DO INDIVIDUALS LEARN TO MAXIMISE EXPECTED UTILITY?

Steven J. Humphrey*

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 a test of whether conditions which facilitate objective probability learning yield decisions better described by expected utility theory than is the case in experiments devoid of learning opportunity. The data show that expected utility maximising behaviour increases with the learning opportunity, but so too do systematic violations. Learning, therefore, may exacerbate choice anomalies.

Keywords: Common consequence effect, monotonicity, probability learning

JEL classification: D81, C91

---

* I am grateful for financial assistance from The Nuffield Foundation, award no. SGS/LB/0295. I would like to thank David Lawson for running the experiments and Chris Starmer for valuable comments. Correspondence to: Steve Humphrey, Centre for Decision Research and Experimental Economics, School of Economics, University of Nottingham, Nottingham NG7 2RD, UK. +44 115 951 5472.

The decision-making under risk literature is abound with experimentally observed deviations from expected utility maximisation (von Neumann and Morgenstern, 1944). Recently investigators have begun to speculate that, despite such 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 this genuine underlying form (e.g. Plott, 1996).[1] This paper reports a test of the impact of a probability learning opportunity on perhaps the best known risky choice anomaly; the common consequence effect.[2] The test investigates two questions. First, does experience in observing the resolution of risk involved in lotteries prior to choice affect the nature of the decision? This type of learning opportunity is economically relevant. Investors can observe the performance (outcomes) of stocks prior to investment choices between risky assets (lotteries). Consumers can observe the incidence of aversive events prior to insurance decisions. Second, if probability learning affects revealed preferences, are those revealed preferences appropriately described by expected utility theory?

The interest in answering these questions lies in the possibility that the nature of decision-making tasks and the type of learning opportunities afforded by the environment shape preferences in different ways. Moreover, evidence relating directly to the influence of specific kinds of learning opportunities on risky choices is quite sparse. In this respect, the suggestion that a lack of appropriate opportunities for learning might create choice anomalies (as purely laboratory phenomena) may be overstated; there is no reason in principle why true preferences need imply expected utility maximisation. It is possible that anomalies are observed because they are features of genuine non-expected utility preferences. So even after learning there is no guarantee that preferences will necessarily home-in on expected utility maximisation (Loewenstein, 1999). It is thus of considerable importance to understand exactly how particular learning opportunities shape preferences, if at all, in specific decision-

[1] Henceforth, 'genuine' or 'true' preferences are taken as those which are free from imprecision attributable to a lack of understanding or experience of the decision-making task, such as not understanding the meaning of stated information, or confusion stemming from experimental procedures. True preferences, therefore, are those which are likely to emerge from learning from experience. Imprecise or noisy preferences are those which are subject to the influence of the above perturbing factors.
[2] The original version of the common consequence effect is called the Allais paradox (Allais, 1953).

making tasks. This knowledge can then be usefully fed into applied studies which involve the elicitation of values from members of the public in order to inform policy.

1. Learning in Experiments

1.1. Discovered preferences and the economists' view

One view of learning in experimental decision-making tasks is called the *discovered preferences hypothesis* (Plott, 1996). This perspective, also supported by Smith (1989), Harrison (1994) and Binmore (1999), is that individuals have a unique and precisely structured set of underlying preferences. In order for these preferences to be elicited in decision-making tasks, however, the individual will first have to discover which action best satisfies their preferences. Learning in experiments is some process which facilitates the discovery of underlying preferences. Examples of learning which are said to facilitate discovery include trial-and error decision-making, repetition (with feedback), deliberation, or a comparison of one's own choices and outcomes with those of others.

Loomes (1999) refers to evidence which supports the discovered preferences hypothesis. Loomes *et al.* (1998) and a yet-to-be-reported experiment conducted by Ken Binmore (see Loomes, 1999, p.F37, note 1) show choices converge on expected utility maximisation as subjects progress through sequences of pairwise lottery choice tasks. This evidence suggests that deviations from expected utility maximisation may be the product of imprecise preferences encountered whilst decision-makers are learning how the task interacts with their basic underlying values. Once this learning has taken place, choices converge on expected utility maximisation. If this interpretation is correct, choice anomalies may be somewhat inevitably observed due to the one-shot nature of (much) laboratory experimentation and/or the use of inexperienced subjects.[3]

[3] It is not entirely clear what this body of evidence is showing. Whereas the former experiment reveals a tendency for subjects to become increasingly inclined to consistently select the safer lottery in pairwise choice, the latter shows the opposite. Both experiments illustrate some kind of learning effect, but these effects appear to be different. For example, it may be that learning in Binmore's experiment caused subjects to realise that the riskier option has a higher expected value than the safer option and therefore engenders expected value maximisation. This could only be considered a rational strategy up to the point where expected value maximisers committed the St. Petersburg paradox. The St. Petersburg game (e.g. see Camerer, 1995) is one in which individuals do not pay an infinite (or even large) amount to participate in a game with infinite expected payoff. This is contrary to what an expected value maximisation rule would prescribe and is the observational seed from which grew expected utility theory.

Additional evidence supporting the discovered preference hypothesis is provided by Friedman (1998). Friedman (1998) shows that irrational behaviour in Monty Hall's three door problem diminishes when subjects are able to keep a record of their performance, take advice on strategy, or compare their performance with that of others.[4] A possible implication of the above evidence, as Friedman (1998, p.941) puts it, is that appropriately structured learning environments render the existence of anomalous behaviour (in the sense of stable and yet non-expected utility choices) unlikely. Experiments which report choice anomalies without appropriate opportunities for learning, therefore, cannot be taken as evidence against expected utility theory. Appropriate learning opportunity may prove anomalies to be transient.

1.2. Mixed evidence

Although the discovered preferences hypothesis does not propose *per se* that discovered preferences are appropriately described by expected utility theory, the hypothesis is often argued to this end (see, for example, Binmore, 1999). There may, however, be grounds upon which to be sceptical about the generality of the discovered preferences school's case for transient anomalies as it currently stands. A famous experiment conducted by Slovic and Tversky (1974) offered subjects who had made a series of choices the chance to change their initial decisions. Some of these initial choices (60%) violated the independence axiom of expected utility theory. Prior to the switching opportunity, subjects were presented with arguments as to why they should change their mind (e.g. why the independence axiom is or is not appealing, depending on whether the initial choice violated it or not). It transpired that the subsequent set of decisions yielded slightly more violations than the original decisions. In this case, learning certainly did not mitigate the anomaly. This evidence suggests that choice anomalies may be genuine and non-transient features of preferences. As such, opportunities for learning which contribute to the precision with which underlying preferences are expressed in behaviour will do little to mitigate their prevalence.

[4] The three door problem asks an individual to choose one of three doors. Behind one door is a good prize and behind the other two are booby prizes. When a door is chosen, one of the other doors is opened to reveal a booby prize and the individual is offered the chance to switch their chosen door for the other unopened door. People rarely switch their original choice, and this represents irrational behaviour because their win probability remains at 1/3 as opposed to 2/3 for switching.

Slovic and Tversky's (1974) experiment appears to involve an appropriate method to facilitate the discovery of true preferences. The 'taking advice on strategy' part of Friedman's (1998) experiment is strikingly similar. Yet it might be argued, as it is by some members of the discovered preferences school (Binmore, 1999), that this type of learning opportunity is not the kind which really interests economists. In economic markets 'bad' decisions often cause the decision-maker to fail to enjoy maximum possible welfare. It is this potential loss of welfare which disciplines decision-makers, possibly as a result of feedback on trial-and-error based decisions, imitation, and so on, into carefully considering their choices. There is evidence that the kind of discipline provided by economic markets may make anomalies disappear. Chu and Chu (1990) money-pumped subjects who violated the transitivity axiom of expected utility theory until, in light of fast-approaching experimental bankruptcy, they ceased to reverse their preferences. In the case of Slovic and Tversky's (1974) experiment financial penalties were not imposed on violators. It might thus be argued that continued violation, in the absence of incentives to do otherwise, is not surprising. But this argument would not explain why, in the similar absence of appropriate incentives, violations of the independence axiom *increased* after the learning opportunity.

The mixed evidence discussed above renders of fundamental importance the question of what types of learning should concern economists? It is certainly the case that experimental economists should be concerned with learning opportunities which facilitate the eradication of imprecision in preferences stemming from confusion about the mechanics of experimental tasks. It is also uncontroversial to suggest that economists should be interested in evolutionary-type learning provided by the market on the basis of feedback on the outcomes of repeated decisions. But what about other types of learning opportunity? The view held here is that economists should subscribe to a broad church. That is, just because Slovic and Tversky's (1974) experiment does not implement analogues of market discipline, and just because it shows divergence from expected utility maximisation, it should not be of diminished interest to economists. The guiding principle to dictate which types of learning should interest economists should be that the learning opportunity is *economically relevant*. Starmer (1999) argues that this includes decisions in one-shot (non-evolved) environments; examples of which are education decisions, employment decisions, major investment decisions, and so on. All of these decisions are often made on the basis of advice analogous to that provided by Slovic and Tversky (1974). The fact that different learning opportunities appear to exert different influences on choices merits further investigation of the nature of these influences. It does not suggest that economists should consider as relevant only those types of learning which result in conventional economic theories working.[5]

2. Some Probability Learning Theory

2.1. Economically relevant probability learning

In the experiment reported below subjects are asked to value a series of lotteries by stating an amount of money which makes them indifferent between the lottery and the valuation attached to it. Prior to making these valuations one group of subjects is provided with the opportunity to learn from the experience of observing a sequence of 10 resolutions of risk in the lottery. Another group is asked to make their valuations on the basis of stated lottery information alone. The observation sequence would probably not satisfy Friedman's (1998) criteria for an appropriately structured learning opportunity. It does not impose market-like discipline on decision-makers, it does not facilitate repeated choices with intermittent feedback, and it does not facilitate imitation. It is, however, economically relevant. As stated in the introduction, investors can observe the performance (outcomes) of stocks prior to investment choices between risky assets (lotteries) and consumers can observe the incidence of aversive events prior to making insurance decisions. There are, moreover, theoretical and empirical reasons why presenting probability information in terms of frequencies (e.g. 6 out 10 draws gave £10) as well as stating it (e.g. £10 with probability 0.6) might be expected to cause decisions to better reflect expected utility maximisation. These reasons lie in the economic theory of risky choice and the behavioural psychology of probability learning. Each will be considered in turn.

2.1. Prospective reference theory

Viscusi's (1989) prospective reference theory proposes that the expected utility of a gamble [EU*(L)] is given by expression (1):

[5] As Starmer (1999, p.F12) puts it, "…if only those phenomena which, *a priori*, we expected to be consistent with standard notions of optimisation are allowed to count as economics, and anything else is sociology or psychology, successful prediction is not much to write home about."

$$\mathrm{EU^*(L)} \quad = \quad W_1\left[\sum_{i=1}^{n} (1/\mathrm{n})U(x_i)\right] + W_2\left[\sum_{i=1}^{n} p_i U(x_i)\right] \qquad (1)$$

$0 \leq W_1 \leq 1,\ W_1 + W_2 = 1$

In expression (1) n is the number of outcomes x_i in a lottery which occur with probability p_i, and $U(.)$ denotes the utility of each outcome. Prospective reference theory views the expected utility of a lottery as a weighted average of (a) the expected utility on the assumption that each outcome is equally probable (the first term) and (b) von Neumann and Morgenstern (1944) expected utility (the second term). The first term in square brackets represents the expected utility of the lottery from the perspective of a reference risk level, which views each outcome as equally likely. In this respect the number of outcomes in a lottery is important in determining the reference risk level. Viscusi (1989) suggests that reference risk levels are important even when experimental probabilities are explicitly stated because subjects may not believe them.

The weights $W_1=\alpha/\alpha+\beta$ and $W_2=\beta/\alpha+\beta$, where α and β are non-negative constants, reflect the relative information content of the reference risk level and the stated lottery. These weights derive from the perceived probability function $P(p_i)=[\alpha(1/\mathrm{n})+\beta p_i]/[\alpha+\beta]$ where the sum of $P(p_i)$ over i=1,…,n outcomes equals 1. The perceived probability function is a Bayesian modification of stated probabilities according to the proportion of total information $\alpha/\alpha+\beta$ attached to the prior reference risk probability of 1/n and the proportion of total information $\beta/\alpha+\beta$ attached to stated probabilities (which can be thought of as a p_i fraction of outcomes i from β trials). If, for example, W_2 is relatively small and W_1 is relatively large, the reference risk level is particularly important in driving choices. Prospective reference theory provides a general reason why individuals who have experience of observing lottery draws prior to choice might behave more consistently with expected utility theory than those who do not. If prospective reference theory is correct, the person who observes lottery draws prior to choice has extra information which (except in the special case where W_1=0) should increase W_2 at the expense of W_1.

2.3. Frequency-based probability learning

In addition to the insights of prospective reference theory, cognitive psychology provides a reason to expect the observation of lottery draws prior to choice to influence decision-making behaviour. In influential work on probability learning Estes (1976a,1976b) argues that human decision-makers do not possess ultra-sophisticated statistical ability in dealing with probabilistic information. Behaviour rather reflects the reliance on simple heuristic devices. One such device is learning on the basis of experience garnered from observing repeated situations, coupled with faith in the uniformity of nature. This yields a basic decision-making heuristic: more frequently observed outcomes are *ceteris paribus* more likely to be future outcomes. If repetitions of similar circumstances give different outcomes, individuals form expectations on the basis of converting absolute event frequencies into relative event frequencies. On this view the basis for probability learning is the learning of absolute frequencies of event occurrence. Estes's (1976a) evidence suggests that the coding of event frequencies in memory and their conversion to decision weights often generates decision-making and judgmental biases (due to limitations in memory capacity or inadequate training). The particular bias Estes (1976a, 1976b) observes is that events which occur more frequently in observation trials than is suggested by their true probability are believed to be more likely to occur in the future than is suggested by their probability.[6]

Estes's (1976a, 1976b) also shows that when event frequency coincides with event probability, individuals behave as though they have accurately and efficiently learned probability information. Additionally, Alba and Marmorstein (1987) conclude that simple frequency heuristics may be used for learning, even under conditions designed to engender processing of all decision-relevant information. Because the information load is high, the overriding concern of individuals is to economise on expenditures of cognitive effort. The use of frequency heuristics contributes to this economy. In a similar vein to Einhorn and Hogarth (1978), they also conclude that frequency information is learned and remembered

[6] Alba and Marmorstein (1987) discuss deviations from optimal choice stemming from the use of frequency heuristics in consumer choice. Humphrey (1999) replicates Estes's (1976a, 1976b) results in an experiment with significant financial incentives for accurate probability learning. Einhorn and Hogarth (1978) are concerned with how frequency-based probability learning produces fallible judgements by both experts and non-experts who, nevertheless, express great confidence in those judgements. They concur that the coding of event frequencies rather than probability results in frequency information being more salient than probability information.

more completely than other information. It is therefore more salient than probability information at the point of choice. Consider the implications of these facets of frequency-based probability learning for decisions also involving stated probability information in the following example.

Recent research has attributed many robust violations of expected utility theory to popular decision-weighting models such as rank-dependent expected utility theory (e.g. Quiggin, 1982; Yaari, 1987; Tversky and Kahneman, 1992). Wu and Gonzalez (1998), for example, attribute the common consequence effect to rank-dependent expected utility with an 'inverse-S' shaped decision weighting function. This weighting function is approximately well-behaved around the 'salient' endpoint probabilities such that $\pi(1)=1$ and $\pi(0)=0$, but is less so elsewhere. If, in the context of these models, frequency-based probability learning is to engender expected utility maximising behaviour, it needs to promote a reduction in the non-linearity of the decision weighting function. By allowing reflection on rank-dependent decision weights formed on the basis of stated probabilities, frequency-based probability learning may achieve exactly this. Following Diecidue and Wakker (in press), consider the lottery (£30, 0.4; £20, 0.5; £10, 0.1). Rank-dependent models suggest 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 decision-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 (in press) point out that rank-dependent decision 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 decision weight $\pi_{10}(0.1)=0.3$ as placing too much emphasis on an unlikely outcome. Alba And Marmorstein's (1987) conclusion that frequency information is more salient in memory at the point of choice than probability information, enhances the likelihood of this type of reflection. Correspondingly, even the pessimist, upon observing £30 to have occurred four times out of ten, might regard the decision 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 decision weights (formed on the basis of stated information) to engender their modification in the direction of linearity. All that is required to form accurate probability assessments is the observation of outcomes over a sufficiently long period such that relative frequency of event occurrence is not greatly at odds with objective event likelihood. If so, rank-dependent decision-weighting based violations of expected utility theory would be diminished.[7]

3. The Common Consequence Effect

The experimental test of whether frequency-based probability learning causes convergence on expected utility maximisation is designed around three pairs (1, 2 and 3) of lotteries as described in table 1.

Table 1: Common Consequence Lotteries*

Lottery		Probability of Outcomes		
		£21	£9	zero
Pair 1	*S1*	-	0.9	0.1
	R1	0.2	0.5	0.3
Pair 2	*S2*	-	0.4	0.6
	R2	0.2	-	0.8
Pair 3	*S3*	0.5	0.4	0.1
	R3	0.7	-	0.3

* *S* lotteries are 'safer' in that they offer a higher probability of winning a non-zero amount and *R* lotteries are 'risker'.

[7] 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 decision weighting function may increase. Associated deviations from expected utility maximisation may accordingly persist or increase. Indeed, Loewenstein's (1999) argues that there should be no presumption that preferences formed on the basis of learning opportunities will necessarily home-in on expected utility maximisation.

Each pair of lotteries is generated from each of the other pairs by shifting a probability mass of 0.5 between the outcomes. 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. The common consequence effect is usually observed by asking subjects to choose between the lotteries in each of the three pairs. These choices can be depicted in the unit probability triangle in figure 1.

Figure 1: Common Consequence Lotteries in the Unit Probability Triangle

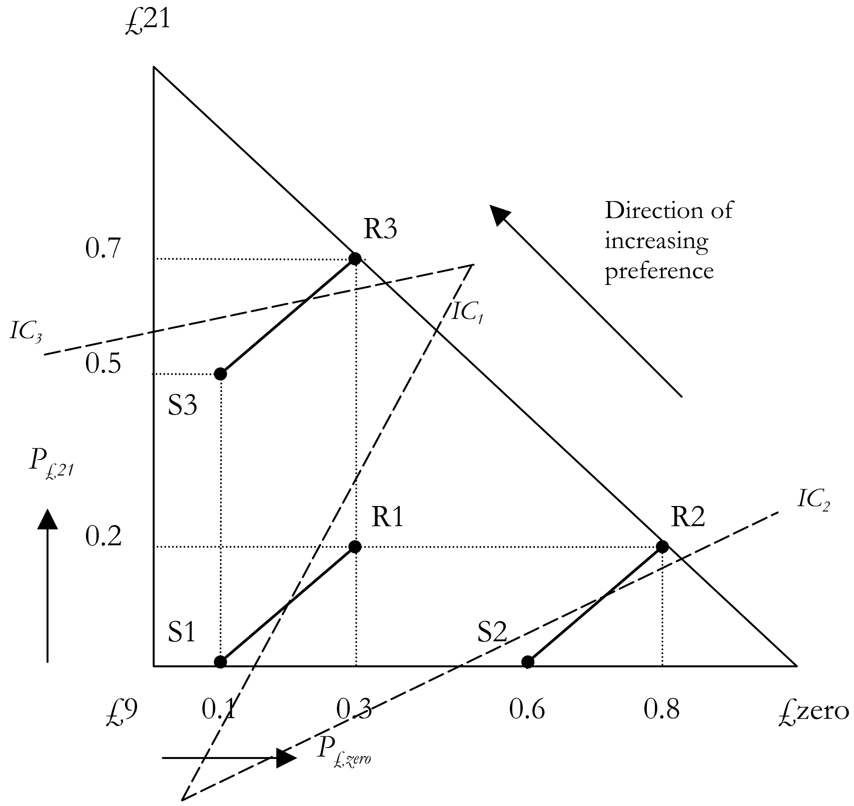

In figure 1 the largest outcome is placed at the top corner, the intermediate outcome at the right-angled corner and the lowest outcome at the bottom right-hand corner. The vertical axis shows the probability of the largest outcome and the horizontal axis the probability of the smallest outcome. The probability of the intermediate outcome is 1 minus the other two. Any point in the triangle therefore represents a lottery. The six lotteries in table 1 have been shown in figure 1. Each pair of lotteries which represents a problem in standard common consequence effect experiment is also shown on the diagram. As can be seen, starting from problem 1 (*S1* vs. *R1*) and shifting the probability mass between outcomes, the problem moves horizontally along the bottom edge of the triangle or vertically towards the top of the triangle. These movements are respectively testing for horizontal and vertical common consequence effects. Comparing problems 2 (*S2* vs. *R2*) and 3 (*S3* vs. *R3*) tests for a north-west common consequence effect.

The differences between all three pairwise decision problems are common to each respective lottery in the problems. Expected utility maximisation therefore demands either the riskier option to be chosen in all three problems (*R1*, *R2* and *R3*), or the safer option to be chosen in all three problems (*S1*, *S2* and *S3*), or indifference to be expressed in all three problems. A common consequence effect is manifest if choices switch systematically between the riskier and the safer options as the problems move horizontally, vertically and in a north-westerly direction. The indifference curves IC_1, IC_2 and IC_3 in figure 1 describe the common pattern of choices observed over problems of this type.[8] Expected utility theory implies the existence of upwards-sloping, linear and parallel indifference curves in the triangle. So in conjunction with the north-westerly direction of increasing preference, IC_1 and IC_2 represent *S1* being preferred to *R1* but *R2* being preferred to *S2*. The shape of these indifference curves gives rise to the horizontal common consequence effect also being referred to as *horizontal fanning-out* (from some point outside the bottom left-hand corner of the triangle). A similar tendency to switch from the safer option to the riskier option in the movements from problem 1 to problem 3 and problem 2 to problem 3 respectively causes IC_1 with IC_3 and IC_2 with IC_3 to 'fan-in' towards some point to the right of the triangle. The overall pattern of preferences illustrated by the set of indifference curves is known as *mixed-fanning*.

The most popular theoretical alternatives to expected utility theory predict mixed-fanning (e.g. rank dependent expected utility theory). The bulk of the experimental evidence regarding common consequence problems indeed suggests mixed-fanning to be robust. There is, however, a substantial amount of evidence which describes common consequence effects which do not conform to this pattern. A number of studies discussed by Humphrey (2000), for example, fail to replicate horizontal fanning-out along the bottom edge of the triangle, or

[8] The indifference curves in figure 1 are illustrated stylistically as being linear. See Camerer (1995) for a discussion of related evidence, including that which suggests non-linear indifference curves.

show horizontal fanning-in.[9] This experiment is concerned with testing common consequence effects in valuations task rather than choice tasks. To my knowledge, there is no empirical guidance as to what pattern of preferences this may reveal. It is therefore necessary to characterise common consequence effects in terms of both fanning-in and fanning-out. This is shown in Table 2.

Table 2: Common Consequence Effects in Valuation Tasks*

	Fanning-out	**Fanning-in**
Horizontal	$V(S1) \geq V(R1)$ and $V(R2) \geq V(S2)$	$V(R1) \geq V(S1)$ and $V(S2) \geq V(R2)$
Vertical	$V(S1) \geq V(R1)$ and $V(R3) \geq V(S3)$	$V(R1) \geq V(S1)$ and $V(S3) \geq V(R3)$
North-West	$V(S2) \geq V(R2)$ and $V(R3) \geq V(S3)$	$V(R2) \geq V(S2)$ and $V(S3) \geq V(R3)$

* V(.) represents the valuation assigned to a lottery. A violation of expected utility theory requires at least one strict inequality in each of the six pairs.

Note that in within-subject valuation tasks a single subject values each of the six lotteries only once. This means that it is perhaps unlikely to observe a single subject exhibiting violations of expected utility theory over all horizontal, vertical and north-west comparisons. For example, it is only possible for a subject to reveal preferences which indicate indifference curves that universally fan-out if $V(R2)=V(S2)$. Similarly, mixed-fanning would require $V(R1)=V(S1)$. Although these patterns of behaviour are entirely possible, it should perhaps be expected that in valuation tasks an expression of indifference between lotteries valued separately (manifest in equality between valuations) is less likely than those valuations indicating a strict preference. Nevertheless, the current design does (for example) allow an individual to either fan-out horizontally and vertically or vertically and north-westerly without expressing indifference. This is sufficient opportunity for individuals to commit the common consequence effect in order to investigate whether pre-decisional probability learning opportunities impact on any such violations which may emerge.

[9] One such study is Starmer (1992).

4. Experimental Design

4.1. Valuation task

The experiment involves 2 conditions. In condition 1 subjects are asked to place a money value on a simple lottery. In condition 2 subjects perform the same task as in condition 1, but first observe a sequence of 10 resolutions of the risk in the lottery. Figure 2 shows an example of a valuation screen from the experimental software. The top 'lottery' box shows the lottery to be valued. All lotteries were expressed in terms of 10 lottery tickets. The lottery in figure 2 corresponds to *S1* in table 1. The 'yardstick' is the vehicle through which the lottery is valued. Subjects were told that they should value the lottery by entering an amount in the small box at the bottom of the screen (which would also appear in the small box within the yardstick) which makes them indifferent between the lottery and the yardstick.[10]

Figure 2: Task Display

Question One

LOTTERY:	Lottery Tickets	1 to 9	pay you	£9
	Lottery Ticket	10	pays you	zero

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

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

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.

[10] Subjects were told that indifference meant that, after they had entered their valuation, they would not mind whether they received either the lottery or the yardstick. The valuation was made by using the 'up' and 'down' cursor keys on the keyboard. Pressing the up key replaced the question marks with £00.00, pressing again incremented this to £00.10 and so on. The down key, generated 10 pence decrements. There was no upper bound on valuations and a zero lower bound. The valuation was confirmed by pressing 'enter', followed by a chance to change it or move on to the next problem.

The second condition in the experiment involves subjects observing the sequence of lottery draws prior to making their valuation. This aspect of the design is captured by the box in figure 2 identified with 'draw' and 'winnings'. Subjects in this group would first see the lottery and then, when ready, press 'enter' to reveal the observation box showing draws 1 to 10 and the empty winnings row. Pressing 'enter' again would start the observation sequence wherein the computer would reveal the outcome of a single draw of the lottery under draw 1 in the winnings row, pause and then repeat the process up to draw 10. So, in terms of figure 2, the first draw gave £9, the second gave £9, the third gave nothing, and so on. After the observation sequence had finished the valuation message would appear on the screen and the subject would proceed to value the lottery.

To provide the strongest test of whether this type of pre-decisional learning opportunity influences the extent violations of expected utility theory associated with non-linear decision weights (by causing decision weights to more precisely reflect stated probabilities), the observation draws were fixed. Each outcome occurred in the sequence of 10 draws in the exact frequency to that which would be suggested by its probability. As can be seen in figure 2, the probability of winning £9 is 0.9 and the number of times £9 occurs in the observation sequence is 9 out of 10. Outcome information was simply provided in a manner analogous to a 'speeding-up' the law of averages. If the frequency-based probability learning hypothesis is correct, it is this which may mitigate any violations of expected utility theory.[11] It is possible that the order in which outcomes occurred during the observation sequence could bias the nature of any probability learning which may occur. To control for this the order of observation outcomes was randomly determined for each subject.

[11] Most probability weighting models suggest individuals overweight small probabilities of positive outcomes. If, in a *genuinely random* observation sequence, such a positive outcome occurs more times than is suggested by the probability, subjective overweighting would be reinforced. Any associated decision anomalies would be exacerbated. This is an equally valid reflection of what might happen in the real world to the 'fixed' observation sequence. But it does not provide a pure test of whether probability learning on the basis of outcome observation mitigates choice anomalies. This is the question dealt with by this paper. The more commonly asked question of whether unrepresentative outcome feedback can distort probability assessments and introduce deviations from expected utility theory, is answered elsewhere (e.g. Humphrey, 1999). Fixing the observation sequence should not regarded as deception. Valuations were made on the basis of genuine and complete information regarding the probability distribution governing outcomes. Subjects were told that the outcome of one of the 20 lotteries which would determine their payment for participation in the experiment, would be determined by a single resolution of the risk in that lottery according to the stated probabilities.

4.2. Incentives

The incentive system employed in the experiment is a variation of what Tversky, Slovic and Kahneman (1990) call the ordinal payoff scheme. With this scheme it is in subjects interests to attach valuations to the set of lotteries which reflect their true preference *ordering* over the set of lotteries. At the end of the experiment two lotteries were randomly selected from the set (by drawing 2 chips from a bag containing 20 consecutively numbered chips). The valuations attached to those lotteries were compared. The risk in the lottery to which the *highest* value was attached resolved (by drawing a chip from a different bag containing 10 consecutively numbered chips). The outcome of the draw would be the subjects payment for participation in the experiment.[12] If valuations do not reflect true preference orderings, a feature of this design is that subjects may play-out one of the randomly selected lotteries for real money when they would have preferred to have played-out the other.

Note that valuations which are above the highest outcome offered by a lottery, and often taken as evidence of irrationality or confusion, are incentive compatible within this design. Valuations are only taken to establish a preference ordering over lotteries. Subjects could therefore perform any monotonic transformation on their valuations and still represent their true underlying preference ordering. It was emphasised to subjects that they would not know which two lotteries are randomly selected (to ultimately determine payment) until the end of the experiment when all tasks had been completed. They were told that they could guarantee playing out their truly preferred lottery from the randomly selected pair, whatever that pair was, by considering each lottery carefully and valuing it genuinely.[13]

The ordinal payoff scheme is favoured over the BDM mechanism (Becker *et al.*, 1964) for the following reason. The BDM scheme in conjunction with the random lottery incentive system elicits absolute valuations by asking subjects to state reservation prices for lotteries. At the end of the experiment one of the lotteries is randomly selected and the reservation price is

[12] In the event that the valuations to the two randomly selected lotteries were the same the payment lottery was determined by flipping a coin. Subjects were informed of this prior to making their valuations.

[13] With this design subjects could effectively eliminate one lottery from the set of potential real payment lotteries by assigning it a zero valuation. It is unclear why subjects would want to do this in any lottery other than that which they valued last. Whilst valuing all previous lotteries they do not know what the subsequent lottery will be and so they can only be sure of their least preferred lottery if that happens to be the last one valued. Of course, if this is the case, elimination of that lottery is entirely consistent with genuine preference.

compared to a randomly generated offer. If the reservation price is below the offer the subject receives the offer. If the reservation price is equal to or above the offer the subject plays out the gamble. It has been shown (Karni and Safra, 1987; Segal, 1988) that experiments which use the BDM mechanism are only theoretically reliable in eliciting genuine absolute valuations if the independence axiom of expected utility theory holds. There is plenty of evidence that it does not (e.g. see Camerer, 1995). As Tversky, Slovic and Kahneman (1990) point out, these concerns can be mitigated by using the ordinal payoff scheme because it does not involve the BDM device. It is not true, however, that ordinal payoff schemes alleviate the concern entirely.

Holt (1986) argues that individuals may treat experiments as a single large decision problem between compound lotteries which have first been simplified by the calculus of probabilities according to the reduction principle. Behaviour is determined by preferences over these simplified compound lotteries. For example, consider the two ways in which common consequence effects can be observed be over the horizontal lotteries in table 1. Let V(.) denote the valuation attached to each lottery. Horizontal fanning-out is observed if $V(S1)>V(R1)$ and $V(R2)>V(S2)$. Fanning-in is observed if $V(R1)>V(S1)$ and $V(R2)>V(S2)$.[14] If, as in Holt's (1986) argument, an individual treats the experiment as a single large decision problem via the reduction of compound lotteries, then fanning-out is equivalent to selecting lottery L_1:

$$L_1 = [S1, 2/n(n-1); R2, 2/n(n-1); Z, 1-(4/n(n-1))]$$

The first term in L_1 is the probability of playing *S1* for real out of *S1* and *R1*. The probability of either *S1* or *R1* being randomly selected as one of the lotteries whose valuations will be compared to determine the payment lottery is 2/n, where n is the number of lotteries to be valued in total. The probability of the lottery from *S1* and *R1* which was not selected as the first comparison lottery being selected as the second comparison lottery is 1/(n-1). Given that $V(S1)>V(R1)$ this yields a probability of *S1* and *R1* being compared and *S1* being played for real money of 2/n(n-1). A similar argument extends to *R2* being played for real money from

[14] This assumes for simplicity that lotteries are not equally valued in any of the pairs.

the pair *R2* and *S2*. *Z* is given by behaviour over all n-4 remaining tasks in the experiment and occurs with the residual probability. Similarly, fanning-in is equivalent to selecting L_2.

L_2 = [*R1*, 2/n(n-1); *S2*, 2/n(n-1); *Z*, 1-(4/n(n-1))]

If it is to be claimed that a systematic preponderance of fanning-out over fanning-in, or vice-versa, represents genuine preferences over the two pairs of lotteries involved, within the context of a broader set of n-4 additional lotteries, then it must be the case that the independence axiom of expected utility holds. If this is not the case, Holt's (1986) argument suggests the common term in *Z* could drive a perturbing wedge between true preferences and the observation of particular valuation patterns in the ordinal payoff scheme described above. However, inserting the parameters from table 1 into L_1 and L_2 and applying the reduction principle, respectively gives L_1* and L_2* (e.g. in L_2, a 2/n(n-1) times 0.5 chance of £9 from R1 and a 2/n(n-1) times 0.4 chance of £9 from S2, gives a 'reduced' overall 1.8/n(n-1) chance of £9):

L_1* = [*£21*, 0.4/n(n-1); *£9*, 1.8/n(n-1); 0, 1.8/n(n-1); *Z*, 1-(4/n(n-1))]

L_2* = [*£21*, 0.4/n(n-1); *£9*, 1.8/n(n-1); 0, 1.8/n(n-1); *Z*, 1-(4/n(n-1))]

Since L_1* and L_2* are identical in all respects, a systematic tendency towards a common consequence effect in one direction (either fanning-out or fanning-in) rather than the other violates the reduction principle. This undermines a central component of Holt's (1986) argument. Systematic common consequence effects, therefore, cannot not be explained by the use of the ordinal payoff scheme and subjects treating the experiment as a single decision problem via the reduction of compound lotteries.

4.3. Implementation

The experiment was conducted at the *Centre for Decision Research and Experimental Economics* (CeDEx) laboratory at the University of Nottingham during February and March 2001. An e-mail was sent to a mailbase of pre-registered volunteers to invite them to reserve a place in one of a number of prearranged sessions. It was randomly determined in advance whether each session would be a group 1 (valuation only) or group 2 (observation and valuation) session. Each group contained 67 subjects. The subject pool comprised students from a broad range of academic departments in the University. In total, 79 (59%) of the subjects were male. Each session lasted approximately one hour. Average subject payment was £12.97. All subjects responded to the same 20 valuation tasks. The only difference between the two groups was whether they saw the pre-valuation observation sequence or not.[15] When subjects arrived at the laboratory they were asked to sit at a computer terminal. Instructions were read out by the experiment organiser and subjects responded to two practice tasks prior to valuing the set of 20 lotteries. The valuations attached to the two practice lotteries were compared to illustrate how winnings would be determined and to provide a vehicle through which the incentive mechanism could be explained in detail. In addition to the randomised order of observation outcomes described above for group 2 subjects, all subjects responded to the valuation tasks in random order. No time limit was imposed.

5. Results

5.1. Hypotheses tested

The first hypothesis test will establish whether the data reveal evidence of common consequence effects as detailed in table 2. These can be manifest in either fanning-in or fanning-out. The null hypothesis will correspond to valuations being made according to expected utility theory; any patterns of valuations which imply common consequence effects are the result of random mistakes. There should, therefore, be no significant difference between the incidences of fanning-in and fanning-out on any two pairs of lotteries. The alternative hypothesis corresponds to the mixed-fanning of indifference curves as illustrated in figure 1. This entails there being significantly more horizontal fanning-out and vertical and north-west fanning-in than horizontal fanning-in and vertical and north-west fanning-out.

The second hypothesis seeks to establish whether there are any differences in behaviour between group 1 and group 2 subjects. The null hypothesis is that the frequency-based probability learning opportunity offered to group 2 subjects does not influence behaviour. The alternative hypothesis is that the observation sequence influences valuations to generate differences between the behaviour each group of subjects. The insights of prospective

[15] 6 of the 20 lotteries are described table 1. 2 additional lotteries investigate whether probability learning influences violation of monotonicity. These are described in the results section. Remaining lotteries were concerned with other hypotheses. Subjects in group 2 saw the observation sequence for all 20 lotteries.

reference theory and frequency-based probability learning suggest that the alternative hypothesis should work in the direction of increasing expected utility maximising behaviour.

5.2. The common consequence effect

Before turning to the hypothesis tests described above it is illustrative to place them in the context of lottery valuations elicited in the experiment. Average valuations and their standard deviations are described alongside the expected value of each lottery in table 3. Note that since the incentives in the experiment do not *require* valuations to represent individuals' reservation prices for each lottery (monotonic transformations are incentive compatible), these values should be interpreted liberally. Nevertheless, table 3 shows average valuations to be broadly similar to expected values and on this basis it seems plausible to suggest that subjects generally assigned values of the order one might expect on the basis of a statement of true certainty equivalents. This should not be taken to imply that valuations were assigned to lotteries on the basis of an expected value maximisation rule.[16]

Table 3: Average Lottery Valuations[*]

Lottery	**E.Val.**	**Valuations**			
		Group 1		Group 2	
		mean	s.dev	mean	s.dev
S1	8.1	7.95	1.93	9.10	4.69[b]
R2	8.7	9.34	3.39	9.51	3.88[c]
S2	3.6	3.54	1.24	4.32	2.43
R2	4.2	3.53	1.27[a]	5.06	3.38
S3	14.1	12.70	3.04	13.13	4.48[d]
R3	14.7	12.91	3.85	13.94	4.35[e]

[*] The 'E.Val.' column shows the expected value of each of the lotteries outlined in table 1. Outliers have been removed as follows: [a] 99.9, [b] 50.5 and 72.0, [c] 48.0, [d] 55.0 and 75.3,[e] 99.9.

Table 4 reports the results of the test for common consequence effects. Taking the group 1 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

[16] With an ordinal payoff scheme the outliers removed from table 3 should not be taken as evidence of irrationality. Also the number of outliers small in relation to the data set. In group 2 for example, there are 6 outliers from 1340 observations (i.e. 67 subjects each making 20 valuations).

always valuing the riskier and safer lotteries identically. This means that over half (54%-57%) of the patterns of valuations assigned by group 1 subjects violate expected utility theory. Despite this overall violation rate, the data in the fanning-out and fanning-in columns show violations to be broadly equally distributed between the two possible directions. There is perhaps a slight tendency towards fanning-in, particularly over the lotteries in the vertical comparison (33% against 24% fanning-out). The *p-value* column shows these differences to be insufficient to yield a rejection of the null hypothesis in favour of the alternative of mixed-fanning. The group 1 data do not rule-out the possibility that violations of expected utility theory are due to imprecise preferences. These preferences may give rise to noisy valuations, possibly due to unfamiliarity with the task or a lack of decision making experience in the absence of opportunities to learn.

Table 4: The Common Consequence Effect[*]

	Number of subjects							
	Group 1 (n=67)				Group 2 (n=67)			
Comparisons of Lottery Pairs	EUT	fanning-out	fanning-in	*p-value*	EUT	fanning-out	fanning-in	*p-value*
Horizontal	31 (46%)	17 (25%)	19 (28%)	0.9953	41 (61%)	12 (18%)	14 (21%)	0.4225
Vertical	29 (43%)	16 (24%)	22 (33%)	0.2088	30 (45%)	10 (15%)	27 (40%)	0.0038*
North-west	31 (46%)	16 (24%)	20 (30%)	0.7975	36 (54%)	10 (15%)	21 (31%)	0.0354*

[*] The EUT 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 $i \neq j$. These valuations do not imply a common consequence effect. The fanning-out and fanning-in columns represent patterns of within-subject valuations as described in table 2. The *p-value* column reports a test based on the binomial distribution of the hypothesis that observations required by mixed-fanning are at least as frequent as the opposite violations. An asterisk denotes a significant common consequence effect in the direction consistent with mixed-fanning at the 5%-level. Percentages may not sum to 100% due to rounding.

One possible reason for failure to observe systematic common consequence effects in group 1 is the parameters of the lotteries employed. Most common consequence studies use sets of problems which involve *S1* and *R1* being 0.1 further to the left than they are in figure 1. *S3* and *R3* are also often located both 0.1 further to the left and 0.1 higher in the vertical direction. This would make *S1* a certainty (as it is in the original Allais paradox) and place it at what Conlisk (1989) calls a 'double boundary point'. *S3* would also lie on the left-hand boundary of the triangle rather than in the interior. Conlisk (1989) argues that since certainties involve only one consequence they may be especially attractive in relation to single-boundary lotteries which entail more than one consequence. Individuals value the simplicity of a sure-thing, and it is this which generates the Allais paradox. He tests this hypothesis by displacing the Allais lotteries such that they lie marginally inside the triangle boundary and reports reduced incidences of, and no longer systematic, violations of expected utility theory. The lotteries employed here do not involve sure-things because this renders the potential for frequency-based probability learning minimal. What is there to learn about a certainty? In terms of rank-dependent theory, the decision weight attached to a certainty is $\pi(1)=1$. Thus, in the context of this model, observation trials provide no potential for reflection over whether this decision weight appropriately captures the importance of the outcome in the lottery. Moreover, common consequence effects have also been observed over lotteries which do not involve sure-things (e.g. Humphrey, 2000).

The group 2 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 higher than that of 45% under group 1. There are a total 94 out of a possible 201 violations of expected utility theory under group 2 compared with 104 under group 1. A test of difference in sample proportions over these pooled valuations yields Z=1.5962. This is significant at 6%. There is tentative evidence therefore to suggest that the probability learning opportunity mitigates violations of expected utility theory. The strongest evidence is provided by the horizontal lottery comparisons. Here 36/67 (54%) of group 1 valuation patterns violate expected utility theory. In group 2 this falls to 26/67 (39%). A one-tailed test of a difference in sample proportions based on the normal distribution yields a Z-value of 1.7326. This is significant at the 5%-level.

The violations of expected utility theory in group 2 vary between 39% and 55% (averaging 47%) of valuation patterns. They are also distributed differently to those under group 1. This distribution yields significant common consequence effects at the 5%-level (respectively greater than 0.5% and 4%) under the vertical and north-west comparisons. In these comparisons the split of valuation patterns is in the direction consistent with mixed-fanning. Significant common consequence effects do not emerge under the horizontal comparison, but this may be due to the fact that for any individual to exhibit all components of mixed-fanning they would have to assign values of $V(R1)=V(S1)$. This degree of precision is perhaps unlikely given the nature of the task.

The implication of a comparison of the group 1 and group 2 data is that frequency-based probability learning influences behaviour. The nature of this influence is to increase overall consistency with expected utility theory whilst simultaneously rendering violations systematic. This result places a slightly different slant on the evidence discussed by Loomes (1999) which shows choices to converge on expected utility maximisation as sequences of pairwise choices are answered. It seems that overall convergence on expected utility maximisation does not rule-out a proportion of systematically anomalous behaviour. This evidence might be interpreted as substantiating Loewenstein's (1999) argument that there should be no presumption that preferences formed on the basis of learning opportunities will home-in on expected utility maximisation. For some individuals at least, it seems that choice anomalies are not necessarily transient.

The group 2 data in relation to group 1 reveal an increased tendency for valuation patterns to exhibit fanning-in alongside a decreased tendency for them to exhibit fanning-out. A test of difference in sample proportions based on the normal distribution does not lead to rejection of the null hypothesis that the probability learning opportunity significantly influences either fanning-in or fanning-out alone, in any of the comparisons.[17] This observation suggests that valuations in group 1 are made on the basis of rather noisy, or imprecise, preferences. Valuation patterns which deviate from expected utility maximisation distribute themselves approximately evenly between fanning-in and fanning-out. But when in group 2 the observation trials facilitate reflection on probability assessments made on the basis of stated information, frequency-based probability learning mitigates the imprecision. Subjects are better able to identify their true preferences and less noisy patterns of valuations emerge. The question that this interpretation poses is exactly why does frequency-based

[17] This test yields Z-values of 1.0489, 1.3103 and 1.3107 for horizontal, vertical and north-west comparisons respectively in group 1, and 1.0025, -0.8968 and -0.1875 in group 2.

probability learning appear 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 *$*-bet (offering a high money prize with low probability) is assigned a higher reservation price than a *P*-bet (offering a lower money prize, but with a higher probability), but is subsequently not chosen in a direct choice between the two. This pattern of behaviour is often attributed to *response mode* effects. One feature of response mode effects is *compatibility*. The compatibility hypothesis states that money is the salient attribute of lotteries in money valuation tasks (the two are compatible). This renders the high prize in the *$*-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. The compatibility hypothesis provides an explanation of why systematic violations of expected utility theory were not observed in group 1, but were in group 2. Assume common consequence effects are the product of how probabilistic biases influence the decision weighting function. The salience of the money attribute of the lotteries in the group 1 valuation tasks may have precluded the emergence of any such probability-driven anomalies. In group 2, the observation sequence may have enhanced the salience of the probability attribute such that some of the imbalance under group 1 was redressed. If common consequence effects are probabilistically based, this would explain their systematic emergence in group 2.[18]

5.3. Monotonicity

The experiment involved two additional tasks to test whether valuations satisfy monotonicity. Monotonic preferences are arguably a fundamental property of rational choice. There is a view that violations of mononicity are not systematic features of genuine preferences,

[18] 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 decision weights.

irrespectively of whether those preferences are described by expected utility theory or alternative non-expected utility theories. In this respect, monotonicity violations are considered erroneous. As such they should diminish given a suitable learning opportunity to facilitate a better identification of preferences. To investigate this possibility each group of subjects faced two valuation tasks where one strictly dominated the other. Lottery *D1* offered a 0.7 chance of £11, otherwise nothing. Lottery *D2* offered a 0.5 chance of £11, a 0.2 chance of £10.50, otherwise nothing. Since *D2* is worse than *D1* (by a 0.2 chance £0.50), monotonic preferences would assign a greater value to *D1* than to *D2*.[19]

Table 5: Violations of Dominance*

Patterns of valuations					
Group 1			Group 2		
Monotonic $V(D1)>V(D2)$	non-monotonic $V(D2)\geq V(D1)$	*p-value*	Monotonic $V(D1)>V(D2)$	non-monotonic $V(D2)\geq V(D1)$	*p-value*
25 (37%)	42 (63%)	0.0249*	33 (49%)	34 (51%)	0.5964

* In the above table $V(D1)>V(D2)$, for example, indicates *D1* being assigned a strictly greater value than *D2*. A *p-value*<0.05 shows 5% significance and is indicated with an asterisk.

Table 5 reports a test based on the binomial distribution of the neutral null hypothesis of random valuations against the alternative of strictly non-monotonic valuations. Table 5 shows systematic violations of monotonicity under group 1 to diminish under group 2 such that they are no longer significant. Test of difference in sample proportions based on the normal distribution of the hypothesis that violations of monotonicity are less frequent under group 2 (34/67) than under group 1 (42/67) allows rejection of the null hypothesis at 6% significance (Z=1.5662).This provides evidence that frequency-based probability learning facilitates a mitigation of the violation of this fundamental property of standard models of choice.

Diminished violations of monotonicity under group 2 suggests that the observation sequence was instrumental in allowing (some) subjects to identify the monotonicity property of their underlying preferences. In this respect individuals appear to be able to learn not to

[19] Note, however, that although *D2* is dominated by *D1* it offers two positive outcomes, whereas *D1* only offers one. Starmer and Sugden (1993) and Humphrey (1995) show that an event-splitting argument can generate indirect violations of monotonicity (manifest as transitivity violations) in pairs of pairwise choices with each lottery (similar to *D1* and *D2*) paired with some other common lottery.

violate expected utility theory. This appears to parallel the data for common consequence effects. The data for monotonicity violations, however, differ slightly from those for common consequence effects. In group 2, the latter effect is systematic and the former is not. The tests for common consequence effects and violations of monotonicity, however, also differ. An increase in expected utility maximisation in the former tests does not preclude the emergence of systematic common consequence effects within the remaining violations. An increase in expected utility maximisation in the latter tests, however, necessarily implies fewer violations of monotonicity.

6. Conclusion

The title of this paper poses a question; can individuals learn to maximise expected utility? As far as frequency-based probability learning is concerned, the evidence presented here suggests the appropriate answer to this question to be yes. This answer, however, must be qualified. The common consequence effect data suggest that whilst the proportion of expected utility maximising choices increases, it does so alongside the emergence of systematic violations of the independence axiom. 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. Systematic common consequence effects are, in fact, introduced by the learning opportunity. The assertion would be sustainable if showing individuals a series of lottery draws prior to choice does not constitute an appropriately structured learning environment. There may be grounds upon which to suspect this to be the case. 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 several reasons to suggest not.

First, frequency-based probability learning has been shown to diminish overall deviations from expected utility maximisation. Second, Estes (1976a, 1976b) and others have shown frequency-based probability learning to be effective in both introducing probabilistic biases *and* engendering accurate probability learning in other tasks. Third, 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. Gambling form guides often provided information on the outcomes of a team's last *n* fixtures (and often not, for example, on who 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 somewhat 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 not 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.

References

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

Allais, M. (1953). Le comportemant de l'Homme rationnel devant le risque, critique des postulates et axiomes de l'ecole Americaine, *Econometrica* 21, 503-546.

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 J.H. Kagel and A.E. Roth (eds.), *The Handbook of Experimental Economics* (Princeton University Press, Princeton).

Chu, Yun-Peng and Chu, Ruey-Ling. (1990). The subsidence of preference reversals in simplified and marketlike experimental settings: A note, *American Economic Review* 80, 902-11.

Conlisk, J. (1989). Three variants on the Allais example, American Economic Review 79, 392-407.

Diecidue, E. and Wakker, P. (in press). On the intuition of rank-dependent expected utility, *Journal of Risk and Uncertainty.*

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* 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 phenomenon, *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. (2000). The common consequence effect: Testing a unified explanation of recent mixed-evidence, *Journal of Economic Behavior and Organization* 41, 239-262.

Karni, E. and Safra, Z. (1987). Preference reversal and the observability of preferences by experimental methods, *Econometrica* 55, 675-685.

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. (1998). A microeconometric test of alternative stochastic theories of risky choice, Economics Research Centre Discussion Paper No. 9806, *The University of East Anglia*.

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).

Quiggin, J. (1982). A theory of anticipated utility, *Journal of Economic Behavior and Organization* 3, 323-343.

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. (1992). Testing new theories of choice under uncertainty using the common consequence effect, Review of Economic Studies 59, 813-830.

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

von Neumann, J. and Morgenstern, O. (1944), *Theory of Games and Economic Behavior*. (Princeton University Press, Princeton).

Wu, G. and R. Gonzalez, 1998, Common consequence conditions in decision making under risk, Journal of Risk and Uncertainty 16, 115-139.

Yaari, M. (1987). The dual theory of choice under risk, *Econometrica* 55, 95-115.