TESTING THE ALLAIS PARADOX IN THE CONTEXT OF HEALTH CARE

DRAFT:- PLEASE DO NOT QUOTE

Adam J. Oliver, BA, MSc, Health Economist

Office of Health Economics, 12 Whitehall, London SW1A 2DY, UK
Tel: +44 (0)20 7930 9203 x1467; Fax: +44 (0)20 7747 1419; E-mail: aoliver@abpi.org.uk

SUMMARY

There has been much experimental testing of the descriptive validity of the axioms of expected utility theory (EU) using monetary outcomes. However, such tests are relatively uncommon with respect to health care outcomes. This is unfortunate, since the standard gamble – widely considered to be the 'gold standard' for health utility assessment – is implied from the axioms of EU. In this article, the classic Allais paradox, which predicts a systematic violation of the independence axiom, is quantitatively and qualitatively tested in the context of health care outcomes. The quantitative results reported in this article demonstrate significant and systematic violations of independence in accordance with the Allais hypothesis. Moreover, the qualitative results show rational and not inconsistent reasoning behind these violations. This evidence challenges the normative status of the independence axiom and underlies the need to test alternative theories of risk and uncertainty in the context of health care to see if any of these theories can incorporate more preference patterns than EU. If an alternative theory is found to accommodate the systematic preference patterns that violate EU, efforts aimed at developing modified or new utility elicitation methods will be warranted.

Utility assessment in health care contexts is undertaken with a variety of techniques. The most common techniques are the standard gamble (SG), the time trade-off (TTO) and the visual analogue scale (VAS). Health care decisions invariably involve a degree of risk; with varying probabilities, medical interventions can have a positive, neutral or negative impact on patient health. TTO and VAS do not explicitly incorporate notions of risk, and thus lack validity in eliciting health state utilities in contexts where outcomes are uncertain. SG, however, can be implied from the axioms of expected utility theory (EU), and thus has a firm basis in the theory of risk and uncertainty. It is for this reason that SG is widely considered to be the gold standard in the measurement of health state utilities.

The axioms of EU were proposed in the 1940s [1]. However, since the 1950s, the descriptive validity of EU has been increasingly discredited through the empirical testing of its axioms [2-10]. This empirical testing has been undertaken in the context of monetary outcomes, and there have been few direct attempts at testing the descriptive validity of EU in the context of health care [11]. The independence axiom has been subject to most of the criticism of EU. Independence proposes that when an individual processes a choice context, the level of any specific outcome will not influence the individual's perception of the probability associated with that outcome and vice versa. The basic concept of independence can best be illustrated with the aid of Table 1.

[Insert Table 1]

In Table 1, gambles A, B, A' and B' are presented in a deconstructed format. It can be seen from Table 1 that, with a probability of 0.89, gambles A and B share a common outcome of $1million and gambles A' and B' share a common outcome of $0. In all other respects A is identical to A' and B is identical to B'. Consider the case where an individual is asked to choose between both A and B, and A' and B'. Under independence, the individual will disregard any common outcomes within choice contexts. Therefore an individual who prefers A (B) should prefer A' (B').

In a major criticism of EU, Allais argued that under certain conditions individuals will systematically violate independence [2]. In a constructed format, A, B, A' and B' in Table 1 can be presented as:

A: $1m for certain
B: (0.89, $1m; 0.10, $5m; 0.01, $0)
A': (0.89, $0; 0.11, $1m)
B': (0.90, $0; 0.10, $5m)

Allais argued that when individuals are faced with choices between A and B, and A' and B', in the above format, many individuals will display a preference for A and B', apparently demonstrating a violation of the independence axiom. Allais' proposition is known as the Allais paradox, and has been supported empirically in many subsequent analyses [3-6, 12-13].

To illustrate how independence between outcomes and probabilities is implicit in SG measurements, consider Figure 1.

[Insert Figure 1]

In health economics, the SG method conventionally adopts the variable-probability-equivalent method [11]. For example, in Figure 1a individuals would be asked to choose the probability, p, for which they are indifferent between the gamble of living for 40 years in full health otherwise immediate death (0yrs), and the certainty of living for 20 years in full health.

Assume that in the example in Figure 1, 40 years in full health represents the best possible outcome, and immediate death represents the worst possible outcome. These outcomes can therefore be assigned a utility score of 1 and 0, respectively. In Figure 1a, assume that an individual specifies a probability of 0.9 for indifference between the gamble and the certainty. The SG utility of living for 20 years in full health can then be estimated as 0.9.u(40yrs) + 0.2.u(0yrs), where u(.) represents the utility scale. By substituting in the utility scores of 1 and 0 for best and worst outcomes, an SG utility of 0.9 is derived. Similarly, if that same individual specifies a probability of 0.6 in Figure 1b, an SG utility score of 0.6 is estimated for 10 years in full health. In Figure 1c, under independence, it is possible to substitute in the individual's previously attained SG utility scores for 10 years and 20 years in full health. If the individual is acting in accordance with the axioms of EU, a probability consistent with 0.6 = p.0.9 (i.e. p = 0.66) would be stated. In the absence of independence between the outcomes and probabilities, the applicability of an SG measurement for a specific outcome to other SG constructs cannot be assumed. In such circumstances, cardinality is compromised.

The objective of this article is to estimate directly the extent to which individuals comply with the independence axiom in Allais-type health care contexts. This analysis represents a rare attempt at testing the descriptive performance of EU in the context of health care outcomes.

METHODS

Experiment design

The results reported in this article are part of a much wider study in which each subject was presented with 20 health care contexts. In an attempt to obtain and maintain the full understanding and interest of the subjects, the experiment was designed with the aim of presenting the contexts with as much clarity as possible. With this in mind, the options in each context were presented in pie chart format. An example of the pie chart format in a typical context is presented in Figure 2.

[Insert Figure 2]

The subjects were asked to imagine that they have an illness from which, without treatment, they would die almost immediately, and that their doctor tells them that there are 2 alternative treatments for their illness. The construct of the 2 pies in each context illustrates the chances of certain outcomes from the 2 available treatments. For example, in the context of Figure 2, the subjects were informed that if they took treatment (a) they would have a 30% chance of living for 4 years in full health and then death, a 30% chance of living for 6 years in full health and then death, and a 40% chance of living for 10 years in full health and then death. Alternatively, if they took treatment (b) they would definitely live for 8 years in full health and then die.

For each context it was also stated that their doctor informs the subject that there is a new treatment being developed for their illness which, if taken, would definitely give the subject 30 years in full health followed by death. The subjects were told, however, that the new treatment would not be available for them.

Each subject was then asked to rate the treatments (a) and (b) on a scale, which had the outcome of the unavailable treatment (30 years in full health followed by death) at the top of the scale, and the outcome of no treatment (immediate death) at the bottom of the scale. Indifference could be indicated by rating both treatment options at the same point on the scale. It was emphasised that this was an ordinal rating; even if it had been considered possible, no effort was made to measure cardinality in this experiment. The subjects were reminded that there are no right or wrong answers to any of the questions.

The reason why the subjects were asked to rate the alternatives in each context on a scale rather than to directly choose between treatments (a) and (b) was because the experiment contained many contexts in which the 2 treatment alternatives shared a common outcome. If the subjects had been asked to directly choose between treatment options, there would have been a danger that they would have learned that many contexts contain options with a common outcome and immediately cancel any consideration of the common outcome. Such an occurrence would represent an immediate focus on a particular outcome and may introduce a cognitive process that would hardly, if ever, be induced in real world settings. It was thus considered important to reduce the possibility of an unnatural and immediate focus upon any particular outcome. Therefore, it was thought that by asking the subjects to rank each treatment on a scale that was also marked with the outcomes of living for 30 years and immediate death, they would be more likely to consider all outcomes in each treatment option before indicating a preference.

The analysis in this article relies on the 2 contexts presented in Figure 3, which test the Allais hypothesis in the context of health care outcomes.

[Insert Figure 3]

The order of the 20 contexts in the full questionnaire were randomised across subjects. Therefore, the order in which the 2 contexts illustrated in Figure 3 were presented varied across subjects. In context 1, treatments (a) and (b) share the common outcome of an 89% chance of living for 12 years in full health then death. Similarly, in context 2, treatments (a') and (b') share the common outcome of an 89% chance of immediate death. According to the independence axiom, subjects, after (carefully) considering the 2 respective sets of treatment options, should disregard the common outcomes. When contexts 1 and 2 are placed together as in Figure 3, it is easy to observe that if the respective common outcomes are disregarded, the 2 contexts are identical. Hence, according to independence, individuals should rate (a) higher than (b) if they rate (a') higher than (b'), and vice versa.

In anticipation that some subjects might find treatment options that include 2 different percentage chances of the same outcome strange (as in option (b) and (a') in Figure 3), all subjects were given 3 practice questions before they answered the full questionnaire. One of the practice questions was identical to context 2 in Figure 3 (which also appeared in the full questionnaire). In order to ensure that subjects fully understood the contexts placed before them, they were allowed to ask questions during the practice session. All subjects were informed that though the presentation of some options may appear strange, they should answer as best they could according to their own individual preferences. The practice questions also served as a test of consistency. All subjects stated that they understood and were happy with the presentation of the contexts before they began the full questionnaire, which they were required to answer without conferring and without asking any questions. During the experiment, subjects were free to return to previous questions in order to revise their answers.

After each context, the subjects were asked to briefly write down in their own words the reason for their (ordinal) rating of the 2 treatment options. This was undertaken to attain some qualitative evidence of the cognitive processes the subjects employed in reaching their decisions.

Subject population

Subjects were recruited in June 1999 on a strictly voluntary basis from the staff of a large health care-related organisation situated in London. A copy of the letter inviting the subjects to participate in the experiment is given in the Appendix.

In experiments that use monetary outcomes, most experimental economists believe that in order to introduce the appropriate incentives, it is necessary to incorporate some form of payment mechanism in the experiment. This is usually done by either paying the subjects to participate, allowing each subject to play out a context for real, or a combination of both. However, in the experiment reported in this article, it was hoped that by informing the subjects that their participation may, in some small way, be a contribution towards health care science, and by offering a presentation of the results, the necessary incentives for them to express their preferences to the best of their abilities would be present. Also, it was deemed that offering a payment to answer health care-related contexts might well introduce inappropriate incentives. Therefore, no payment was offered to the subjects in this experiment.

Thirty-eight people agreed to participate in the experiment. In order to obtain agreement to undertake the experiment, the author had to guarantee that the subjects would remain strictly anonymous. Therefore, scant information on the construct of the subject population can be given. However, it can be stated that subjects were recruited from all grades within the organisation (i.e. from the general office staff to the directors) and that 55% (21/38) were women. Subjects undertook the experiment in 9 groups of between 2 and 6 people during July and August 1999.

RESULTS

Quantitative results

In presenting the results, it is necessary to refer to contexts 1 and 2 in Figure 3. Table 2 shows the preference patterns observed after subjects had rated treatment (a) against (b), and (a') against (b').

[Insert Table 2]

In Table 2, preference pattern (a) (a') refers to a situation where treatment (a) is rated higher than (b) in context 1 of Figure 3, and (a') is rated higher than (b') in context 2. All other preference patterns can be read similarly. The symbol (I) refers to a context where the subject is indifferent between the 2 treatment options. Of the observed preference patterns, (a) (a') and (b) (b') are consistent with the independence axiom. Preference patterns (a) (b') and (b) (a') represent strict violations of independence, and the preference patterns (a) (I') and (I) (a') represent weak violations of this axiom.

Twenty of the 38 subjects (52.6%) violated independence. Of the 18 subjects whose preferences were consistent with independence, 14 (77.8%) exhibited preference pattern (a) (a'). Therefore, within this experimental design, those that were consistent with independence were more likely to exhibit risk seeking behaviour, since most appeared to accept a 1% chance (or a 1% increase in the chance) of immediate death in exchange for a 10% chance of living for 18 years in full health.

If violations of independence are to be accepted as errors, it is expected that violations of this axiom would be (roughly) equally split in both directions. That is, subjects would be expected to exhibit patterns (a) (b') and (a') (b) in (roughly) equal numbers. However, in Table 2, this is clearly not the case. Fourteen of the 20 subjects (70%) who violated independence (14 of the 17 (82.4%) who demonstrated strict violations) exhibited preference pattern (b) (a'). This preference pattern is consistent with choices that have been observed in the context of monetary outcomes (and accords with the Allais hypothesis). This preference pattern implies that people will often be risk averse when presented with treatment options where one option involves a small chance of death and the alternative option presents no chance of death (and when the former option also presents a chance of a better outcome than can be realised with the alternative treatment), but that they will be risk seeking when both options present a

chance of death, with the probability of death being slightly higher in the former option than in the latter. Ultimately, the preponderance of subjects (amongst those who violated independence) exhibiting preference pattern (b) (a') demonstrates a systematic violation of the independence axiom that cannot be explained away as error and that should not be ignored.

Another important point to note is that in many of the monetary contexts that have been used to test for the Allais paradox, the options have been presented in a constructed form, as explained at the beginning of this article. In these constructed forms, it is not immediately obvious that the alternatives within a context share a common outcome, which may be used as an explanation as to why violations of independence have been observed. However, in the current experiment, the independence axiom was systematically violated even though the presentational design made it obvious that (a) and (b) (and (a') and (b')) share a common outcome. This experiment, therefore, strengthens the argument that violations of independence are deliberate.

As mentioned earlier, partly as a test of consistency, one of the practice questions given to the subjects before they completed the main questionnaire was identical to context 2 in Figure 3. Twenty-six subjects (68.4%) gave consistent answers for these 2 questions. There are many possible explanations for this seemingly low level of consistency. For example, many of the subjects may have been almost indifferent between the treatment options in this context, their preferences may have be highly transient, and some of them could have made errors, particularly in the practice questions which they undertook very early in the experiment.

Qualitative results

As noted in the methods section, in an effort to obtain some qualitative explanations for the quantitative results, the subjects were asked to write down the reasons for their responses. The subjects were informed at the beginning of the experiment that they were not compulsorily required to give an explanation for their answers; that is, if they found it difficult to write in their own words the reason underlying their responses, or if they simply did not want to give an explanation, then they could ignore the qualitative section. This was done in order to minimise the possibility of exhausting the subjects' attention span.

The qualitative evidence was surprisingly well defined, and the reason for all responses could be identified from the following explanations:

Explanation i:	The 10% chance of 18 years outweighs the 1% chance of death.
Explanation ii:	Difference between 10% and 11% negligible; go for 18 years.
Explanation iii:	Overall expectation of life years is greater under (a') than (b').
Explanation iv:	Avoid the possibility of immediate death at all costs.
Explanation v:	The 1% chance of death outweighs the 10% chance of 18 years.
Explanation vi:	The certainty of 12 years is overriding.
Explanation vii:	No explanation given.
Explanation viii:	Explanation is inconsistent with rating.

Explanation ix: No preference between alternatives.

A summary of the explanations that each subject gave for their ratings in contexts 1 and 2 is given in Table 3.

[Insert Table 3]

In Table 3, subject numbers have been assigned according to preference patterns; i.e. the subjects who conformed with independence (preference patterns (a) (a') and (b) (b')) have been listed from numbers 1 to 18, and the independence violating subjects are numbered from 19 to 38. The column headed 'Explanation' accords with the explanations listed above, and indicates the reason that each subject gave for their respective ratings in contexts 1 and 2 of Figure 3.

In terms of the particular preference patterns, 9 of the 14 subjects with the dominant independence conforming preference pattern of (a) (a') gave a common reason for preferring (a). The reason was that, after apparently cancelling the common outcome across (a) and (b), they considered the 1% chance of immediate death worth taking for the 10% chance of living for 18 years in full health, as compared to the 11% chance of living for 12 years in full health in (b). The 1% chance of death was therefore not significant enough to prevent these subjects from being risk seekers in this context. The remaining 5 subjects with this preference pattern did not give an explanation for preferring (a).

For the 12 subjects with this preference pattern who gave an explanation for preferring (a'), 6 gave the same explanation as that given for (a). However, only 4 subjects (subjects 3, 4, 6 and 7) gave the same explanation for both (a) and (a'). Five subjects preferred (a') because they felt that the difference in the chance of a successful outcome between (a') and (b') was so small (10% versus 11%) that they simply preferred the option with the highest outcome; i.e. the 10% chance of 18 years in (a'). The implication of this is that these subjects engaged in apparently risk seeking behaviour, though after approximating the probability of a successful outcome as identical across options, the subjects themselves are unlikely to have perceived their preference as relatively risky. All 5 of these subjects had given explanation i for (a); therefore, whilst the motivation for their preference for (a) was given in terms of the probabilities, their motivation for preferring (a') was given in terms of the outcomes. One subject (subject 10) preferred (a') because the overall expectation of life years in (a') is greater than that for (b').

Only 2 of the 4 subjects with the less common independence conforming preference pattern of (b) (b') gave any explanation for their responses. These subjects (subjects 17 and 18) preferred both (b) and (b') due to a strong aversion to a (higher) chance of experiencing immediate death.

Of the 10 subjects who conformed with the Allais paradox preference pattern of (b) (a'), 4 preferred (b), because they did not think it worth risking the 1% chance of immediate death in (a) for the 10% chance of living for 18 years in full health

(explanation v), and 6 preferred (b) because they laid emphasis on the certainty of living for 12 years in full health (explanation vi). Explanations v and vi do seem related, and both could in the absence of any other information be categorised under the more generic explanation of a focus on avoiding immediate death (explanation iv). However, explanation v implies that there may be a better outcome than 18 years in full health for which subjects would be willing to take the 1% chance of immediate death, whereas explanation vi implies a higher degree of risk aversion in that subjects giving this explanation focus explicitly on the certain outcome.

For preference pattern (b) (a'), 10 subjects (though not the same 10 subjects who gave an explanation for preferring (b)) also gave an explanation for preferring (a'). Four of these subjects preferred (a') because, after cancelling the common outcome, they thought the 1% chance of immediate death worth the 10% chance of living for 18 years in full health. One of these subjects (subject 19) had, however, given the opposite explanation of the 1% chance of death not being worth the risk for a 10% chance of 18 years when rating (b) higher than (a) in context 1. Subject 19's preferences therefore seem inconsistent across contexts. Two of the 4 subjects (subjects 23 and 25) had focussed on the certainty of 12 years in context 1. In context 2, this certainty had been removed, and it does not therefore seem inconsistent that they should wish to rate the more risky option higher in this context since a different cognitive process was now governing their decision. It does demonstrate, however, that these individuals would not necessarily wish to minimise the chance of death, as may have been interpreted from their responses in context 1.

The remaining 6 subjects preferred (a') due to the difference in the chance of a successful outcome between (a') and (b') being so small that they might as well go for the option with the highest outcome (explanation ii). As stated above, these subjects gave 1 of 2 explanations for preferring (b) in context 1: they either considered that the 1% chance of death outweighs the 10% chance of 18 years, or that the certainty of 12 years is overriding. These subjects therefore appeared to ultimately focus on the probabilities in reaching their decision in this context, which may have caused them to be relatively risk averse; however, the framing of context 2 may have led them to ultimately focus on the outcomes in this context, offering a possible explanation for them demonstrating relative risk seeking behaviour in preferring (a').

In general, the explanation behind the systematic violations of independence that emerges from this qualitative evidence is that subjects often use reasonable and rational thought processes that induce them to be risk averse when there is a low probability of death over treatment options, and use different, but equally reasonable, rational and not necessarily inconsistent thought processes that induce them to be risk seeking when the probability of death is high in both contexts.

The qualitative evidence is less illuminating for the relatively minor independence violating preference patterns. For the 3 subjects who demonstrated preference pattern (a) (b'), 2 gave an explanation for their responses. Both preferred (a) because, after cancelling the common outcome, they considered the 1% chance of immediate death a risk worth taking for the 10% chance of living 18 years in full health. Incidentally, for those offering an explanation for their answers, this explanation was given by all subjects who preferred (a) throughout the experiment. For those demonstrating

preference pattern (a) (b'), 1 subject preferred (b') over (a') due to the lower chance of immediate death (an explanation that was inconsistent with their explanation for preferring (a)), and 1 subject wrote that their preference (for (b')) was due to the 10% chance of living for 18 years in full health outweighing the 1% chance of immediate death (an explanation which implied that they had rated (b') higher than (a') by mistake). The 2 subjects with preference pattern (a) (I') preferred (a) for the reason stated above. One of these subjects gave an explanation for rating (a') and (b') at the same level, which was simply to say that they had no preference between these alternatives. The single subject who demonstrated preference pattern (I) (a) gave no explanation for their preferences.

From the column headed 'Practice response and explanation' in Table 3, the response and explanation for the practice context by those subjects who gave a different answer to the practice context than they did to the identical context 2 can be observed. In this column, the superscript 'p' simply denotes the answer to the practice context. Over these 2 contexts, quantitative consistency required the subject to respond (a') if (a^p), (b') if (b^p) or (I') if (I^p).

Eight (67%) of the 12 subjects who appeared to give inconsistent answers demonstrated preferences in the direction of (b^p) (a'). Therefore, in this experiment, there was a slight tendency for the subjects who demonstrated inconsistent preferences to do so because they began the experiment in a relatively risk averse manner, but became more risk seeking by the time they reached context 2. However, concerning this issue, a sample of 12 is too small to reach definite conclusions.

The reason for an apparently inconsistent answer could only be assessed for subjects who gave a reason for their preferences in *both* the practice context *and* in context 2. Six of the 12 subjects (subjects 9, 19, 26, 27 33 and 37) gave this information. Two subjects (subjects 9 and 26) stated in the practice context that the extra 1% risk of immediate death (from 89% to 90%) was not worth taking for the possibility of extra life years. However, in context 2, they rationalised that an 89% and a 90% chance of immediate death are nearly identical and almost certain to occur with both options, and so went for the relatively small chance of winning 18 years of full health. Therefore, these subjects used 2 perfectly reasonable rules of thumb to reach 2 different conclusions when presented with the same context twice in the space of 1 hour.

Subject 19 also stated in the practice context that the extra 1% risk of immediate death was not worth taking for the possibility of extra life years, but then reversed this conclusion in context 2. This may indicate that this subject's preferences are highly transient. Similarly, subject 27 indicated that they preferred to avoid the higher chance of immediate death in the practice context, but then rationalised so as to accept this higher chance in context 2.

Subject 33, though having appeared to have demonstrated inconsistent preferences in the quantitative analysis, was found to have given an ordinal rating of options inconsistent with their underlying explanation in context 2 (as described earlier). Correcting for this in the qualitative analysis thus results in this subject preferring (a^p) and (a'). Similarly, and finally, subject 37 stated that they had no preference over

options in the practice question, but proceeded to rate (a^p) higher than (b^p). Correcting for this would remove the inconsistency in this subject's responses.

Overall, the qualitative analysis of the consistency check, though limited by sample size, seems to suggest that quantitatively observed inconsistency is a mixture of reasonable and rational thought processes, and error.

DISCUSSION

To summarise the quantitative and qualitative results presented above, 20 (52.6%) of the 38 subjects who participated in the experiment violated independence, with 14 (36.8%) of these violating, in accordance with the Allais hypothesis, in the direction of (b) (a'). The violations were therefore systematic and occurred even though the presentation of a common outcome across options was made explicit. Twenty-eight (73.6%) of the 38 subjects followed preference pattern (a) (a') or (b) (a') and, thus, within the context of this experiment, an acceptable theory of decision making behaviour under risky circumstances should (at least) be able to incorporate these 2 preference patterns. For subjects who followed preference pattern (a) (a'), the sets of explanations they gave for their preferences were as follows:

The reason for preferring (a) over (b): The risk of the 1% chance of death is worth taking for the 10% chance of living for 18 years in full health (explanation i);
The reason for preferring (a') over (b'): The risk of the 1% chance of death is worth taking for the 10% chance of living for 18 years in full health (explanation i).

Or:

The reason for preferring (a) over (b): The risk of the 1% chance of death is worth taking for the 10% chance of living for 18 years in full health (explanation i);
The reason for preferring (a') over (b'): The difference between a 10% and an 11% chance of a positive outcome is negligible, so go for 18 years (explanation ii).

Both of these 2 sets of explanations gave preferences that accord with the independence axiom. Subjects who gave the first set of explanations gave identical, and therefore perfectly consistent, explanations for preferring (a) and (a'); subjects who gave the second set of explanations used different thought processes for preferring (a) and (a'), both of which appear to be perfectly acceptable ways of rationalising preferences, and reached conclusions that are consistent with EU. Subjects who demonstrated preference pattern (b) (a'), gave the following sets of explanations for their preferences:

The reason for preferring (b) over (a): The risk of the 1% chance of death is not worth taking for the 10% chance of living for 18 years (explanation v);
The reason for preferring (a') over (b'): The risk of the 1% chance of death is worth taking for the 10% chance of living for 18 years in full health (explanation i).

Or:

The reason for preferring (b) over (a): The risk of the 1% chance of death is not worth taking for the 10% chance of living for 18 years (explanation v);
The reason for preferring (a') over (b'): The difference between a 10% and an 11% chance of a positive outcome is negligible, so go for 18 years (explanation ii).

Or:

The reason for preferring (b) over (a): The certainty of a positive outcome is overriding (explanation vi);
The reason for preferring (a') over (b'): The risk of the 1% chance of death is worth taking for the 10% chance of living for 18 years in full health (explanation i).

Or:

The reason for preferring (b) over (a): The certainty of a positive outcome is overriding (explanation vi);
The reason for preferring (a') over (b'): The difference between a 10% and an 11% chance of a positive outcome is negligible, so go for 18 years (explanation ii).

The first set of explanations for this preference pattern are clearly inconsistent with one another, but only 1 subject gave these explanations for their preferences. In the remaining sets of explanations, a different cognitive process is used to reach a preference for (b) than for (a'). Therefore, it cannot be concluded that these cognitive processes are inconsistent; they are merely different. In fact, what appeared to happen in this experiment is that when the subjects were asked to rate an option that gives a positive outcome with certainty against an option that has a small probability of immediate death (as in context 1 of Figure 3), they often focused upon the small probability of immediate death, which may have induced risk averse behaviour. However, with an identical difference in the percentage chance of immediate death between the 2 options, but with a large chance of death in both options (as in context 2 of Figure 3), the subjects often appeared to attach less weight to the probability of death and were more likely to base their preference on the option that gave the best possible outcome, which resulted in (seemingly) risk seeking behaviour.

If people demonstrate preference pattern (b) (a') they have violated independence and, hence, EU. The above explanations for demonstrating (b) (a') thus offer some qualitative evidence for descriptive violations of independence. However, the question also arises as to whether the above explanations challenge the normative status of the independence axiom. That is, though many subjects *did* violate independence, do their explanations suggest that they *should not have* violated independence?

As stated earlier, the first set of explanations for demonstrating (b) (a') given above appear to be inconsistent with one another; it can therefore be concluded that the subject who gave these explanations should not have rationalised their preferences in this manner, which preserves the normative status of independence. However, in the 3 remaining sets of explanations, the subjects appeared to give perfectly reasonable, and not inconsistent, explanations for preferring (b) and (a'). Their explanations were not inconsistent because they used different thought processes for preferring (b) and (a'). Had they used the same thought process in reaching their preferences for (b) and (a'),

or if it could be concluded with confidence that they *should* demonstrate preferences that are consistent with independence after employing different thought processes, then it could be stated indisputably that the normative status of independence remains in tact. However, within the context of this experiment, it cannot be concluded with confidence that subjects *should not* prefer to be risk averse if they can avoid any possibility of immediate death *and* risk seeking if the possibility of immediate death is large and unavoidable across treatment options. In short, the normative status of the independence axiom is challenged.

Should health economists be concerned with significant and systematic violations of the independence axiom? As explained earlier in the context of Figure 1, systematic violations of the independence axiom imply that cardinal utilities cannot necessarily be accurately elicited from the SG method. If cardinal utilities are not being generated from the SG method, then summing the values derived from the SG across individuals, health states and treatment options may give misleading results. For example, consider the simple example in Table 4. Assume that a policy maker is faced with 3 treatments, A, B, and C, which cost the same and are available for use against a particular disease. Also assume that society contains 3 individuals, X, Y and Z. The decision maker has the available resources to fund only 1 of the treatments and, as an aid to decision, measures the worth of each treatment by summing SG generated values obtained from the 3 individuals in society (note that this example can be extended to cases where the value of health states rather than treatment options is measured). The non-parenthesised figures in Table 4 refer to the (hypothetical and ordinal) SG values, and the slightly different figures in parenthesis refer to the cardinal utility that each individual attaches to the treatment options, assuming for the moment that cardinal utility is measurable (note that the SG values derived from individuals who perfectly conform with EU would be cardinal utilities).

[Insert Table 4]

After summing the SG values across X, Y and Z, the 'Total' column in Table 4 illustrates that treatment A is the best option. However, treatment B generates the most cardinal utility. Two points should be noted. Firstly, the differences in total SG values and total cardinal utility between A and B are small, and it could therefore be argued that with the appropriate application of confidence intervals and sensitivity analysis the difference between using SG values and cardinal utility (if measurable) is negligible. However, the example in Table 4 involves the valuation of only 3 options over 3 individuals. In a more realistic setting, involving the valuation of many more treatment options (and health states) over hundreds of individuals, the scope for generating misleading results by using a method that is based on assumptions that people significantly and systematically violate, is much greater. Secondly, it may be argued that there is no point in arguing about 'true' cardinal utility unless it is measurable, and that the SG values are a good enough, and indeed the best, available approximation. However, can this argument be accepted if individuals do indeed systematically violate (at least 1 of) the axioms of the theory on which the method is based? In these circumstances, can it really be concluded that the SG adequately

reflects people's preferences? Efforts should clearly be made to develop value elicitation methods that are based on theories that are not systematically violated.

The modification of the SG method to align with rank-dependent utility theory (RDU) has been suggested [11]. RDU generalises EU by weakening the independence axiom. An essential feature of RDU is that it assumes that people will apply decision weights to probabilities. For example, when people are faced with a treatment option that involves, say, a 1% chance of death, RDU specifies that they may apply a weight to this probability and perceive the chance of death as being greater than 1%. It is important to note, however, that RDU significantly deviates from EU only for gambles (or treatments) that involve small or large probabilities of outcomes. If RDU reflects preferences better than EU, an argument can be made to transform the probabilities derived from the SG method for the purpose of eliciting utility measurements.

The axioms of RDU have been tested under risk and uncertainty in contexts using monetary outcomes, but the results have been negative [14-15]. The theory has yet to be tested in the context of health care and, therefore, it is not yet known if an RDU-modified SG method, even under the assumption that an appropriate probability weighting function can be derived, would provide improved approximations of health-related cardinal utility. However, the systematic and rational violations of independence reported in this article places a question mark against the validity of the EU-based SG method. Therefore, the quantitative and qualitative testing of the underlying axioms of RDU and other generalisations of and alternatives to EU in the context of health care, and the development of utility elicitation methods based on theories that are found to better incorporate preference patterns, provides an important research agenda for experimental and health economists to follow.

ACKNOWLEGEMENTS

The author is grateful to the volunteers from the Association of the British Pharmaceutical Industry who agreed to participate in this experiment. The author is also grateful for support in the form of ESRC award number R00429834596.

APPENDIX

Many researchers are currently trying to study the way people make choices regarding health care. The purpose of them doing this is so that they can advise policy makers on how to make decisions that better reflect the preferences of the population.

I have designed an experiment that will help me to analyse the choices that people make when they are faced with health care-related choices. The questions are purely hypothetical and therefore do **not** require you to disclose **any** information about your own health. The experiment itself will consist of about 20 similar questions, and should take 45-60 minutes to complete. There are no right or wrong answers to any of the questions – I am merely interested in people's preferences, which will be different for different people.

I am hoping to conduct the experiment on several groups of 5-10 people. For each group, I will explain the design of the experiment and go through some example questions, and will then ask each of you to complete the experiment. I am therefore asking you if you would be willing to undertake the experiment. **All of your answers will of course remain strictly confidential.** If you are willing to undertake the experiment, please e-mail me so that I can sort out some times.

After the experiment, I shall invite you to attend a short presentation given by me on the general pattern of preferences observed in the experiment.

Thank you for your co-operation.

REFERENCES

1. von Neumann, J. and Morgenstern, O. *Theory of games and economic behavior.* Princeton: Princeton University Press, 1944.
2. Allais, M. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'ecole Américaine. *Econometrica* 1953; **21**: 503-546.
3. Conlisk, J. Three variants on the Allais example. *American Economic Review* 1989; **79**: 392-407.
4. MacCrimmon, K.R. and Larsson, S. Utility theory: axioms versus paradoxes. In: Allais M, Hagen O, editors. *Expected utility hypotheses and the Allais paradox.* Dordrecht: Reidel, 1979: 27-145.
5. Slovic, P. and Tversky, A. Who accepts Savage's axiom? Behavioral Science 1974; **19**: 368-373.
6. Kahneman, D. and Tversky, A. Prospect theory: an analysis of decision under risk. *Econometrica* 1979; **47**: 263-291.
7. Machina, M.J. Choice under uncertainty: problems solved and unsolved. *Economic Perspectives* 1987; **1**: 121-154.
8. Tversky, A. and Kahneman, D. Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 1992; **5**: 297-323.
9. Camerer, C. and Weber, M. Recent developments in modelling preferences: uncertainty and ambiguity. *Journal of Risk and Uncertainty* 1992; **5**: 325-370.
10. Camerer, C. and Ho, T. Violations of the betweenness axiom and non-linearity in probability. *Journal of Risk and Uncertainty* 1994; **8**: 167-196.
11. Wakker, P. and Stiggelbout, A. Explaining distortions in utility elicitation through the rank-dependent model for risky choices. *Medical Decision Making* 1995; **15**: 180-186.
12. Moskowitz, H. Effects of problem representation and feedback on rational behavior in Allais and Morlat-type problems. *Decision Sciences* 1974; **5**: 225-242.
13. Morrison, D.G. On the consistency of preferences in Allais' paradox. *Behavioral Science* 1967; **12**: 373-383.
14. Wakker, P.P, Erev, I. And Weber, E. Comonotonic independence: the critical test between classical and rank-dependent utility theories. *Journal of Risk and Uncertainty* 1995; **9**: 195-230.
15. Fennema, H. and Wakker, P. A test of rank-dependent utility in the context of ambiguity. *Journal of Risk and Uncertainty* 1996; **13**: 19-35.

Table 1: Demonstrating the concept of independence

Probability	0.89	0.10	0.01
Gamble A	$1m	$1m	$1m
Gamble B	$1m	$5m	$0
Gamble A'	$0	$1m	$1m
Gamble B'	$0	$5m	$0

Table 2: Quantitative results from the test of the Allais paradox

Preference pattern	Number of subjects
(a) (a')	14 (36.8%)
(b) (b')	4 (10.5%)
(a) (b')	3 (7.9%)
(b) (a')	14 (36.8%)
(a) (I')	2 (5.3%)
(I) (a')	1 (2.6%)

Table 3: Qualitative results from the test of the Allais paradox

Preference pattern	Subject	Explanation			Practice response and explanation	
(a) (a')						
	1	i	ii			
	2	i	ii			
	3	i	i			
	4	i	i		(b^p)	vii
	5	i	ii			
	6	i		i		
	7	i		i		
	8	i	ii			
	9	i	ii		(b^p)	v
	10	vii	iii			
	11	vii	vii			
	12	vii	i		(b^p)	vii
	13	vii	vii			
	14	vii	i			
(b) (b')						
	15	vii	vii			
	16	vii	vii			
	17	iv	iv			
	18	iv	iv		(a^p)	vii
(b) (a')						
	19	v	i		(b^p)	v
	20	v	ii			
	21	v	ii			
	22	v	vii			
	23	vi	i			
	24	vi	ii			
	25	vi	i			
	26	vi	ii		(b^p)	v
	27	vi	ii		(b^p)	iv
	28	vi	vii			
	29	vii	vii		(b^p)	vii
	30	vii	vii		(b^p)	vii
	31	vii	ii			
	32	vii	i			
(a) (b')						
	33	i	viii		(a^p)	i
	34	i	iv			
	35	vii	vii			
(a) (I')						
	36	i	vii			
	37	i	ix		(a^p)	viii
(I) (a)						
	38	vii	vii		(I^p)	vii

Table 4: The implication of systematic violations of the independence axiom

Treatment	Individual X	Individual Y	Individual Z	Total
A	0.8 (0.7)	0.9 (0.8)	0.6 (0.5)	2.3 (2.0)
B	0.6 (0.5)	0.8 (0.7)	0.8 (0.9)	2.2 (2.1)
C	0.3 (0.2)	0.3 (0.2)	0.1 (0.2)	0.7 (0.6)

Figure 1: Independence in the standard gamble

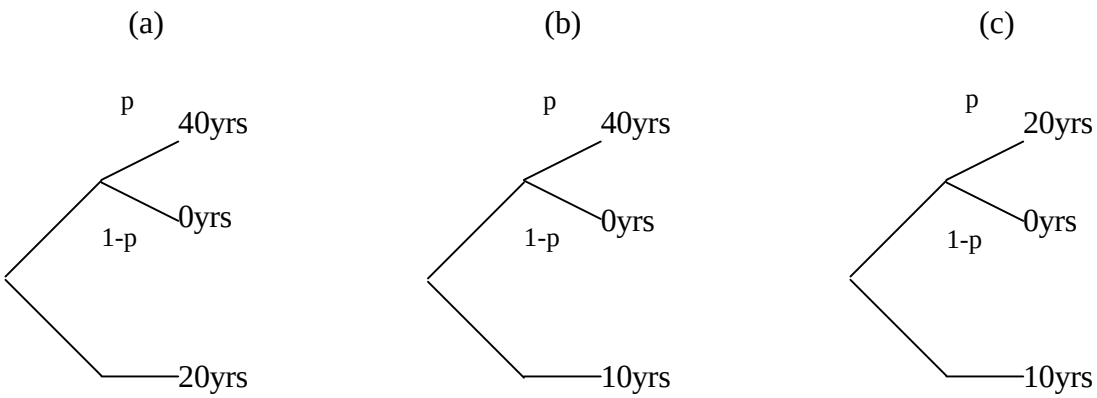

Figure 2: The context design

(a) (b)

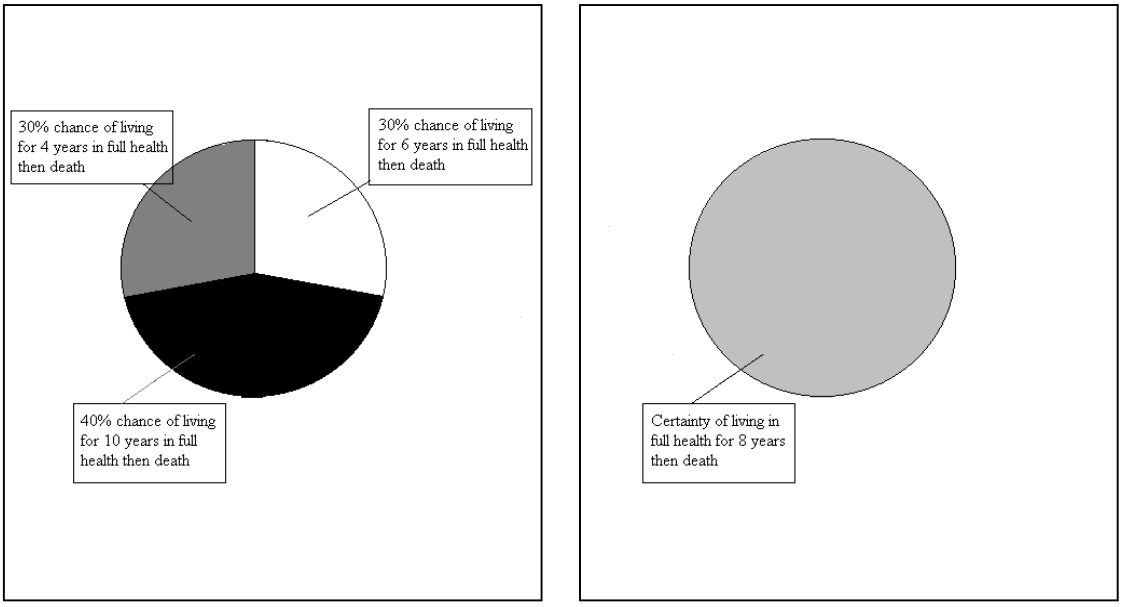

Figure 3: Allais-type contexts

Context 1

(a) (b)

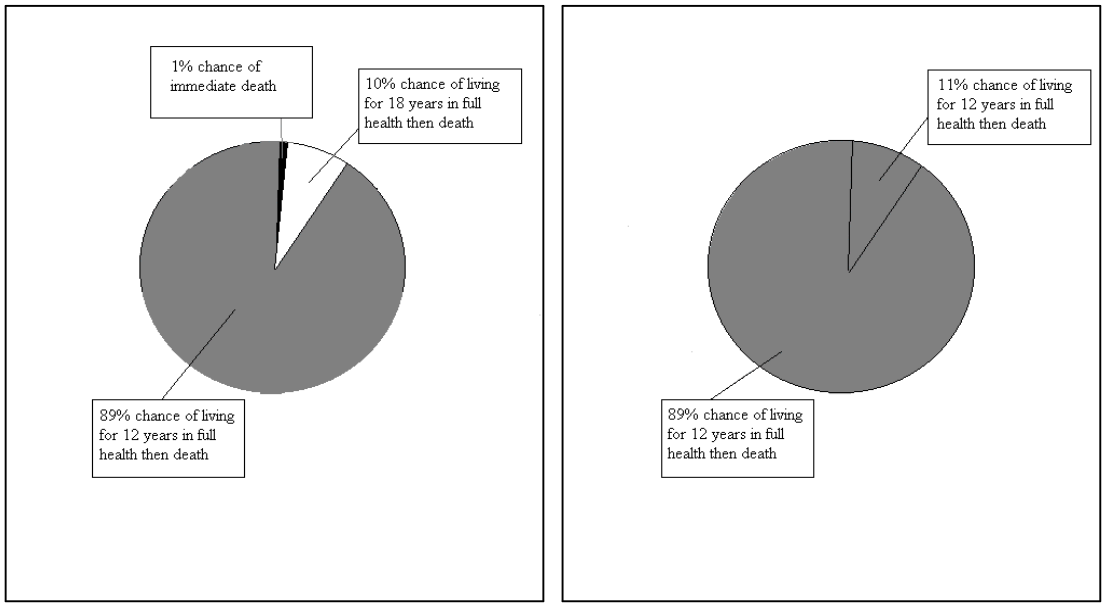

Context 2

(a') (b')