may be very little regret for them to anticipate[17]. Given the importance of the transitivity axiom in theories of single-agent risky decision-making, further research is warranted to investigate the antecedents of systematically non-transitive choices over standard triples. Whilst the volume of evidence now available is unlikely to engender an abandonment of the transitivity axiom, it does merit addressing questions such as whether there exist market-like analogues of standard triples and, moreover, whether such potentially exploitable and welfare damaging behaviour persists therein.

---

[17] See Zeelenberg *et al.* (1996) who show that anticipation of resolving risk involved in the counterfactual choice increases the impact of anticipated regret. An experiment which investigates the impact of resolving counterfactual risk in standard triples of the kind discussed here will be reported soon.

References

Chu, Y. and R. Chu, 1990, The subsidence of preference reversal in simplified and marketlike experimental settings: A note, *American Economic Review 80*, 902-911.

Harless, D.W., 1992, Actions versus prospects: The effect of problem representation on regret, *American Economic Review 82*, 634-649.

Hey, J.D., 1991, *Experiments in Economics* (Blackwell, Oxford).

Holt, C.A., 1986, Preference reversals and the independence axiom, *American Economic Review 76*, 508-515.

Humphrey, S.J., 1995, Event-splitting effects or regret-aversion: More evidence under risk and uncertainty, *Journal of Risk and Uncertainty 11*, 263-274.

Humphrey, S.J., 1998, Non-transitive choice and the prediction of frames, *unpublished manuscript*.

Humphrey, S.J., in press, More mixed results on boundary effects, *Economics Letters*.

Kahneman, D. and A. Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica 47*, 263-291.

Loomes, G. and R. Sugden, 1987, Some implications of a more general form of regret theory, *Journal of Economic Theory 41*, 270-287.

Loomes, G., C. Starmer and R. Sugden, 1989, Preference reversal: Information-processing effect or rational non-transitive choice?, *Economic Journal 99*, 140-151.

Loomes, G., C. Starmer and R. Sugden, 1991, Observing violations of transitivity by experimental methods, *Econometrica 59*, 425-442.

Starmer, C., 1992, Violations of monotonicity – regret effects or probability processing effects? An experimental test, *Economics Research Centre Discussion Paper No. 9218*, University of East Anglia.

Starmer, C. and R. Sugden, 1991, Does the random-lottery incentive system elicit true preferences? An experimental investigation, *American Economic Review 81*, 971-978.

Starmer, C. and R. Sugden, 1993, Testing for juxtaposition and event-splitting effects, *Journal of Risk and Uncertainty 6*, 235-254.

Starmer, C. and R. Sugden, in press, Testing alternative explanations of cyclical choices, *Economica*.

Sugden, R., 1993, An axiomatic foundation for regret theory, *Journal of Economic Theory 60*, 159-180.

Thaler, R., 1987, The psychology of choice and the assumptions of economics, in A.E. Roth (ed.)., *Laboratory experimentation in economics: Six points of view* (Cambridge University Press, Cambridge).

von Neumann, J. and O. Morgenstern, 1944, *Theory of games and economic behavior* (Princeton University Press, Princeton).

Zeelenberg, M., Beattie, J., van der Plight, J. and N.K. de Vries, 1996, Consequences of regret aversion: Effects of expected feedback on risky decision making, *Organizational Behavior and Human Decision Processes 65*, 148-158.