# Do individuals learn not to make irrational choices?

Steven J. Humphrey*
*Centre for Decision Research and Experimental Economics (CeDEx),*
*School of Economics, University of Nottingham, UK*

Choice anomalies are often 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 mitigate the inherent imprecision in preferences and yield more rational decisions than is the case in experiments devoid of learning opportunity. Contrarily, the data show that violations of expected utility theory do not always emerge in the absence of learning, but emerge systematically after learning. Learning, therefore, does not necessarily engender rational choice even when it might be expected to do so, and 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.

## 1. Introduction

The decision-making under risk literature is abound with 'anomalous' experimentally observed, systematic and apparently robust deviations from expected utility maximisation (von Neumann and Morgenstern, 1944). Recently investigators have begun to speculate that, despite such anomalies, under certain circumstances expected utility theory is a pretty good approximation of individuals' true preferences, but that some kind of learning and familiarisation with decision tasks is required before preferences 'firm-up' to settle on this genuine underlying form (e.g. Plott, 1996). This paper investigates whether conditions which enhance accurate probability learning, or at least are likely to foster a more precise formation of decision weights, by providing a particular kind of cognitively important learning opportunity, are successful in causing decisions to better reflect expected utility preferences.[1]

The motivation for the study is that despite the fact that a lack of appropriate decision-making experience and opportunity for learning are often suggested as ways in which anomalies are 'created' as purely laboratory phenomena, evidence relating directly to the influence of specific kinds of learning opportunities on risky choices is pretty thin on the ground. Moreover, there is no reason in principle why true preferences need imply expected utility maximisation. It is entirely 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). 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? Second, if probability learning affects revealed preferences, are those revealed preferences appropriately described by expected utility theory?

---

[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).

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. It is thus of considerable importance to understand exactly how particular learning opportunities shape preferences, if at all, in specific decision-making tasks. This knowledge can then be fed directly into applied studies which involve the elicitation of values from members of the public. In these types of studies, many of which are concerned with the evaluation of health risks and where, according to Loomes (1999, p.F39), behaviour may be significantly influenced by learning from different kinds of experience, ignorance of how decision-making tasks may engender different types of learning and how this affects those decisions can, quite literally, have life or death implications.

## 2. Related Literature

### *2.1. Learning in risky choice experiments*

Loomes (1999) refers to evidence generated by himself and collaborators (Loomes *et al*., 1998) and a yet-to-be-reported experiment conducted by Ken Binmore (see Loomes, 1999, p.F37, note 1) which both show choices to converge on expected utility maximisation as subjects progress through the respective experiments. This evidence supports the argument (Plott, 1996; Friedman, 1998; Binmore, 1999) that irrational behaviour may be the product of imprecise preferences encountered whilst decision-makers are learning how the task interacts with their basic underlying values, but once this learning has taken place, choices converge on expected utility maximisation. In this respect, choice anomalies are somewhat inevitably observed due to the one-shot nature of (much) laboratory experimentation and/or the use of inexperienced subjects but, with sufficient learning opportunity, will prove to be transient.[3] Additional evidence in support of the transient anomaly perspective is provided by Friedman (1998). Friedman (1998) shows that irrational behaviour on Monty Hall's three door problem

[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.

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]

There may be grounds, however, upon which to be sceptical about the case for transient anomalies, or at least the generality of the case, as it currently stands. Consider, for example, a famous experiment conducted by Slovic and Tversky (1974) which offered experimental 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, suggesting that subjects found the arguments against the independence axiom more persuasive than those in favour of it. In this case, learning certainly did not mitigate the anomaly. This evidence suggests that choice anomalies are 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.

It may be argued that the learning opportunity in Slovic and Tversky's (1974) experiment did not mitigate the anomaly because it was not appropriate to facilitate sufficiently disciplining reflection upon anomalous decisions. Presenting detailed arguments to subjects in the form of simple explanations and then allowing them to consider those arguments prior to reflecting on whether they are happy with their previous behaviour, however, appears to be an appropriate method to facilitate learning. The subtext to the suggestion that the learning opportunity was not appropriate to mitigate the anomaly is often that violators continued to violate because they were not sufficiently punished. There is indeed evidence that the kind of discipline provided by economic markets may make anomalies disappear. Chu and Chu (1990), for example, money-pumped subjects who violated the transitivity axiom of expected

[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.

utility theory until, in the light of fast-approaching experimental bankruptcy, they ceased to reverse their preferences. In the case of Slovic and Tversky's (1974) experiment, it is true that financial penalties were not imposed on violators, and so it might be argued that it is not surprising that violators continued to violate in the absence of incentives to do otherwise. But this argument would not explain why, in the similar absence of appropriate incentives, violations of the independence axiom *increased* after the learning opportunity. More generally the practical importance of observations such as Chu and Chu's (1990) in terms of how they translate to other contexts is unclear. It may, for example, be perfectly acceptable to ask cancer patients to consider how much they would be willing to pay for beneficial changes in their health state for specified periods of time, and explain ways in which they might think about the task to facilitate the arrival at genuine valuation (taking care to ensure the explanations did not prime the valuations). But ethics committees would be unlikely to approve valuation protocols which involved patients being told that the slightest inconsistency in their reported valuations would result in increased charges for their treatment, or some other penalty, until they stopped being inconsistent.

The mixed evidence regarding the impact of learning on choice behaviour substantiates the suggestion that although learning might facilitate the eradication of imprecision in preferences which stems from confusion about the mechanics of experimental tasks, there should be no presumption that the remaining confusion-free preference structure will imply expected utility maximisation (Loewenstein, 1999). The fact that different learning opportunities appear to exert different influences on choices merits further investigation of the nature of these influences.

### *2.2. Probability learning theory and the observation of lottery outcomes*

Attention is now turned to the particular type of learning opportunity with which the rest of this paper is concerned.

> **Hypothesis:** Prior to choice a sequence of resolutions of the risk in the lotteries comprising the problem will firm-up inherently imprecise subjective evaluations of stated probabilities by enhancing understanding of the meaning of those probabilities. Decisions will then reflect true preferences more accurately due to a

reduction, or eradication, of imprecision encountered whilst working-out how the task interacts with basic underlying values to yield a response.

In the experiment reported below the above hypothesis is tested as follows. Subjects are asked to value a series of lotteries by stating an amount of money such that they would be indifferent between the lottery and the valuation attached to it. Prior to making these valuations one group of subjects will observe a sequence of 10 resolutions of risk in the lottery. Another group will not see the outcome of 10 lottery draws prior to making their decision, but will be asked to make their valuations on the basis of stated information alone. Although such an observation sequence represents a simple and non-punishing (in the sense of market discipline) learning opportunity there are, from a psychological perspective, both 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 influence choices by contributing to the precise formation of decision weights from stated probabilities.

In influential work, Estes (1976a,1976b) addresses the what he sees as a deficiency in the risky choice literature; the question of how outcomes are coded in memory and how these codings affect probability assessments and subsequent choice behaviour in predictive decision-making tasks. He points out (Estes, 1976a, p.40) that the behaviour of animals, young children and adults alike does not reflect ultra-sophisticated statistical ability, but rather reflects learning on the basis of experience garnered from observing repeated situations and faith in the uniformity of nature. This yields what might be described as a basic decision-making heuristic: more frequently observed outcomes are *ceteris paribus* more likely to be future outcomes. In other words, the basis for probability learning lies in the learning of absolute frequencies of event occurrence. Estes (1976a, p.40) further posits that when repetitions of similar circumstances give different outcomes, individuals form expectations on the basis of converting absolute event frequencies into relative event frequencies.

Estes's (1976a) evidence suggests that the coding of event category frequencies in memory and their conversion to decision weights introduces a bias (due to limitations in memory capacity or inadequate training) such that the likelihood of events which occur more frequently than suggested by their true probability is overestimated in relation to that probability. Einhorn and Hogarth (1978) concur that the coding of event category frequencies rather than probability results in frequency information being more salient in memory than probability information. Frequency information, therefore, constitutes a basic building block in the process which ultimately yields a decision.[5] Imagine, for example, that an individual is asked to value a 50-50 chance of £100 or zero. Prior to making this valuation, and in full knowledge of the parameters of the lottery, they observe a sequence of ten resolutions of the uncertainty. This sequence yields seven £100 outcomes and three zero outcomes. The individual encodes the occurrence frequency of the £100 event as $f(7)$, and by forming a relative frequency converts this to a subjective probability assessment of the form $p(f(7)/10)$. Finally, the probability assessment forms the basis of the decision weight attached to the event such that $p(f(7)/10)=p(0.5)>0.5$.[6] In this instance frequency-based probability learning causes an overweighting of the £100 event in relation to its true likelihood.

The major concern of research into the role of event frequencies in probability learning has been the types of decision-making anomalies or biases this might be responsible for introducing, particularly when frequency information and probability information are inconsistent (as in the example in the previous paragraph). Einhorn and Hogarth (1978), for example, are concerned with how frequency-based probability learning can produce fallible judgements by both experts and non-experts who, nevertheless, express great confidence in those judgements. Alba and Marmorstein (1987) are concerned with the (mis)use of frequency information by marketers to shape consumer decisions. Humphrey (1999) replicates Estes's (1976a, 1976b) results in an experiment with significant financial incentives for accurate probability learning. What all of these studies have in common is that they show the manipulation of frequency information to affect decision-making behaviour such that it no longer conforms to standard models of rationality.

---

5 Frequency-based probability learning has been substantiated in many studies, some of which are discussed by Alba and Marmorstein (1987). They conclude that 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 and the overriding concern of individuals is to economise on expenditures of cognitive effort. They also agree that frequency information is learned and remembered more completely than other information.

Estes's (1976a, 1976b) original work, however, also documents extremely efficient and accurate probability learning from frequency information when the two are not in conflict. In terms of the above example, if the observation sequence yielded five £100 outcomes and five zero outcomes the decision weight attached to the £100 outcome would take the more accurate form $p(f(5)/10)=p(0.5)\approx 0.5$. If deviations from expected utility theory are mediated by the impact of probabilistic biases on the decision weighting function, possible due to a lack of understanding of the meaning of stated probability information, as is encapsulated in more popular alternatives to expected utility theory such as rank-dependent expected utility theory (e.g. Quiggin, 1982; Yaari, 1987; Tversky and Kahneman, 1992), the frequency-based probability learning argument suggests there to be good reason to expect the observation of lottery outcomes prior to choice or judgement to mitigate such anomalies. 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 frequency-based probability learning works in mitigating choice anomalies, it is quite different to beating anomalous choosers with the stick of financial (or other) punishment until they become consistent. In this respect, there may be implications for how information on probabilities is presented to members of the public or other participants in value-elicitation trials when it is not feasible or practical to 'enforce' genuine values.

In a theoretical sense, the pre-decisional experience discussed above is the type of process which may substantiate the (post learning) accuracy of expected utility theory in describing preferences. In testing this hypothesis it is necessary to recognise that preferences may initially be imprecise but, once firmed-up, are most appropriately represented by some non-expected utility theory. For example, learning may purge preferences of biases attributable to a lack of understanding of the meaning of stated probability information, but perfectly understood probability information may still be transformed into decision weights in a manner which implies deviations from expected utility theory. In this respect, if frequency-based probability learning contributes to the firming-up of decision weights, genuine properties of those weights may cause choice anomalies to persist or be exacerbated.

---

[6] A weight would also be formed for the zero outcome. Einhorn and Hogarth's (1978) evidence that 'successes' (£100) elicit more attention than 'failures' (zero) suggests the choice would be driven by the weight attached to

*1.3. The common consequence effect*

The experiment is designed around three pairs (1, 2 and 3) of lotteries as described in table 1. Each pair 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. In the diagram 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.

*Table 1*: Common consequence lotteries [a]

| **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 |

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

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

£100.

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.

*Figure 1*: Common consequence lotteries in the unit probability triangle

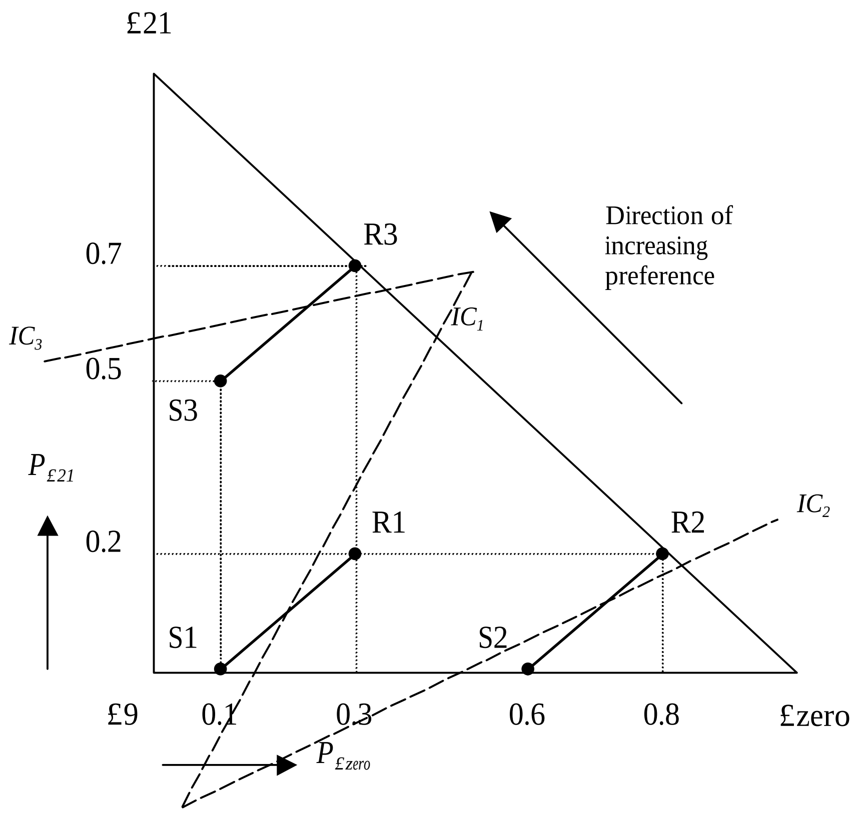


Since the only differences between all three pairwise decision problems are common to each respective lottery in the problems, expected utility maximisation demands either the riskier option to be chosen in all three problems (*R1*, *R2* and *R3*), or the safer option to be chosen in all three problems (*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.[7] Expected utility theory implies the existence of upwards-sloping, linear and parallel indifference curves in the triangle and so in

[7] 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.

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*.

It should be pointed out that although the bulk of the experimental evidence regarding common consequence problems suggests mixed-fanning to be robust, and the most popular theoretical alternatives to expected utility theory predict mixed-fanning (e.g. rank dependent expected utility theory), there is 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 show horizontal fanning-in. Since this experiment is concerned with testing common consequence effects in valuations task rather than choice tasks and, to my knowledge, there is no empirical guidance as to what pattern of preferences this may reveal, it is 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 [a]

| | **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 $V(S3) \geq V(R3)$ |

[a] 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 only values each of the six lotteries once. This means that it is perhaps unlikely to observe a single subject exhibiting consistent 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.

## 3. Experimental design

### *3.1. Valuation task*

The experiment involves 2 conditions; the first is where subjects are asked to place a money value on a simple lottery and the second where the subjects are asked to do the same, but having first observed a sequence of 10 resolutions of the risk in the lottery. Figure 2 shows an example of a valuation screen from the experimental software. In the top 'lottery' box is the lottery which the subject is being asked to value. All lotteries were expressed in terms of 10 lottery tickets and so that in figure 2 corresponds to lottery *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 would make them indifferent between the lottery and the yardstick. That is, after they had entered their valuation, were they to be offered either the lottery or the yardstick they would not mind which one they received.[8]

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

[8] 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

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.

*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.

To provide the strongest test of whether this type of pre-decisional learning opportunity causes decision weights to more precisely reflect stated probabilities, and influence the observation of common consequence effects associated with biases in the decision weighting function, the observation draws were fixed. That is, 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

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 observation sequence is 9 out of 10. This does not mean that subjects were deceived, valuations were made on the basis of complete information about the genuine probability distribution governing outcomes. Outcome information was simply provided in a manner analogous to a 'speeding-up' the law of averages because, if the frequency-based probability learning hypothesis is correct, it is this which may mitigate any violations of expected utility theory.[9] 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.

### *3.2. Incentives*

The incentive system employed in the experiment is a variation of what Tversky, Slovic and Kahneman (1990) refer to as 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 compared, and 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) to give an outcome which would be the subjects payment for participation in the experiment.[10] 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, for example, above the highest outcome offered by a lottery and often taken as evidence of irrationality or confusion, are incentive compatible within this design. Since valuations are only taken to establish a preference ordering over

---

9 Consider the possibility that subjects overweight a small probability associated with a positive outcome, as is suggested by most probability weighting models. Then, in a *genuinely random* observation sequence, if that positive outcome occurs more times than is suggested by the probability, this would serve to reinforce the subjective overweighting. Any associated decision anomalies would then be exacerbated. Although this is an equally valid reflection of what might happen in the real world to the 'fixed' observation sequence, it does not provide a pure test of whether probability learning (maybe over the long-run) on the basis of outcome observation can successfully mitigate choice anomalies. This is the question dealt with by this paper as opposed to the commonly asked question of whether unrepresentative outcome feedback can distort probability assessments to introduce deviations from rationality (e.g. Humphrey, 1999), which is already known to be true.

lotteries, subjects could perform any monotonic transformation on their valuations and still represent their true underlying preference ordering. Subjects were told, however, that because they would not know which two lotteries would be randomly selected to have their valuations compared to determine the payment lottery until after all tasks were complete, one way in which they could guarantee playing out their truly preferred lottery from the randomly selected pair, whatever that pair was, would be to consider each lottery carefully and value it genuinely.[11]

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. Then at the end of the experiment one of the lotteries is randomly selected and the reservation price is compared to a randomly generated offer. If the reservation price is below the offer the subject receives the offer and if the reservation price is equal to or above the offer the subject plays out the gamble. It has been shown, however (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, and 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. If

---

[10] 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.

[11] 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, because 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.

*V*(.) denotes the valuation attached to each lottery then horizontal fanning-out is observed if *V*(*S1*)> *V*(*R1*) and *V*(*R2*)> *V*(*S2*), and fanning-in is observed if *V*(*R1*)> *V*(*S1*) and *V*(*R2*)> *V*(*S2*).[12] 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 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 (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), respectively gives $L_1$* and $L_2$*.

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

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

$L_2^* = [\ £21,\ 0.4/n(n\text{-}1);\ £9,\ 1.8/n(n\text{-}1);\ 0,\ 1.8/n(n\text{-}1);\ Z,\ 1\text{-}(4/n(n\text{-}1))]$

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, which is central to 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.

*3.3. 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 of which 79 (59%) were male. Each session lasted for approximately one hour and average subject payment was £12.97. All subjects responded to the same set of 20 valuation tasks, with the only difference between the two groups being whether they saw the pre-valuation observation sequence or not.[13] When subjects arrived at the laboratory they were asked to sit at a computer terminal. Detailed instructions were read out by the experiment organiser and subjects responded to two practice valuation 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.

[13] 6 of the 20 lotteries in the experiment are described table 1. Another 2 lotteries investigate whether probability learning influences the rate of violation of monotonicity. These are described in the results section. The remaining lotteries were concerned with other hypotheses.

## 4. Results

*4.1. 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 and that 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 in group 2 does not influence behaviour. If the null is accepted, appropriate interpretation depends on whether significant common consequence effects are observed in both groups or in neither. If the former is true, then the implication is that either frequency-based probability learning does not occur or that it does not cause the formation of decision weights which more accurately reflect stated probabilities. If the latter is the case then it would appear that probability learning opportunities in valuation tasks are of little practical importance in that significant common consequence effects do not emerge as broadly documented in pairwise choice tasks.

The alternative hypothesis is that the observation sequence influences valuations to generate differences between the behaviour each group of subjects. Appropriate interpretation of a rejected null depends on whether that rejection is based on more or less common consequence effects being observed in group 2 than in group 1. The former would suggest that although the observation sequence influences the contribution of probability learning to the formation of decision weights, those weights need not accurately reflect stated probabilities. Decision weights might rather retain the properties postulated in alternatives to expected utility theory which accommodate the common consequence effect. The latter possibility would perhaps be good news for proponents of expected utility theory. By diminishing common consequence effects one interpretation would be that violations of expected utility theory are due to biases in the decision weighting function, possibly stemming from a lack of understanding of stated probability information, which can be ironed-out by an opportunity to learn such as that provided to group 2.

### 4.2. *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.

*Table 3*: Average lottery valuations [a]

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

[a] The 'E.Val.' column shows the expected value of each of the lotteries outlined in table 1. Outliers have been removed as follows: [b] 99.9, [c] 50.5 and 72.0, [d] 48.0, [e] 55.0 and 75.3,[f] 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. Over each of the sets of two lottery pairs compared (as described in table 1 and illustrated in figure 2), 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 group 1 subjects violated expected utility theory in the direction consistent with common consequence effects. 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), but the *p-value* column shows these differences to be insufficient to yield a rejection of the null hypothesis of consistent choice in favour of the alternative of mixed-fanning. The group 1 data alone, therefore, do not rule-out the possibility that violations of expected utility theory are due to imprecise preferences giving rise to noisy valuations due to unfamiliarity with the task and/or a lack of decision making experience in the absence of opportunities to learn.

*Table 4:* The common consequence effect [a]

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

[a] 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.

The group 2 data show patterns of valuations consistent with expected utility theory to vary between 45% and 61%. The average consistency rate of 53% is slightly higher than that of 45% under group 1.[14] The violations of expected utility theory of between 39% and 55% (averaging 47%) of valuation patterns, however, are 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 $V(R1) = V(S1)$ and that 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, and that this generally operates in the direction of introducing systematic violations of expected utility theory. Note that the emergence of significant systematic common consequence effects in group 2 occurs alongside the tendency for a greater proportion valuations patterns to be consistent with expected utility maximisation. This is an important result. First, it is contrary to the evidence discussed by Loomes (1999) showing choices to converge on expected utility maximisation as sequences of pairwise choices are answered. Second, it is contrary to Friedman's (1998) evidence. Third, the evidence directly substantiates 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. Choice anomalies, it seems, are not necessarily transient.

Perhaps the best evidence that probability learning causes irrational choices to diminish is provided by the horizontal lottery comparisons. Here 36/67 (54%) of group 1 valuation patterns violate expected utility theory, whereas under group 2 this falls to 26/67 (39%). If the alternative hypothesis was that probability learning causes such violations to diminish then 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. On the basis of a two-tailed test,

---

[14] Although there is a total 94 out of a possible 201 violations of expected utility theory under group 2 compared with 104 under group 1, there is no overall evidence to suggest that the probability learning opportunity mitigates violations of expected utility theory *per se* (i.e. irrespective of the direction of the violation). A test of difference in sample proportions over these pooled valuations yields $Z$=1.5962 which is not significant at 5%. See also the discussion of differences in the frequency of particular common consequence effects below.

however, with an alternative hypothesis based on the possibility that probability learning firms-up imprecise decision weights which might cause them to more closely reflect stated probabilities, but equally might allow the more accurate identification of non-expected utility preferences, this difference would not be significant at the 5%-level.

What seems to be happening in group 2 in relation to group 1 is an increased tendency for valuation patterns to exhibit fanning-in alongside a decreased tendency for them to exhibit fanning-out. Indeed, 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.[15] This observation is consistent with the suggestion that imprecision in genuine preferences under group 1 results in valuation patterns which deviate from expected utility maximisation distributing themselves approximately evenly between fanning-in and fanning-out, but when probability learning allows a mitigation of this inherent imprecision in group 2, and subjects are better able to identify their true preferences, a more discriminating pattern of valuations emerges. The question that this interpretation poses is exactly why might one expect this improved discrimination (on the basis of enhanced precision in decision weights and a better identification of genuine preferences) to give rise to 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 so-called *$*-bet (offering a high money prize with low probability) is assigned a higher reservation price than a *P*-bet (offering a lower money prize, but with a higher probability), but is subsequently not chosen in a direct choice between the two. This pattern of behaviour is often attributed to *response mode* effects, one feature of which is *compatibility*. The compatibility hypothesis states that money is the salient attribute of lotteries in money valuation tasks (the two are compatible) and this renders the high prize in the *$*-bet particularly influential in driving the valuation. This engenders a higher money valuation for the *$*-bet than for the *P*-bet. 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), and so preferences are reversed in favour of the *P*-bet. The importance of the compatibility hypothesis in terms of the experiment reported here is that it provides an explanation of why systematic violations of expected utility theory were not observed in group 1, but were in group 2. If common consequence effects are the product of how probabilistic biases influence the decision weighting function, then the salience of the money attribute (in relation to the probability attribute) of the lotteries in the group 1 valuation tasks may have precluded the emergence of any such probability-driven anomalies. Then, under group 2, the observation sequence may have influenced 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.[16]

### *4.3. Monotonicity*

In interpreting the above data as evidence of learning contributing to the formation of non-expected utility preferences it is necessary to recognise the possibility that alternative vehicles for learning may have a different impact on behaviour. Although the experiment reported here is not concerned directly with investigating this possibility, the evidence relating to the influence exerted by probability learning on common consequence effects can be placed in a broader context. The experiment involved two additional valuation tasks to test whether valuations satisfied monotonicity. Monotonic preferences are arguably a fundamental property of rational choice and there is a view that violations of mononicity are not systematic features of genuine preferences, whether those preferences are described by expected utility theory or alternative non-expected utility theories. In this respect, non-monotonic choices or valuations are considered as erroneous, possibly stemming from an inherent imprecision in preferences, which would presumably diminish or disappear given a suitable learning opportunity to facilitate a better identification of those preferences.

---

[15] 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.

[16] 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.

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, and lottery *D2* offered a 0.5 chance of £11, a 0.2 chance of £10.50, otherwise nothing. It is clear that 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*.[17] 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 and significant violations of monotonicity under group 1 to diminish under group 2 such that they are no longer significant.[18] This provides evidence that frequency-based probability learning does facilitate a mitigation of this fundamental violation of expected utility maximisation and its more popular alternatives.

*Table 5:* Violations of dominance [a]

| **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 | 42 | 0.0249* | 33 | 34 | 0.5964 |

[a] 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.

When taken alongside the data relating to common consequence effects, the data in table 5 substantiate the argument that there should be no presumption that learning opportunities cause a better identification of preferences which will more closely reflect those postulated by expected utility theory. The mitigation of 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, and by so doing implies that some violations of expected utility theory may be attributable to an inherent initial imprecision in those preferences. In this respect individuals appear to be able to learn not to violate expected

[17] 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.

[18] Although a 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) does not allow rejection of the null hypothesis that they are equally frequent with 5% significance ($Z$=1.5662).

utility theory. The fact that the very same instrument has the opposite impact on common consequence lotteries is open to the interpretation that it was imprecision in preferences which precluded systematic common consequence effects in group 1 but, when learning has facilitated the better identification of underlying preferences, revealed them in group 2 because they are genuine features of those preferences. In this respect, genuine preferences elicited in this experiment are monotonic but do not satisfy independence. Note that this interpretation of the above evidence draws parallels with Slovic and Tversky's (1974) observation that explaining the independence axiom increased the frequency with which it was violated.

Another important implication of the data relating to common consequence effects and monotonicity violations taken together is that they suggest violations of different principles of rational choice to be differentially affected by the same learning opportunity within the same type of decision-making task. It is of course possible that this differential impact stems from the speed with which the particular learning opportunity allows subjects to identify particular properties (such as monotonicity or independence) of their underlying preferences. Given the similarity of the lotteries which test for violations of monotonicity and common consequence effects in this experiment, however, it is unclear why this should be the case. Also, any convincing support for this argument in the context of the current data would need to explain why a longer observation sequence would initially introduce common consequence effects (after 10 observations as is the case here), but then subsequently cause their mitigation (after, say, another 10 observations).

A more sustainable interpretation of the data would appear to be simply that some choice anomalies are transient and some are not. In order to gain an understanding of whether this is the case the challenge may not only be, as Loomes (1999, p.F39) puts it, to investigate how behaviour is changed by learning from experience, but also to investigate the implications of those behavioural changes over a broad range of choice tasks and decision-making contexts. This is a potentially enormous task and, for many, it would be more convenient to be able to invoke evidence which shows experience-based learning of all types to cause convergence towards expected utility maximisation. The data reported here provide evidence that just because learning is considered to be good, it does not necessarily follow that choices made on the basis of learned preferences will be better reflect expected utility maximisation, even when every chance is afforded them to do so.

## 5. Discussion and Conclusion

The title of this paper poses a question; do individuals learn not to make irrational choices (from the perspective of expected utility theory)? As far a frequency-based probability learning is concerned, the evidence presented suggests the appropriate answer to this question to be both yes (for violations of monotonicity) and no (for common consequence effects). Yet in his investigation of anomalous behaviour in Monty Hall's three doors problem, Dan Friedman (1998, p.941) rather provocatively asserts that, "Every choice 'anomaly' can be greatly diminished or entirely eliminated in appropriately structured learning environments." In light of the current data this assertion can only be true if showing individuals a series of lottery draws prior to choice does not constitute an appropriately structured learning environment. Indeed, 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 choosers. But does this render it inappropriate? The data suggest not.

First, frequency-based probability learning mitigates deviations from expected utility maximisation manifest in violations of monotonicity. 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. For example, the time-series of stock performances is often observed prior to periodic portfolio decisions and, in independent betting form guides, information is often provided 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). Fourth, economic decisions are often made where there is not opportunity to observe and imitate more successful choosers 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, not just those which might a priori be suspected of yielding the best chance of convergence on expected utility maximisation.

In fairness, Friedman (1998, p.42) does not prescribe an ignorance of anomalies because they will eventually disappear, but he does argue the lack of need to modify, criticise or reject rational choice 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 choices conform to expected utility theory. Moreover, how would this argument respond to the data on mitigated violations of monotonicity (approximately complete and appropriate learning successfully rendering an anomaly insignificant and non-systematic) to be subsequently confronted with the data on common consequence effects (learning which introduces an anomaly is clearly incomplete and/or inappropriate)? This is not to suggest that the transient anomaly school is guilty of wanting it's cake and eating it, but rather as emphasising that just as there is the need to empirically investigate the impact of learning on choice anomalies, there is also the need to engage in a theoretical, and possibly philosophical, debate regarding the status of expected utility theory in the context of different types of learning opportunity. By providing a small amount of evidence which is potentially troublesome to the transient anomaly perspective the research reported in this paper will hopefully contribute to provoking just such a debate.

# References

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

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

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

Binmore, K. (1999), 'Why experiment in economics?, *Economic Journal*, vol. 109, pp.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*, vol. 80, pp.902-11.

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

Estes, W.K. (1976a), 'The cognitive side of probability learning', *Psychological Review*, vol. 83, pp.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*, vol. 88, pp.933-946.

Grether D. and Plott, C. (1979), 'Economic theory of choice and the preference reversal phenomenon', *American Economic Review*, vol. 69, pp.623-638.

Holt, C. (1986), 'Preference reversals and the independence axiom', *American Economic Review*, vol. 76, pp. 508-515.

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

Humphrey, S.J. (1999), 'Probability learning, event-splitting effects and the economic theory of choice', *Theory and Decision*, vol. 46, pp.51-78.

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

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

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

Loomes, G. (1999), 'Some lessons from past experiments and challenges for the future', *Economic Journal*, vol. 109, pp.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*, vol. 3, pp.323-343.

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

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

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

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

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

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

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