Theory and Decision (2006) 61: 93–128
DOI 10.1007/s11238-006-8047-x
© Springer 2006

STEVEN J. HUMPHREY

DOES LEARNING DIMINISH VIOLATIONS OF INDEPENDENCE, COALESCING AND MONOTONICITY?

ABSTRACT. Violations of expected utility theory are sometimes attributed to imprecise preferences interacting with a lack of learning opportunity in the experimental laboratory. This paper reports an experimental test of whether a learning opportunity which engenders accurate probability assessments, by enhancing understanding of the meaning of stated probability information, causes anomalous behaviour to diminish. The data show that whilst in some cases expected utility maximising behaviour increases with the learning opportunity, so too do systematic violations. Therefore, there should be no presumption that anomalous behaviour under risk is transient and that discovered preferences will be appropriately described by expected utility theory.

KEY WORDS: discovered preferences, event-splitting effects, independence, monotonicity, probability learning.

1. INTRODUCTION

The decision-making under risk literature reports numerous experimentally observed violations of the axioms of expected utility theory. These observations challenge the predictions of broader economic theories which assume expected utility maximization and undermine the policy assessments emerging from them. More recently investigators have begun to speculate that apparently anomalous behaviour may stem from an inherent imprecision, or incompleteness, in preferences, coupled with unfamiliarity with the types of decision task faced in the experimental laboratory (Loomes, 1999). Plott (1996), for example, has argued that despite experimentally observed choice anomalies, expected utility theory is a descriptively appropriate approximation of individuals' true preferences. Given some kind of learning process and familiarisation with decision tasks, preferences will 'firm-up' and settle on expected utility maximization.[1] If Plott's (1996) argument is correct, the problems posed to economic analysis by experimentally observed violations of expected utility theory may be less serious than previously thought.

This paper reports a test of the impact of a probability learning opportunity on frequently observed choice anomalies: common consequence effect violations of the independence axiom, violations of the coalescing principle and violations of monotonicity. The experiment employs a lottery valuation task whereby conformity with or violation of expected utility maximisation is manifest in patterns of valuations over sets of individually valued lotteries. The vehicle though which the learning opportunity is offered is the revelation of the outcomes of 10 resolutions of the risk in each lottery prior to that lottery being valued. This learning opportunity has three important properties: (i) For reasons outlined below, it may be expected to promote the formation of accurate probability weights, and thereby give expected utility maximising behaviour a fair chance of being observed. (ii) It does not directly financially reward decision-makers for conformity with expected utility theory (as is the case in Chu and Chu's (1990) disciplining of transitivity violations). In this sense the experiment can be regarded as a test of whether decision-makers voluntarily learn to avoid anomalous behaviour. This is often the most economically relevant question when market discipline is either not present or not strong enough to enforce behaviour modifications. (iii) It is behaviourally relevant. Investors, for example, can observe the performance of stocks prior to investment decisions.

The test investigates two questions. First, does experience in observing the resolution of risk involved in lotteries prior to choice affect behaviour? Second, if probability learning affects revealed preferences, are those revealed preferences appropriately described by expected utility theory? Seeking an answer to these questions is important. One possibility is that Plott (1996) is right: learning opportunities may reveal behavioural anomalies to be transient laboratory-generated nuisances to the otherwise sound theory of rational choice. On the other hand, even if learning occurs in experimental decision-making tasks, there is no reason in principle why the discovery of genuine underlying preferences need imply expected utility maximisation (Loewenstein, 1999). If anomalies are genuine features of initially imprecise non-expected utility preferences, and if learning and experience reduce imprecision, deviations from expected utility maximisation may be expected to persist or even be exacerbated.

2. THREE VIOLATIONS OF EXPECTED UTILITY THEORY

2.1. *Independence violations: The common consequence effect*

The first panel of Table I describes a set of three decision problems. Each pair of lotteries is generated from each of the other pairs by shifting a probability mass of 0.5 between outcomes in both lotteries in a pair. For example, $S2$ and $R2$ are respectively generated from $S1$ and $R1$ by replacing a 0.5 chance of £9 with a 0.5 chance of zero. Since the 0.5 probability event contains outcomes common to each lottery within a pair, the independence axiom of expected utility theory states that either the riskier option is preferred in all three problems ($R1$, $R2$ and $R3$), the safer option is preferred in all three problems ($S1$, $S2$ and $S3$), or that there is indifference in all three problems. A common consequence effect occurs when preferences switch systematically between the riskier lottery and the safer lottery in a comparison of two lottery pairs.

Wu and Gonzalez (1998) discuss a large amount of experimental evidence which shows that the usual pattern of observed common consequence effects is as described in Table II.

2.2. *Coalescing violations: Event-splitting effects*

Event-splitting effects are observed by Starmer and Sugden (1993) and Humphrey (1995) over pairs of decision problems

TABLE I
Lotteries

Panel 1: Common Consequence Effect Lotteries

	Lottery	Probability	0.1	0.1	0.3	0.5
Pair 1	*S1*	Outcomes	0	9	9	9
	R1		21	21	0	9
Pair 2	*S2*	Outcomes	0	9	9	0
	R2		21	21	0	0
Pair 3	*S3*	Outcomes	0	9	9	21
	R3		21	21	0	21

Panel 2: Event-Splitting Effect Lotteries[a]

	Lottery				
Split Problem		Probability	p	q	r
	P1	Outcomes	a	0	0
	P2		0	b	b
Coalesced Problem		Probability	p	$1-p$	
	P3	Outcomes	a	0	
	P4		0	b	

Panel 3: Monotonicity Lotteries[b]

Lottery				
A	Probability	p	$1-p$	
	Outcomes	x	0	
B	Probability	q	$1-q$	
	Outcomes	y	0	
C	Probability	r	s	$1-r-s$
	Outcomes	y	$y-\varepsilon$	0

[a] $p+q+r=1$, $a>b>0$. Set 1 parameters are $p=r=0.3$, $b=£21$. Set 2 parameters are $p=0.5$, $q=0.3$, $r=0.2$, $b=£10$.
[b] $p<q$, $q=r+s$, $x>y>\varepsilon>0$. Set 1 parameters are $p=r=0.5$, $s=0.2$, $x=£18$, $y=£11$, $\varepsilon=£0.5$. Set 2 Parameters are $p=r=0.4$, $s=0.2$, $x=£24$, $y=£13$, $\varepsilon=£1$.

TABLE II
Common consequence effects*

Pairs 1 and 2	$v(S1) \geqslant v(R1)$ and $v(R2) \geqslant v(S2)$
Pairs 1 and 3	$v(S1) \geqslant v(R1)$ and $v(R3) \geqslant v(S3)$
Pairs 2 and 3	$v(S2) \geqslant v(R2)$ and $v(R3) \geqslant v(S3)$

*$v(.)$ represents the valuation assigned to a lottery. A violation of expected utility theory requires at least one strict inequality in each of the three rows.

of the type illustrated in the second panel of Table I. Over these problems a violation of expected utility theory is manifest in a preference for P2 over P1 and P3 over P4, despite the fact that each problem offers identical probabilities of identical outcomes in each respective lottery. Splitting the event which offers a $1-p$ chance of b in P4, such that it offers a q chance of b and a r chance of b in P2, renders P2 relatively more attractive than P4. This is a violation of the *coalescing principle*, which states that identical outcomes (b) will be combined by adding their probabilities ($q+r=1-p$) prior to choice.[2]

Event-splitting effects are unlike many choice anomalies because there exists a range of real world decision-making contexts where one observes analogous behaviour. For example, insurance policies often describe risks such that they are split into their most detailed 'sub-risks'. Insurance cover for having your wallet stolen from your house *and* having your wallet stolen whilst out of the house may appear more attractive than cover for simply having your wallet stolen. By enhancing the impression of cover, splitting risks may allow insurance companies to charge higher premiums (see Johnson et al. (1993) for related evidence). Alba and Marmorstein (1987) discuss the use of frequency information by marketers to shape consumer decisions through influencing the salience of object attributes (splitting events increases the frequency with which outcomes are represented in the decision problem). In related work, Weber et al. (1988) report *attribute-splitting effects* in multiattribute utility evaluation studies and

Bateman et al. (1997) present evidence of *part-whole bias* in contingent valuation studies. Common to the above examples is that the frequency with which attributes are represented in decision tasks can influence behaviour. Tversky and Koehler (1994) refer to this phenomenon as an *unpacking effect*, and allude to the possibility that it may operate at a fundamental level where no risk is present.

2.3. *Monotonicity violations: Preference for dominated lotteries*

Starmer (1999a) reports violations of monotonicity and transitivity over lotteries described in the third panel of Table I. When subjects were asked to choose between lottery A and lottery B, lotteries B and C, and lotteries A and C, Starmer (1999a) observes choice cycles involving A being chosen over B, B over C and C over A (hence abbreviated to ABC). Note that this (predicted) cycle, as opposed to the (counter-predicted) cycle (BCA), involves choosing B over C, which it dominates by a s chance of ε.

Predicted cycles and event-splitting effects appear related. Assuming that individuals recognise the dominance relation in the choice between B and C (94% of Starmer's (1999a) subjects do), they are driven by the fact that, in the choice between A and B, lottery A is relatively less attractive than it is in the choice between A and C. Lottery C offers two positive outcomes whereas B offers only one, despite the fact that the chances of winning positive amounts are identical in each and that B dominates C. To investigate this, Starmer (1999a) conducts a test which compares the frequency of what he calls pattern 1 choices (A chosen over B and C chosen over A) with pattern 2 choices (B chosen over A and A chosen over C). He observes a greater incidence of pattern 1 choices, which favour the predicted cycle, than pattern 2 choices, which favour the counter-predicted cycle. Therefore, although Starmer's (1999a) subjects obey monotonicity in the transparent choice between B and C, it is indirect violations of monotonicity generated by an event-splitting type argument which appear to be driving predicted choice cycles.

In the experiment reported here it is not possible for a set of three valuations assigned to lotteries A, B and C to indicate a non-transitive preference ordering. It is, however, possible to observe violations of monotonicity in valuation in a manner which would imply non-transitive choices of the type observed by Starmer (1999a), as well as violations of monotonicity which do not.[3] For example, if $v(.)$ denotes the money value attached to a lottery, then a violation of monotonicity would be manifest in $v(C) \geqslant v(B)$. If the preference ordering implied by these valuations was replicated in choices it need not imply cycles, because it may represent only part of a more general ordering such as $v(A) > v(C) \geqslant v(B)$ or $v(C) \geqslant v(B) > v(A)$. If the latter ordering, for example, was applied to choices, cycles would not be observed because A would be chosen over neither B nor C. If, however, it was the case that $v(C) > v(B)$ and this formed part of a more general ordering $v(C) > v(A) > v(B)$, then the implication is that A is preferred to B, and C is preferred to A, exactly as is the case for Starmer's (1999a) pattern 1 choices. By comparing the incidence of all valuation patterns involving $v(C) \geqslant v(B)$ with $v(C) \geqslant v(A) \geqslant v(B)$ (with at least one strict inequality), it is possible to derive an indication of the extent of monotonicity violations which might imply non-transitivities in choice, without allowing non-transitive preferences to emerge *per se*.

3. LEARNING IN RISKY CHOICE EXPERIMENTS

3.1. *The discovered preference hypothesis*

As Cubitt et al. (2001) point out, Plott's (1996) *discovered preference hypothesis* is supported by Smith (1989), Harrison (1994) and Binmore (1999), and holds that individuals have a unique and precisely structured set of underlying preferences, but in order for them to be elicited in decision-making tasks the individual will first have to discover which action best satisfies those preferences. To this end some kind of learning, possibly trial-and error, or deliberation, is required. Learning from experience transforms initially imprecise preferences into 'firmed-up' expected utility preferences. Experiments which report choice anomalies without appropriate opportunities for learning cannot therefore be taken as evidence against expected utility theory. The discovered preference hypothesis gives rise to the possibility that the plethora of experimentally observed choice anomalies are the product of inexperienced decision-makers making novel choices in an unfamiliar environment whilst learning how their imprecise preferences interact with the task. As individuals become more experienced and their preferences become less imprecise, one might expect choice anomalies to diminish or disappear altogether. If risky choice anomalies are transient, as the discovered preference hypothesis suggests, there are (at least) two important implications. First, the descriptive challenges which have undermined expected utility theory, broader models based upon expected utility theory, and the welfare and policy judgements which follow, may turn out to be misplaced in 'evolved' economic contexts. Second, the risky choice research agenda should prioritise modelling the learning process (Friedman, 1998) and investigating how different learning opportunities help to form preferences (Loomes, 1999).

There is evidence which supports the discovered preference hypothesis. For example, Cox and Grether (1996) observe the decay of preference reversals in a Vickrey second price auction. Loomes et al. (2002) find that as subjects progress through sequences of pairwise choice problems, their decisions converge upon expected utility maximisation. Friedman (1998) experimentally replicates Monty Hall's three door problem. In this task subjects are asked to select one of three doors. Behind one door is the 'grand' prize and behind the two others are virtually worthless prizes. When the subject has selected a door, one of the other doors is opened to reveal a worthless prize, and the subject is offered the chance to switch their selected door with the remaining unopened door. Switching is rarely observed despite the fact that this would increase the probability of winning the grand prize from 1/3 to 2/3. When subjects are offered the chance to learn by keeping a written record of the outcomes of repeated trials, taking written advice, and comparing their winnings with those of others, irrational behaviour is greatly diminished. Friedman (1998, p. 941) interprets his evidence as suggesting that appropriately structured learning environments render the existence of anomalous behaviour (in the sense of stable and yet irrational choices) unlikely.

3.2. *Frequency-based probability learning*

Estes (1976a) notes that individuals do not process probability information with advanced statistical ability, but rather rely on simple heuristic devices. Of particular importance is that learning by experience, in the shape of observing the outcomes of repeated situations, coupled with faith in the uniformity of nature, yields the fundamental heuristic that more frequently observed outcomes are more likely to be future outcomes. If repeated situations yield different outcomes then probability judgements are generated by converting absolute event frequencies into relative event frequencies. On Estes's (1976a,b) view, therefore, probability learning is based on the learning of absolute event frequency. This conception of probability learning is hence termed the *frequency-based probability learning hypothesis*.[4] To support his view Estes (1976a) provides experimental evidence from a predictive decision-making task which shows that if in observation trials an event occurs more frequently than is suggested by objective probability, individuals judge the future occurrence of that event to be more likely than it actually is. Further evidence of deviations from rationality stemming from the use of frequency heuristics is provided by Einhorn and hogarth (1987). Humphrey (1999) replicates Estes's (1976a) results in a computerised experiment with financial incentives for accurate probability learning.

Although the above evidence documents irrational behaviour attributable to frequency-based probability learning, Estes (1976a,b) also shows that if individuals observe sequences of outcomes where event occurrence matches true probabilities, probability learning is accurate and extremely efficient. It is this evidence with which the experiment reported here is primarily concerned. The learning opportunity provided to one group of subjects involves showing them the outcomes of 10 resolutions of the risk in the lottery, alongside stated probability information, prior to making their valuation. These resolutions of the risk will yield event occurrences exactly as suggested by the stated probabilities. Consider how this type of learning opportunity might promote expected utility maximisation.

All of the anomalous behaviours discussed in Section 2 can be explained by probability weighting models such as prospect theory (Kahneman and Tversky, 1979). Assume an individual is faced with a choice between prospects P_i, where $i = 1, 2$. Each prospect is represented by a vector of probabilities p_{ij} for $j = 1, \ldots, n$ where p_{ij} is the probability that P_i results in consequence x_j. Individuals maximise the value of:

$$\sum_{j=1}^{n} \pi(p_{ij}) v(x_j) \tag{1}$$

Where $\pi(p_{ij})$ is the probability weight attached to probability p_{ij} and $v(.)$ is the monotonically increasing utility function assigned to increments or decrements of wealth relative to the individual's current asset position. $v(.)$ is unique up to multiplication by a positive constant with $v(0) = 0$ at the point of reference.

If the decision rule in expression (1) is applied to the common consequence effect between pairs 1 and 2 in Table II, the violation of coalescing discussed in Section 2.2 (assuming that there is no editing stage) and the predicted cycles in Section 2.3 (assuming the dominated lottery in the choice between B and C is ignored) expressions (2), (3) and (4) are respectively yielded:

$$\pi(0.9) > \pi(0.4) + \pi(0.5) \tag{2}$$

$$\pi(p) + \pi(q) > \pi(p + q) \tag{3}$$

$$\pi(r)v(y) + \pi(s)v(y - \varepsilon) > \pi(q)v(y) \tag{4}$$

The probability weighting function is assumed to be well-behaved around the salient endpoint probabilities such that $\pi(1) = 1$ and $\pi(0) = 0$, but is less so elsewhere. A subadditive probability weighting function would render expression (3) true and be a necessary condition for expression (4). In the absence of a probability learning opportunity, probability weights are more likely to display non-linearity, such as subadditivity, which generates anomalous behaviours.

If, in the context of probability weighting models, frequency-based probability learning engenders expected utility maximising behaviour, it would do so by promoting a reduction in the non-linearity of the probability weighting function. By allowing reflection on probability weights formed on the basis of stated probabilities, frequency-based probability learning may achieve exactly this. Following Diecidue and Wakker (2001), consider the lottery (£30, 0.4; £20, 0.5; £10, 0.1). Cumulative prospect theory (Tversky and Kahneman, 1992) suggests that the importance of an outcome in the evaluation of the lottery depends not only on its probability, but also how good it is in relation to the other outcomes. If the decision-maker is a pessimist, for example, the probability weight attached to the worst outcome (£10) will be such that $\pi_{10}(0.1) > 0.1$, say 0.3. Similarly, being a pessimist, more than half of the remaining attention will be paid to the next worst outcome (£20) such that $\pi_{20}(0.5) > 0.5 = 0.6$. This leaves $\pi_{30}(0.4) = 0.1$. Diecidue and Wakker (2001) point out that rank-dependent probability weights of the type described may represent an irrational belief that relatively aversive events tend to happen more often than suggested by their true probability.

Now imagine that the decision-maker observes ten resolutions of the risk in the lottery, which yields £10 only once, £20 five times and £30 four times (i.e. event occurrence exactly matching stated probabilities). It is plausible to suggest that experience of observing these outcomes causes the pessimist to regard the probability weight $\pi_{10}(0.1) = 0.3$ as placing too much emphasis on an unlikely outcome. Correspondingly, even the pessimist, upon observing £30 to have occurred four times out of ten, might regard the probability weight $\pi_{30}(0.4) = 0.1$ as under-representing the importance of the best outcome in the evaluation of the lottery. Thus it seems that the opportunity to engage in frequency-based probability learning might cause sufficient reflection upon probability weights (formed on the basis of stated information) to engender their modification and purge the weighting function of the properties which imply anomalous behaviour. All that is required to form accurate probability assessments is the observation of outcomes over a sufficiently long period such that relative frequencies of event occurrence converge on objective event likelihood. If so, any imprecision in preferences rooted in a lack of experience in dealing with and/or understanding the meaning of probabilities may be diminished.[5]

4. EXPERIMENTAL DESIGN

4.1. *Valuation tasks, learning and incentives*

The experiment involved subjects assigning money valuations to a total of twenty lotteries, fifteen of which were concerned with problems described in Table I. Each lottery was expressed in terms of tickets numbered consecutively from 1 to 10 as illustrated in Figure 1. The valuation was attached to each lottery through what was termed the 'yardstick'. Subjects were told that they should increment or decrement the amount shown in the boxes containing question marks (the 'up' cursor key changed the question marks to £00.00, pressing it again gave £00.10, etc.) until the box displayed an amount such that should they be offered the lottery or the yardstick they would not mind which they received.[6]

Subjects were assigned to one of two conditions. 'Control' treatment subjects simply valued each lottery before moving on to the next. 'Learning' treatment subjects saw the lottery, but before valuing it saw an observation sequence involving ten resolutions of the risk in the lottery. This is captured in the box in Figure 1 showing 'draw' (1 to 10) and 'winnings' (the outcome of each of the draws). When subjects started the observation sequence the computer showed the outcome of the first draw (£21 in Figure 1), paused for a second, showed the outcome of the second draw, and so on. The lottery draws in the observation sequence were such that events occurred exactly according to their stated probabilities in a manner analogous to a 'speeding-up' of the law of averages.[7] Subjects were told that in a genuinely random draw according to the stated probabilities, any of the outcomes in the lottery could occur, but that the observation sequence shows them what might happen over a series of ten resolutions of the risk. Investigating whether violations of independence, coalescing and monotonicity replicate in valuation tasks, as opposed to pairwise choice tasks, is an interesting test of their robustness. However, the primary motivation for employing a valuation task is that it provides a fair and simple test of the frequency-based probability learning hypothesis. By allowing subjects to concentrate on one lottery at a time, a valuation task precludes subjects in the learning treatment having to evaluate two lotteries at a time alongside assimilating an observation sequence for each, as would be the case in a pairwise choice task.

The incentive system used in the experiment made it in subjects' financial interests to attach valuations to the lotteries which reflected their genuine preference *ordering* over the set of lotteries. Subjects were told that at the end of the experiment two lotteries would be randomly selected by drawing two numbered disks from a bag containing twenty consecutively numbered disks. The valuations assigned to these lotteries would be compared and the lottery which was valued highest would be played out for real money (by drawing a disk from a bag containing ten consecutively numbered disks) to determine their payment for participation in the experiment.[8] Since subjects would not know which two lottery valuations would be compared to determine the payment lottery

Question One

LOTTERY:	Lottery Tickets	1 to 4	pay you	£21
	Lottery Tickets	5 to 7	pay you	£21
	Lottery Tickets	8 to 10	pay you	zero

YARDSTICK:	Lottery Tickets	1 to 10	pay you	??.??

Draw:	1	2	3	4	5	6	7	8	9	10
Winnings:	21	0	21	21	21	0	0	21	21	21

Enter a value for ??.?? which would make the LOTTERY and the YARDSTICK equally attractive to you.
The up and down cursor keys select a value. Press <ENTER>to confirm.

Figure 1. Valuation task display.

until after all lotteries had been valued, they could only be sure of playing out their truly preferred lottery from the pair by assigning valuations which reflected their genuine preference ordering over the set of twenty lotteries. Subjects were told that one way in which they could guarantee playing out their truly preferred lottery from the randomly selected pair would be to consider each lottery carefully and assign it a genuine valuation.[9]

The system of incentives described above is a hybrid of the random lottery incentive system and what Tversky et al. (1990) call the *ordinal payoff scheme*. The ordinal payoff scheme is favoured over the Becker et al.(1964) incentive mechanism (BDM) in valuation tasks because it has been shown that the latter can only be theoretically relied upon to elicit true valuations if the independence axiom of expected utility theory holds (e.g. Segal, 1988).[10] Camerer (1995) discusses a wealth of evidence which shows that it does not. Avoiding the BDM device mitigates these concerns. Holt (1986), however, shows that if subjects treat experiments as a single large choice between compound lotteries which have first been simplified according to the reduction principle, then it cannot be inferred that behaviour observed in a comparison of two lotteries can be taken as indicating genuine preference over those lotteries. For example, take the tests of the coalescing principle. Holt's (1986) argument suggests that observing a higher valuation of P2 than of P4 cannot be taken as a genuine preference for P2 over P4. According to Holt (1986), valuing P2 over P4 is equivalent to selecting lottery H_1:

$$H_1 : [P2, \lambda; Z, 1 - \lambda]$$

Where λ is the probability that P2 and P4 are selected as the two lotteries whose valuations are compared to determine the lottery in which the risk will be resolved to determine the subject's payment, and Z is determined by behaviour over all other tasks in the experiment. Similarly, valuing P4 greater than P2 is equivalent to selecting lottery H_2:

$$H_2 : [P4, \lambda; Z, 1 - \lambda]$$

Holt's (1986) argument is that P2 being valued above P4 cannot be taken as a genuine preference for P2 because there is a potentially perturbing wedge driven between observation and inferred preference by the term in Z. If the independence axiom of expected utility does not hold then this term may render revealed preference in an experiment involving additional tasks different to that which would be expressed, for example, in a single direct choice between P2 and P4. Relative valuations may have been contaminated by the other tasks. If, however, a preponderance of valuations patterns which indicate P2 to be preferred to P4 is observed, Holt's (1986) argument would need to explain why violations of independence systematically favour this valuation pattern rather than the opposite one. Since, from the perspective of conventional theory, P2 and P4 are identical, any such explanation would appear to require an account of event-splitting effects. These tests, therefore, control for Holt's (1986) hypothesis. A similar argument can be extended to the other anomalies discussed in Section 2.

Holt's (1986) argument can neither explain any differences which may be observed in behaviour between the control and learning treatments. Any such difference could not be explained by any generalized expected utility theory which does not involve independence and views preferences between lotteries as dependent only on probabilities and outcomes, even if independence does not hold and relative valuations are contaminated. Any theory of this type would view the incentive system faced by our control and learning groups as equivalent. And so this could not be responsible for any observed differences in behaviour.

4.2. *Hypotheses*

Violations of independence, coalescing and monotonicity in valuation tasks are manifest in particular patterns of valuations. Table III describes the hypothesis tests which will be conducted to establish whether violations of independence, coalescing and monotonicity emerge. In the tests for common consequence effects and event-splitting effects the null hypothesis is that patterns of valuations which violate expected utility theory in the direction consistent with the predicted anomaly (common consequence effects as described in Table II or an event-splitting effect) are no more frequent than the opposite patterns of valuations which violate expected utility theory.[11] Implicit in this null hypothesis is the neutral assumption that patterns of valuations which deviate from expected utility maximisation are due to random errors. Random errors, of course, may be indicative of imprecision in preferences. In the tests for monotonicity violations the null hypothesis is that patterns of valuations which violate dominance are no more frequent than monotonic valuation patterns. A failure to reject the null hypothesis would mean that it cannot be ruled out that patterns of anomalous valuations are the result of random mistakes. In the case of monotonicity violations which might imply non-transitivities in choice, the null hypothesis is that patterns of valuations consistent with the choice cycles observed by Starmer (1999a) are no more frequent than valuations which imply a choice cycle in the opposite direction. The alternative hypotheses are that patterns of valuations consistent with the common consequence effects described in Table II, or which violate coalescing and monotonicity, are observed with greater frequency than the opposite patterns of valuations.

Once it has been established whether violations of independence, coalescing or monotonicity are respectively observed within the control and learning groups, attention will turned towards investigating any differences in behaviour between treatments. This investigation will be two-fold. Firstly, the null hypothesis that learning does not occur and violations of expected utility theory occur with similar frequency under both conditions will be pitted against a one-tailed alternative hypothesis provided by the discovered preference and the frequency-based probability learning hypotheses; that behaviour observed in the learning group is more consistent with expected utility maximization than behaviour in the control group. Secondly, the data will be interrogated to reveal whether frequency-based probability learning affects the distribution of any observed violations of expected utility theory. For example, if significant violations of coalescing observed in the control group are not observed in the learning group, this could be due to learning generating fewer event-splitting effects, more violations of expected utility theory in the opposite direction, or a combination of both.

4.3. *Conduct of the experiment*

The experiment was conducted at the Centre for Decision Research and Experimental Economics (CeDEx) laboratory at the University of Nottingham. A total of 203 subjects were recruited by e-mail shot to the CeDEx mailbase of pre-registered volunteers from across the undergraduate population and asked to reserve a place in one of a number of pre-arranged sessions. It was determined randomly in advance whether each session would be a control treatment or a learning treatment, with 67 subjects taking part in the former and 136 in the latter. 60% of subjects were male. Each session

TABLE III

Hypotheses[a]

Test	Null Hypothesis	Alternative Hypothesis
Common Consequence Effect[b]	$n[v(S1) \geqslant v(R1)$ and $v(R2) \geqslant V(S2)] = n[v(R1) \geqslant v(S1)$ and $v(S2) \geqslant V(R2)]$	$n[v(S1) \geqslant v(R1)$ and $v(R2) \geqslant V(S2)] > n[v(R1) \geqslant v(S1)$ and $v(S2) \geqslant V(R2)]$
	$n[v(S1) \geqslant v(R1)$ and $v(R3) \geqslant V(S3)] = n[v(R1) \geqslant v(S1)$ and $v(S3) \geqslant V(R3)]$	$n[v(S1) \geqslant v(R1)$ and $v(R3) \geqslant V(S3)] > n[v(R1) \geqslant v(S1)$ and $v(S3) \geqslant V(R3)]$
	$n[v(S2) \geqslant v(R2)$ and $v(R3) \geqslant V(S3)] = n[v(R2) \geqslant v(S2)$ and $v(S3) \geqslant V(R3)]$	$n[v(S2) \geqslant v(R2)$ and $v(R3) \geqslant V(S3)] > n[v(R2) \geqslant v(S2)$ and $v(S3) \geqslant V(R3)]$
Coalescing	$n[v(P2) > v(P4)] = n[v(P4) > v(P2)]$	$n[v(P2) > v(P4)] > n[v(P4) > v(P2)]$
Monotonicity	$n[v(C) \geqslant v(B)] = n[v(B) > v(C)]$	$n[v(C) \geqslant v(B)] > n[v(B) > v(C)]$
Monotonicity/ Transitivity[c]	$n[v(C) \geqslant v(A) \geqslant v(B)] = n[v(B) \geqslant v(A) \geqslant v(C)]$	$n[v(C) \geqslant v(A) \geqslant v(B)] > n[v(B) \geqslant v(A) \geqslant v(C)]$

[a]n$[v(P2) > v(P4)]$, for example, should be interpreted as the number of valuation patterns which involve P2 lotteries (as in Table I) being assigned a higher value that P4 lotteries.

[b]The set of three null and alternative hypotheses respectively refer to the common consequence effects tests which compare lottery pairs 1 and 2, pairs 1 and 3, and pairs 2 and 3, as in Table II.

[c]Each term in square brackets requires at least one strict inequality. The term on the left hand side of the equality under the null hypothesis represents the number of valuation patterns which violate monotonicity in a manner consistent with Starmer's (1999a) pattern 1 choices. The term on the right-hand side of the equality represents valuation patterns consistent with Starmer's (1999a) pattern 2 choices.

lasted for approximately one hour, including detailed instructions from the experiment organiser and working through two practice valuations to illustrate the task and the incentive mechanism. Average payment for participation was £12.97. The order in which the twenty lotteries were valued was randomised for each subject to control for order effects. For learning treatment subjects, the order in which each outcome occurred during the observation sequence was also determined randomly. There was no time limit for completion of the tasks.

5. RESULTS

5.1. *Independence violations: The common consequence effect*

Table IV reports the results of the test for common consequence effects. Taking the control treatment data first, the EUT column shows 43–46% (with an average of 45%) of valuation patterns to be consistent with expected utility maximisation. These subjects always valued the riskier lottery higher, always valued the safer lottery higher, or expressed indifference by always valuing the riskier and safer lotteries identically. This means that over half (54–57%) of the patterns of valuations assigned by control group subjects violate expected utility theory. Despite this overall violation rate, the data in the predicted and counter-predicted anomaly columns show violations to be broadly equally distributed. The *p-value* column shows that the null hypotheses in Table III cannot be rejected in favour of the alternative hypothesis of systematic common consequence effects. The most sustainable explanation of violations of expected utility theory in the common consequence effect tests is therefore random noise. The source of this noise may be imprecision in preferences.

The learning treatment data show patterns of valuations consistent with expected utility theory to vary between 45% and 61%. The average consistency rate is 53%. This is slightly

TABLE IV
Results[a]

Patterns of valuations

	Control treatment (valuation only)				Learning treatment (observation then valuation)				Z
	EUT	Anomaly		$p-value$	EUT	Anomaly		$p-value$	
		Predicted	Counter			Predicted	Counter		
Common Consequence Effect									
Pairs 1, 2	31/67 (46%)	17/67 (25%)	19/67 (28%)	0.9953	41/67 (61%)	12/67 (18%)	14/67 (21%)	0.4225	−1.7523††
Pairs 1, 3	29/67 (43%)	22/67 (33%)	16/67 (24%)	0.2088	30/67 (45%)	27/67 (40%)	10/67 (15%)	0.0038**	−0.1740
Pairs 2, 3	31/67 (46%)	20/67 (30%)	16/67 (24%)	0.7975	36/67 (54%)	21/67 (31%)	10/67 (15%)	0.0354**	−0.8663
Coalescing									
Set 1	15/67 (22%)	24/67 (36%)	28/67 (42%)	0.3389	20/67 (30%)	29/67 (43%)	18/67 (27%)	0.0719*	−0.9868
Set 2	32/67 (48%)	22/67 (33%)	13/67 (19%)	0.0877*	23/69 (33%)	27/69 (39%)	19/69 (28%)	0.1510	1.7313
Monotonicity									
Set 1	25/67 (37%)	42/67 (63%)	–	0.0249**	34/67 (51%)	33/67 (49%)	–	0.5964	−1.5807†
Set 2	26/67 (39%)	41/67 (61%)	–	0.0432**	37/69 (54%)	32/69 (46%)	–	0.7648	−1.7525††

Montonicity /Transitivity[b]									
Set 1	56/67 (84%)	7/67 (10%)	4/67 (6%)	0.2744	51/67 (76%)	13/67 (19%)	3/67 (4%)	0.0106**	1.0815
Set 2	51/67 (76%)	10/67 (15%)	6/67 (9%)	0.2272	58/69 (84%)	4/69 (6%)	7/69 (10%)	0.8867	−1.1635

[a]The ‘EUT’ column denotes valuation patterns corresponding to consistent preferences as described by expected utility maximisation. For the common consequence effect tests this column shows valuations which indicate consistent within-subject preferences encapsulated in $V(S_i) > V(R_i)$ and $V(S_j) > V(R_j)$, or $V(R_i) > V(S_i)$ and $V(R_j) > V(S_j)$, or $V(S_i) = V(R_i)$ and $V(S_j) = V(R_j)$ for $i, j = 1, 2, 3$ and $i \neq j$. The ‘anomaly’ column shows valuation patterns which are inconsistent with expected utility preferences. The ‘predicted’ anomalies are the common consequence effects described in Table II, event-splitting effects, violations of monotonicity, and violations of monotonicity which might imply non-transitive choices, for each of the four tests respectively. For the monotonicity tests, there is no counter-predicted anomaly as monotonic preferences imply expected utility maximisation. The *p-value* column reports a test statistic based on the binomial distribution that predicted anomalies occur with a greater frequency than counter-predicted anomalies or, in the monotonicity tests, the selection of the dominant lottery. An asterisk denotes 10% significance. Two asterisks denote 5% significance. Percentages may not sum to 100% due to rounding. The Z column shows the results of test of difference in sample proportions based on the normal distribution of the null hypothesis that expected utility maximizing behaviour is the same under the control and learning treatments. The alternative hypothesis is that learning increases expected utility maximization. Support for the alternative hypothesis requires a negative Z-value. ‘†’ and ‘††’ respectively indicate rejection of the null hypothesis at the 10% and 5% levels.

[b]In the tests for monotonicity/transitivity violations some patterns of valuations are placed in the in the ‘EUT’ despite violating monotonicity. These patterns include, for example, $v(C) \geqslant v(B) \geqslant v(A)$, which values C at least a high as B, but would not imply choices cycles because B and C would both be at least as preferred as A.

higher than that of 45% under the control group. The Z column in Table IV contains the results of a test of difference in sample proportions based on the normal distribution. Three negative Z-values show greater consistency with expected utility theory amongst the learning group in all three common consequence effect tests. This difference is significant at the 5% level in the pairs 1 and 2 comparison ($Z = -1.7523$). There is weak evidence therefore that the probability learning opportunity mitigates violations of the independence axiom of expected utility theory.

Violations of expected utility theory in the learning group are distributed differently to those in the control group. Whereas the latter indicate randomly noisy violations, the former reveal significant common consequence effects at the 5% level in both the pairs 1 and 3 and pairs 2 and 3 comparisons. The implication of these data is that whilst frequency-based probability learning increases consistency with expected utility theory, it also distils noisy violations of expected utility theory such that they become systematic. It seems that overall convergence on expected utility maximisation does not rule-out a proportion of systematically anomalous behaviour.

The data from the common consequence effect tests are consistent with Friedman's (1998, p. 942) recognition that some learning environments may actually encourage choice anomalies. They are also consistent with Slovic and Tversky's (1974) findings in a common consequence effect test which used pairwise choices and involved a different type of learning opportunity. After subjects had made their choices (some 60% of which violated independence), they were presented with arguments in favour of and against the logic of the independence axiom. They were then offered the chance to change their initial decisions. Slovic and Tversky (1974) observe that rather than mitigating the choice anomaly, this learning opportunity increased violation of the independence axiom. There appears, therefore, to be a growing body of support for Loewenstein's (1999) view that although there is a role to be played by learning in reducing confusion which may stem from experimental procedures, it should not be presumed that confusion-free preferences will be those dictated by expected utility theory. *Prima facie*, this argument appears to be as reasonable as the discovered preference hypothesis. If choice anomalies are genuine features of underlying non-expected utility preferences, then learning opportunities which give rise to a better identification of those preferences are likely, if anything, to cause the emergence or exacerbation of anomalous behaviour.

The question that this interpretation poses in the context of the experiment reported here is exactly why frequency-based probability learning appears to facilitate systematic violations of expected utility theory when no such systematic violations were previously in evidence? One answer to this question is provided by the preference reversal literature (e.g. Grether and Plott, 1979). Preference reversals are observed when a $-bet (offering a high money prize with a low probability) is assigned a higher reservation price than a *P*-bet (offering a relatively low money prize with a high probability), but is subsequently not chosen in a direct choice between the two. This pattern of behaviour is often attributed to *response mode* effects. One feature of response mode effects is *compatibility*. The compatibility hypothesis states that money is the salient attribute of lotteries in money valuation tasks (the two are compatible). This renders the high prize in the $-bet particularly influential in driving the valuation. A higher money valuation for the $-bet than for the *P*-bet is the result. In the choice task there is no such compatibility with money outcomes (and possibly one operating in favour of the *P*-bet because of the potentially enhanced salience of the probability of winning). So preferences are reversed in favour of the *P*-bet. Assume common consequence effects are the product of how probabilistic biases influence the probability weighting function, as in expression (2). The salience of the money attribute of the lotteries in the control group valuation tasks may have precluded the emergence of any such probability-driven anomalies. In the learning group, however, the observation sequence may have enhanced the salience of the probability attribute such that probability-based common consequence effects emerged.[12]

5.2. *Coalescing violations: event-splitting effects*

Considering the control treatment data first, the tests of coalescing reveal 22% of valuation patterns to be consistent with expected utility maximisation in parameter set 1 and 48% in set 2. The valuations patterns which violate coalescing are distributed differently under each set. The set 1 violations are approximately evenly distributed such that the null hypothesis of consistent valuations cannot be rejected. Under set 2, however, violations of coalescing are distributed in the direction of event-splitting effects such that the null hypothesis can be rejected at the 10% level. This, to my knowledge, is the first evidence of event-splitting effects in a valuation task. It is interesting that the systematic violation of the coalescing principle should emerge within the parameter set which also indicates the highest proportion of expected utility maximising behaviour. The difference in behaviour between the two sets of parameters is that the counter-predicted violation of coalescing is more prevalent in set 1 than in set 2. The set 2 parameters split a 0.5 chance (of £10) into 0.3 and 0.2 chances, and the set 1 parameters split a 0.7 (of £21) into 0.4 and 0.3 chances. Since the former split probabilities are smaller than the latter, this observation is consistent with subadditive probability weights, as in expression (3), stemming from the overweighting of small probabilities.

Turning to the learning treatment data, parameter sets 1 and 2 respectively reveal 30% and 33% of valuation patterns to be consistent with expected utility maximisation. The patterns which violate coalescing are split in the direction consistent with event-splitting effects in both parameter sets. In set 1 event-splitting effects are sufficiently frequent in relation to the opposite violation to reject the null hypothesis at the 10% level. This is not the case in set 2.

In parameter set 1, the impact of the learning opportunity on violations of coalescing between treatments is that expected utility maximisation increases from 22% to 30%. The Z-value in Table IV reveals this difference to be insignificant ($Z = -0.9868$). In a similar manner to that under the common consequence effect tests, however, frequency-based probability learning introduces a systematic violation of the coalescing principle. This is the case despite the fact that event-splitting effects do not significantly increase between treatments (36% and 43% of valuation patterns in the control and learning treatments, respectively, giving $Z = -0.88$). Increasingly systematic violations of coalescing appear to be driven by the reduction in counter-predicted violations. These fall from 42% in the control treatment to 27% in the learning treatment. This yields $Z = 1.82$ and allows rejection of the null hypothesis that learning does not affect behaviour with 5% significance in a one-tailed test. One interpretation of this observation is that counter-predicted violations of the coalescing principle stem from an inherent imprecision in preferences which frequency-based probability learning diminishes. Thus whilst probability learning significantly affects neither the proportion of expected utility maximization nor violations of coalescing in the direction of event-splitting effects *per se*, it does help to weed-out certain patterns of behaviour. On this interpretation, valuation patterns consistent with event-splitting effects should not be considered as transient features of imprecise preferences. If they were, it would be necessary to explain why subjects learned to commit fewer violations of the counter-predicted violation, but were not similarly inclined to commit fewer event-splitting effects.

The impact of learning on violations of coalescing in parameter set 2 between the control and learning treatments appears contrary to that in set 1. In parameter set 2 learning reduces patterns of valuations consistent with expected utility theory from 48% to 33% ($Z = 1.7313$). This may lead to the *prima facie* conclusion that violations of expected utility theory are features of genuine non-expected utility preferences which learning opportunities exacerbate by reducing initial imprecision in those preferences. The structure of these genuine preferences, however, would apparently not include subadditivity as stated in expression (3), because learning in set 2 also causes systematic event-splitting effects to diminish. The source of which appears to be a disproportionately large increase in counter-predicted violations of coalescing in relation to event-splitting effects. The latter increases from 33% to 39% and the former from 19% to 28% (neither are significant, respectively yielding $Z = -0.76$ and $Z = -1.12$). This is contrary to the set 1 observation.

The differential impact of learning on behaviour in each parameter set is not easy to reconcile. What appears to be happening is that violations of coalescing in set 2 are (insignificantly) greater in the learning treatment than in the control treatment (by similar order to that in parameter set 1), but a proportion of expected utility maximising behaviour (48% of valuation patterns falling to 33%) is being replaced with counter-predicted violations (19% of valuation patterns increasing to 28%). It may be that the set 2 lotteries are easier to value than set 1 lotteries because they involve what might be considered more salient parameters (a 0.5 chance of £10 as opposed to a 0.7 chance of £21 in set 1). But why this, in conjunction with learning, would shift expected utility maximizing valuations towards counter-predicted violations of expected utility theory is not clear. One explanation would be that subjects in both treatments were pretty sure of their preferences over the set 2 lotteries, and therefore of the appropriate valuations to assign, but the inclusion of the observation sequence *per se* (rather than what was learned from the observation sequence) disrupted this sureness such that expected utility maximising behaviour diminished. Any confidence in this explanation, however, would require an auxiliary account of why such disruption should favour counter-predicted violations over event-splitting effects.

The fact that systematic violations of coalescing are not observed in the learning treatment in parameter set 2, does not critically undermine the conclusions derived from parameter set 1. Violations of expected utility theory do not seem to be transient. The conclusion might be different if the lack of systematic coalescing violations in the set 2 learning treatment was observed alongside an increase in expected utility maximization. But this is not the case. Deviations from expected utility maximisation in set 2 are more prevalent in the learning treatment than they are in the control treatment. This observation is consistent with the conclusion that learning does not increase expected utility maximization, despite the reasons why it might be expected to, and may exacerbate features of preferences which generate anomalies.

5.3. *Monotonicity violations: Preference for dominated lotteries*

The results of the tests for violations of monotonicity appear more clear cut than those for violations of coalescing. In the control condition, in both parameter sets, valuation patterns indicate significant violations of monotonicity at the 5% level. These non-monotonic valuations may have their roots in an event-splitting type argument resting on the subadditivity of decision weights as in expression (3). The data support Starmer's (1999b) suggestion that despite the status of monotonicity as a fundamental canon of rational choice, when the dominance relation is not transparent, significant violations emerge. In this respect the extent of non-monotonic behaviour may be better reflected by the data at hand than it is by pairwise choice experiments which demand consideration of dominant and dominated lotteries simultaneously. Starmer's (1999b, p. F9) comment that individuals generally do not have monotonic preferences, however, must be judged in light of the learning treatment data.

The learning treatment data reveal that frequency-based probability learning appears to facilitate a better identification of monotonicity in preferences such that systematic violations no longer emerge under either parameter set. In parameter set 1, 63% of valuation patterns violate monotonicity in the control treatment, and this falls to 49% in the learning treatment. Similarly, in parameter set 2, 61% violation falls to 46%. As Table IV shows, on the basis of a one-tailed test of difference in sample proportions, the former reduction yields $Z = -1.5807$ and the latter yields $Z = -1.7525$.These respectively allow rejection of the null hypothesis that learning does not increase expected utility maximization at the 10% and 5% levels.

When set alongside the results of the tests for violations of coalescing, the tests for non-monotonic preferences suggest that particular types of learning opportunity successfully diminish some violations of expected utility theory but not others. This might be surprising since event-splitting effects and monotonicity violations appear to be related in that they are both consistent with a subadditive probability weighting function as described in expressions (3) and (4). However, the subadditivity of probability weights may be embedded in preferences such that learning will only modify those preferences if the vehicle through which learning is applied is a task which has properties that act as a catalyst for behavioural modifications. The fact that one of the small probability events is dominated in the tests for monotonicity, but is not in the tests for event-splitting effects, may provide the required catalyst.

5.4. *Violations of monotonicity which imply non-transitive choices*

Finally, Table IV reports the extent of monotonicity violations which might imply cyclical choices of the type reported by Starmer (1999a) in pairwise choice tasks. In both parameter sets the control treatment data reveal no significant differences between the predicted and counter-predicted violations of expected utility theory. In the learning treatment, parameter set 2 reveals that total violations of expected utility theory fall from 24% to 16%. This reduction is primarily driven by the fall in the incidence of predicted violations from 15% to 6% ($Z = 1.75$). However, since significant predicted violation patterns are not observed under either treatment, this should not be taken as strong evidence that probability learning mitigates this particular violation of expected utility theory.

The parameter set 1 data show that an insignificant incidence of predicted violations in the control treatment increases (insignificantly, $Z = -1.45$) alongside a marginal reduction in counter-predicted patterns to yield a significant effect in the learning treatment ($p = 0.0106$ in Table IV). Frequency-based probability learning appears to engender an identification of preferences which violate monotonicity in a manner which might be considered analogous to that which generated the cyclical choices observed by Starmer (1999a).

6. CONCLUSION

The experiment reported here tests whether frequency-based probability learning causes violations of expected utility theory to diminish or be exacerbated. The motivation for conducting this test is the possibility that choice anomalies stem, at least in part, from an inherent imprecision in risk preferences attributable to a lack of understanding of the meaning of stated probability information. The answer to this question has implications for both the status of rational choice theory and for practitioners involved in the elicitation of risk preferences for the purpose of informing public policy.

The overall results of the experiment are mixed. For example, the tests for coalescing and monotonicity/transitivity violations show significant violations of expected utility theory to diminish with learning under one parameter set, but to emerge where none previously existed under the other parameter set. However, the data also contain clear patterns. The common consequence effect test data for pairs 1 and 3 and pairs 2 and 3, the coalescing test data for parameter set 1 and the monotonicity/transitivity test data for parameter set 1 all reveal learning to engender the emergence of systematic violations of expected utility theory. This occurs alongside the overall incidence of expected utility maximization being insignificantly different between treatments. If, as these data show, enhancing the understanding of probability information through learning causes a known anomaly to emerge, then an interesting question is whether there are as yet undiscovered anomalous features of preferences which this or other learning opportunities might bring to light in other risky choice tasks?

In his investigation of anomalous behaviour in Monty Hall's three doors problem, Dan Friedman (1998, p. 941) asserts that, "Every choice 'anomaly' can be greatly diminished or entirely eliminated in appropriately structured learning environments." The present data suggest this assertion to be only partly sustainable. The assertion would be sustainable if showing individuals a series of lottery draws prior to the elicitation of their preferences does not constitute an appropriately structured learning environment. There may be grounds upon which to suspect this to be the case. For example, frequency-based probability learning does not involve market discipline to punish ineffective learners. Nor does it allow the opportunity to imitate more successful decision-makers. But does this render it inappropriate? There are reasons to suggest not.

First, Estes (1976a, b) and others have shown frequency-based probability learning to be effective in both introducing probabilistic biases *and* engendering accurate probability learning in other tasks. Second, the beneficial information content of frequency-based probability learning opportunities enjoys anecdotal support from real world observations. The time-series of stock performances is often observed prior to periodic portfolio decisions. Betting form guides often provide information on the outcomes of a team's last *n* fixtures (and often not, for example, on whom the opponents were, the location of the game, the weather, injured players, and a variety of other potentially decision-relevant information).

Friedman (1998, p. 42) does not prescribe an ignorance of anomalies because they will eventually disappear. He does, however, argue the lack of need to modify, criticise or reject expected utility theory on the basis of anomalies stemming from incomplete learning. How the present experiment bears on this conclusion depends on how one defines incomplete learning. It would seem tautological to defend expected utility theory on the grounds of a learning argument where the definition of complete learning is when all choices conform to expected utility theory. Moreover, economic decisions are often made where there is no opportunity to observe and imitate more successful decision-makers, or where market forces are not strong enough to discipline behaviour. In this respect it is important to investigate the full range of economically-relevant learning opportunities. It would be dubious practice to concentrate research effort solely on investigating learning opportunities which might a priori be suspected of yielding the best chance of convergence on expected utility maximisation.

However, it is of course important to recognise that a different type of learning opportunity, or process to facilitate the discovery of genuine preferences might cause all of the violations of expected utility theory considered here to disappear. This is a matter for future investigation. But other evidence (e.g. Chu and Chu, 1990; Cox and Grether, 1996) and argument (e.g. Smith, 1989; Plott, 1996) suggest that market forces might successfully facilitate the discovery of expected utility preferences. Yet it should be pointed out that the basis for broader market theory is only one of the roles played by the theory of risky choice. For example, it is also used to structure the process by which values elicited in survey studies are fed into public policy decisions. In survey studies market forces are not present to discipline behaviour, and so policy decisions may be made on the basis of assumed expected utility preferences and values elicited from members of the public whose preferences may not be appropriately described by expected utility theory. The potential for wrongly informed policy decisions is obvious. Moreover, expected utility theory does not restrict its domain of applicability (although some authors such as Binmore (1999) have argued the case for this). It therefore seems unreasonable to attempt to validate the theory solely on the basis of observations that it seems to be approximately right in only one of the domains in which it is applied. In this respect the experiment reported here contributes to a richer picture of how different learning opportunities affect the formation of preferences in different, but equally economically relevant, domains.

Aside from contributing to the emerging debate between the discovered preference school and what Binmore (1999, F17) refers to as "the school of Kahneman and Tversky", the results of this experiment show that the impact of learning on similar choice anomalies can be quite subtle. The prevailing explanation of both the monotonicity violation considered here and event-splitting effects is the subadditivity of probability weights. Yet the former tests show learning to converge on expected utility theory whilst the latter tests show that increased expected utility maximization does not preclude the emergence of a proportion of systematic anomalies. Thus, if Friedman's (1998, p.942) call is answered and the "sterile" debate on anomalies is replaced by the modelling and testing of learning processes, it seems that this will not be an easy route to pursue. The transience or persistence of similar anomalies may not be robust to the same learning opportunity. By contributing some evidence contrary to the discovered preference perspective, the experiment reported here will hopefully assist in stimulating this potentially fruitful new research agenda.

ACKNOWLEDGEMENT

I am grateful for financial assistance from The Nuffield Foundation (award no. SGS/LB/0295), to David Lawson for running the experiments and to Chris Starmer for valuable comments.

NOTES

1. Henceforth, 'genuine' or 'true' preferences are taken as those which are free from imprecision attributable to a lack of experience or understanding of the decision-making task. Possible sources of imprecision are not understanding the meaning of stated information, confusion stemming from experimental procedures, and a lack of learning opportunity.
2. The coalescing violation is coupled with a transitivity violation as the effect over P2 and P4 is mediated through respectively pairing them with P1 and P3. Note that P1 and P3 are not strictly identical since, in relation to the latter lottery, the former splits the event offering zero. In most decision theories zero consequences are carri-

ers of zero utility and do not affect the decision. Humphrey (2001) reports a test of whether an aversion to greater frequencies of zero outcomes contributes to event-splitting effects and concludes that it does not.

3. This assumes that, contrary to the preference reversal phenomenon (e.g. Grether and Plott, 1979), there is no divergence between preference orderings elicited using choices and valuations.
4. Humphrey and Bleaney (2006) report evidence that frequency-based probability information is less ambiguous than stated likelihoods, and that when probabilities are represented in this way, individuals' valuations of risky prospects increase.
5. Frequency-based probability learning has been discussed in terms of how it may contribute to the (post-learning) accuracy of expected utility theory in describing preferences. It is, however, necessary to recognise that (ex post learning opportunity) preferences may be appropriately represented by some non-expected utility theory. In terms of the previous example, the pessimistic decision-maker may, on the basis of observing the least favourable £10 outcome occur once in ten observation trials, view $\pi_{10}(1/10) = 3/10$ as under-representing the importance of this outcome in their evaluation of the lottery. If so, non-linearity in the probability weighting function may increase. Associated deviations from expected utility maximisation may accordingly persist or increase.
6. Negative valuations were not allowed and no upper bound was imposed.
7. This controls for the possible bias introduced into lottery valuations by genuinely random sequences of observations which turn out to be non-representative. This question has been examined elsewhere (e.g. Humphrey, 1999).
8. In the event that these two lotteries were valued identically, the payment lottery was determined by flipping a coin. To clarify these incentives subjects worked through two example valuations. The valuations assigned to the two example lotteries were used to illustrate the incentive mechanism.
9. This incentive system controls for behaviour which, in other experiments, might be taken as evidence of irrationality, such as assigning valuations above the highest outcome offered by the lottery. Within this design it is legitimate to perform any monotonic transformation on valuations. Subjects can effectively eliminate one lottery from the set which could determine their payment by assigning it a zero valuation. This would only make sense, however, if it were their genuinely least preferred lottery out of the twenty.
10. The random lottery incentive system involves randomly selecting a task at the end of the experiment and resolving the risk involved

in that task to determine subject payment. This device controls for wealth effects and motivates subjects to consider all tasks carefully. The BDM device provides incentives to reveal true absolute valuations of lotteries by comparing reservation prices with a randomly generated offer. If the offer is below the valuation attached to a lottery then the subject plays out the lottery, but if it is equal to or above the valuation the subject gets the offer.

11. Opposite violations are those on the right-hand side of the strict inequalities in the alternative hypotheses in Table III.
12. Given that the observation sequence displayed a series of money outcomes, one might question why this would enhance the probability attribute of lotteries? An explanation is offered by the proposition that outcome frequency information is a basic building block in the formation of subjective *probability* assessments which ultimately contribute to the formation of probability weights.

REFERENCES

Alba, J. W. and Marmorstein, H. (1987), The effects of frequency knowledge on consumer decision making, *Journal of Consumer Research* 14, 14–25.

Bateman, I., Munro, A., Rhodes, B., Starmer, C. and Sugden, R. (1997), Does part-whole bias exist? An experimental investigation, *Economic Journal* 107, 322–331.

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

Binmore, K. (1999), Why experiment in economics, *Economic Journal* 109, F16–F24.

Camerer, C. (1995), Individual decision making, in Kagel, J. H. and Roth, A. E. (eds.), *The Handbook of Experimental Economics*, Princeton University Press, Princeton.

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

Cox, J. C. and Grether, D. M. (1996), The preference reversal phenomenon: Response mode, markets and incentives, *Economic Theory* 7, 381–405.

Cubitt, R. P., Starmer, C. and Sugden, R. (2001), Discovered preferences and the experimentally observed violations of expected utility theory, *Journal of Economic Methodology* 8, 385–414.

Diecidue, E. and Wakker, P. (2001), On the intuition of rank-dependent expected utility, *Journal of Risk and Uncertainty* 23, 281–298.

Einhorn, H. J. and Hogarth, R. M. (1978), Confidence in judgment: Persistence in the illusion of validity, *Psychological Review* 85, 396–416.

Estes, W. K. (1976a), The cognitive side of probability learning, *Psychological Review* 83, 37–64.

Estes, W. K. (1976b), Some functions of memory in probability learning and choice behavior, in Bower, G. H. (ed.), *The Psychology of Learning and Motivation*, Vol. 10. Academic Press, New York.

Friedman, D. (1998), Monty Hall's three doors: Construction and deconstruction of a choice anomaly, *American Economic Review* 88, 933–946.

Grether, D. and Plott, C. (1979), Economic theory of choice and the preference reversal rhenomenon, *American Economic Review* 69, 623–638.

Harrison, G. (1994), Expected utility and the experimentalists, *Empirical Economics* 19, 223–253.

Holt, C. (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. (1999), Probability learning, event-splitting effects and the economic theory of choice, *Theory and Decision* 46, 51–78.

Humphrey, S. J. (2001), Are event-splitting effects actually boundary effects? *Journal of Risk and Uncertainty* 22, 79–93.

Humphrey, S. J. and Bleaney, M. (2006), An experimental test of generalized ambiguity aversion using lottery pricing tasks, *Theory and Decision* 60, 257–282.

Johnson, E. J., Hershey, J., Meszaros, J. and Kunreuther, H.. (1993), Framing, probability distortions, and insurance decisions, *Journal of Risk and Uncertainty* 7, 35-51.

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

Loewenstein, G. (1999), Experimental economics from the vantage point of behavioural economics, *Economic Journal* 109, F25–F34.

Loomes, G. (1999), Some lessons from past experiments and challenges for the future, *Economic Journal* 109, F35–F45.

Loomes, G., Moffatt, P. and Sugden R. (2002), A microeconometric test of alternative stochastic theories of risky choice, *Journal of Risk and Uncertainty* 24, 103–130.

Plott, C. R.: (1996), Rational individual behaviour in markets and social choice processes: The discovered preference hypothesis, in Arrow, K. J., Colombatto, E., Perlman, M., and Schmidt, C. (eds.), *The Rational Foundations of Economic Behaviour*, MacMillan, Basingstoke.

Segal, U.: (1988), Does the preference reversal phenomenon necessarily contradict the independence axiom?, *American Economic Review* 78, 233–236.

Slovic, P. and Tversky, A. (1974), Who accepts Savage's axiom?, *Behavioral Science* 19, 368–373.

Smith, V. L. (1989), Theory, experiment and economics, *Journal of Economic Perspectives* 3, 151–169.

Starmer, C. (1999a), Cycling with rules of thumb: An experimental test for a new form on non-transitive behaviour, *Theory and Decision* 46, 141–158

Starmer, C. (1999b), Experimental economics: Hard science or wasteful tinkering? *Economic Journal* 109, F5–F15.

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

Tversky, A. and Kahneman, D. (1992), Advances in prospect theory: A cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 291–323.

Tversky, A. and Koehler, D. J. (1994), Support theory. A nonextensional representation of subjective probability, *Psychological Review* 101, 547–567.

Tversky, A., Slovic, P. and Kahneman, D. (1990), The causes of preference reversal, *American Economic Review* 80, 204–217.

Weber, M., Eisenführ, F. and von Winterfeldt, D. (1988), The effects of splitting attributes on weights in multiattribute utility measurement, *Management Science* 34, 431–445.

Wu, G., and Gonzalez, R. (1998), Common consequence conditions in decision making under risk, *Journal of Risk and Uncertainty* 16, 95–115.

Address for correspondence: Steven J. Humphrey, Centre for Decision Research and Experimental Economics, School of Economics, University of Nottingham, Nottingham NG7 2RD, UK. Tel.: +44-115-951-5472
E-mail: steve.humphrey@nottingham.ac.uk