Pfanzagl Exchanges Diagnose a Continuity Anomaly Pertinent to Allais' Problem

Paul Snow

P.O Box 6134
Concord, NH 03303-6134 USA
paulusnix@cs.com

Abstract

The principle of choice often expressed as a continuity axiom recommends a single ordering of both riskless and risky prospects. An analytical device contributed by Pfanzagl is extended here to show that riskless *versus* risky monetary comparisons ought to be more nuanced, even from a perspective sympathetic to the usual axioms. Allais' example illustrates a related point to oppose the usual theory. Although the continuity difficulty is genuine, the consequences of conceding frank normative tolerance of some riskless-risky continuity violations on the specific grounds revealed are not catastrophic for the orthodox axioms.

Introduction

Von Neumann and Morgenstern's (1953) axiomatic development of expected utility for decision making under risk finds artificial intelligence application in the widely fielded technique of influence diagrams. Normative expected utility also looms large in a motivation of orthodox Bayesianism, descended from Ramsey (1926) and Savage (1972), prevalent among uncertaintists in the artificial intelligence community and in the world at large.

Criticism of the received decision axioms has often been based on puzzle problems, where some people's intuitions about rational risk-taking behavior conflict with expected utility principles. One of the best known and longest-lived of these problems was posed by Maurice Allais (1953).

Some people prefer a lottery offering a 10% chance of some positive amount of money (call it *great*) and a complementary chance of no gain nor loss (call that 0) rather than a lottery offering an 11% chance of a smaller-than-*great* positive amount (call that *good*) and a complementary chance of 0. The same people also prefer to take the *good* amount rather than a lottery with an 89% chance of *good*, a 10% chance of *great*, and a 1% chance of 0. There is no utility function on { 0, *good*, *great* } that agrees with both of these choices.

From its beginning, what specific aspect of the received theory Allais' puzzle threatens has been murky. The editor of *Econometrica* documented that difficulty in a note which ran with Allais' paper. Allais portrayed himself as lancing Savage's sure-thing principle, but that principle intentionally subsumes several strands of the von Neumann and Morgenstern argument, and other ideas as well.

Ellsberg's (1961) three-colors puzzle shares Allais' sure-thing target. Ellsberg, however, tells us of his more specific doubt, whether a decision maker selects a single probability distribution when given a set of possible ones.

Only two distinct outcomes appear among all of Ellsberg's options. Presumably, you would pick the option with the greater probability of the better outcome, *if you had chosen a specific probability* to guide your choices. Someone who had could reject all of the usual utility theory except for probability dominance (which even Allais accepted, see page 518 of his 1953 paper), and still give fully orthodox answers in Ellsberg.

In Allais, precise probabilities are given. Probability dominance does not decide any of Allais' choices. The two puzzles thus raise different concerns, despite their common sure-thing conflict. A target more specific than sure-thing might help, both to settle whether Allais really did reveal a normative defect in orthodox theory, and if so, what should be done about it.

The thesis of this paper is that there is an authentic difficulty in the received theory's treatment of choices between certainties and risky lotteries with money outcomes, like the choice between *good* and the three-way lottery in Allais. To establish the authenticity of the difficulty, we recall some theoretical results of Johann Pfanzagl (1959), who is fairly described as working within the von Neumann and Morgenstern tradition.

Pfanzagl's investigations into the certainty equivalents of lotteries led to the regularization of some theretofore informal decision modeling practices, which was an important advance. Nevertheless, loose threads remain. Those are examined here, and found sufficient to recommend toleration of behavior of which Allais violations furnish familiar examples.

The anomaly contradicts what is typically expressed in a continuity axiom. Axiomatizations differ, and elements of the same principle might occur in other axioms. The establishment of one difficulty hardly excludes the possibility of others. This much, however, is anomalous, and it is that which will be discussed here.

Decision theory has two roles: as a foundational component of other theories, and as an applied discipline in its own right. Unexplained non-conforming but rationally tolerable behavior would imperil both.

It emerges that the portion of decision theory which appears largely untouched by the anomaly, the comparison of risks with other risks, looks much the same as the familiar theory, and accomplishes almost all of the same things in much the same ways.

Moreover, there is no practical cause to refrain from the comparison of certainties with risks, only reason to take some care. The requisite degree of care is already exercised by many practitioners during utility curve and subjective probability assessment, and only slightly exceeds existing precautions in sensitivity analysis.

Decision Axioms

Many restatements of von Neumann and Morgenstern's axioms have appeared. The following account is typical, except that the continuity axiom is broken into two parts.

Lottery specifications are denoted as triples, *e.g.* (p: a, b) The first element *p* is the probability of receiving the second element, with a complementary probability, 1-*p*, of receiving the third element. The second or third element may be a grant or another lottery. The enclosing device may be parentheses, square brackets, or braces.

As is usual, the axioms and a theorem they imply are written for dichotomous lotteries. The easy standard extension from two to many alternatives is omitted here.

In writing about the axioms, we shall use the symbol "≥", as in $A \geq B$. The symbol "~", as in $A \sim B$, is shorthand for $A \geq B$ *and* $B \geq A$. The symbol ">", as in $A > B$, is shorthand for $A \geq B$ *but not* $B \geq A$.

$A \geq B$ is often read "*B* is not strictly preferred to *A*." While that is a fine reading, consider also "It could happen that the decision maker, holding *B* and being able to exchange it freely for *A*, would do so." Exchanges figure prominently in Pfanzagl's work soon to be discussed, and advice about exchanges is the point of the theory.

Ordering of Outcomes. For any outcomes a, b, c: a ≥ b or b ≥ a or both; if a ≥ b and b≥ c, then a ≥ c; and a ≥ a.

Since our current concern is exclusively with money, this axiom seems uncontroversial.

Transitivity. For any lotteries A, B, C: if A ≥ B and B ≥ C, then A ≥ C.

This axiom summarizes one view of the force of the received theory's prescriptions. Suppose that one of the lotteries *A* and *C* is to be chosen. Suppose we also notice that there is a series of exchanges beginning with *C* and ending with *A* which, if offered, we would accept at each step of the series, and there is at least one of those steps which we would not willingly undo.

The existence of such a series is a plausible rebutter to the proposition that we ought to choose *C* in the actual choice before us. When we choose *A* instead of *C*, then the series' existence furnishes a plausible explanation and justification of that choice.

The axiom also says that having found a strict, "one way" series of free exchanges leading from *C* to *A*, then we assume that there is no other series of exchanges leading from *A* back to *C*. If that assumption is correct, then the force of the arguments of the preceding paragraph is that much stronger. If not, however, then that would contradict this axiom. The defeat of that assumption will be salient in the upcoming discussion of Pfanzagl exchanges, and how well exchange arguments fare without the assumption's bolstering will influence the recommendations for repairs.

Probability Dominance. For any outcomes a, b where a ≥ b, and any probabilities p, q: p ≥ q ⇔ (p: a, b) ≥ (q: a, b).

There are some lotteries for which the direction of free exchange seems compelling. The axiom also justifies our notation for lotteries, in which only the probabilities, and not the propositions, appear.

Compound Probability. For any outcomes a, b and any three probabilities p, q, r: [p: (q: a, b), (r: a, b)] ~ (pq + r - pr: a, b).

Independence. For any probability p and any four lotteries or outcomes A, B, C, and D: if A ≥ C and B ≥ D, then (p: A, B) ≥ (p: C, D).

These two axioms express and apply the notion that the scope of the theory is restricted to situations where all that matters to the decision maker are the grants and whether or not he or she attains them, as gauged by probabilities. That idea has already been woven into the notation chosen for lotteries, as just noted.

Compound probability extends the idea in a blunt and straightforward way. Independence has historically been controversial, but it has a simple exchange interpretation. If we owned (p: C, D), then we would be willing to swap *C* to get *A* if the proposition whose probability is *p* comes true. We could commit to do so in advance of knowing the truth. Similarly, we could commit in advance to swap *D* to get *B*. The original lottery plus these two commitments is constructively identical to acquiring the lottery (p: A, B) on the same propositions. But we do not care about the propositions, only the probabilities, and so the axiom is assumed as stated.

That the axioms entail a restriction on the scope of the theory was fully acknowledged by von Neumann and Morgenstern. Casinos plainly operate on other principles, emphasizing the recreational possibilities of gambling. Other occasions of risk may be unpleasant.

Continuity of Risk. For any outcomes a ≥ b > c ≥ d, and for probability q in the interval (0, 1), then there is a unique probability p in the interval (0, 1) such that (q: b, c) ~ (p: a, d).

The axiom is phrased so as to confine its scope to comparisons among risky lotteries, and to leave open for the moment the question of comparison of riskless to risky prospects. The uniqueness of p could be replaced with a simple existence assumption; transitivity and probability dominance would establish its uniqueness if p existed.

Expected Utility Theorem. There exists a function U() whose domain is the set of outcomes in the interval [x, y], such that for lotteries A = (p: a, b) and C = (q: c, d) with outcomes in the interval (x, y): $A \geq C \Leftrightarrow pU(a) + (1 - p)U(b) \geq qU(c) + (1 - q)U(d)$.

A sketch proof appears in the appendix. Except for its steps (1) and (2), readers will recognize the method of proof as a close paraphrase of the usual contemporary proofs of the corresponding received theorem, *e.g.* compare with that of Pearl (1988). The exceptional steps (1) and (2) revive Ramsey's "ethically neutral proposition of probability one-half."

The method of proof is visibly to construct a series of acceptable free exchanges leading from C to A. An expected utility calculation, then, expeditiously detects the existence of an exchange series without the bother of actually constructing one. If the outcome of the calculation is a strict inequality, then we assume from transitivity that there is no series of free exchanges in the other direction.

Finally for this section, here is a version of usual, received continuity axiom.

Continuity. For any outcomes $a \geq b \geq c$: there is a unique probability p for which (p: a, c) ~ b.

Pfanzagl

If the received version of the continuity axiom is accepted, then it is easily understood how the quantity b mentioned in the axiom statement came to be called the "certainty equivalent" of the lottery (*p: a, c*). Johann Pfanzagl (1959) accepted the received axioms, and proposed a further axiom of his own,

Consistency. For amounts of money a, b, c, and x: if (p: a, c) ~ b, then (p: a+x, c+x) ~ b+x.

Among Pfanzagl's arguments was one based upon what will be called here a "Pfanzagl exchange." Someone who owned the lottery (*p: a+x, c+x*) would freely trade it for the lottery (*p: a, c*) on the same events if also given a side payment of x. The decision maker's situations before and after the exchange are identical if the payment coincides with the resolution of the lottery, and the situation improves (assuming the usual ideas about the direction of time preference) if $x > 0$ and the money is paid now, while the lottery is held to maturity.

We shall denote the situation of receiving a lottery L and a side payment of x as { *L, x* }. Unless stated otherwise, we shall mean that x is paid now, and that lottery L is resolved sometime later. The two elements are not otherwise bound to one another. The money (if $x > 0$) may be spent and the lottery may be exchanged as the decision maker wishes.

Pfanzagl went on to show that only two families of utility functions (*i.e.* the functions aU() + b, $a > 0$; all of which reach the same expected utility decisions as U() itself) conform to both the usual axioms and his consistency axiom. One family, exponential utility, displays *constant risk aversion, i.e.* the certainty equivalent of a lottery does not change if the decision maker's wealth changes. The other, linear utility, displays a special kind of constant risk aversion. It always values lotteries as equal to their expected money value, regardless of existing wealth.

A full appreciation of the role of risk aversion in decisions came after Pfanzagl's studies, based on the seminal work of Pratt (1964). It soon became apparent that constant risk aversion was descriptively hopeless, and that any normative case for obligatory constant risk aversion was untenable.

People appear typically to be decreasingly positive risk averse, especially when facing actuarially fair or favorable risks. That is, people seem to value lotteries at less than their actuarial value (the *positive* part). Further, the effect decreases as people get richer, *e.g.* perhaps one rejects (.5: \$120, - \$100), a favorable lottery, when poor, but maybe accepts it when richer. There is nothing irrational about that pattern of behavior.

Pfanzagl's proposed axiom was therefore rejected as an obligatory feature of rational choice, but is not held to be irrational if that is how you feel about risk. Practitioners were cautioned that if a situation like { (*p: a, c*), *x* } appeared in a decision problem, then it was to be recast as the lottery (*p: a+x, c+x*), and never the other way around. Indeed, all outcomes were to be coded as the "terminal wealth" for utility computation purposes if there was any chance of ambiguity (although lottery outcomes routinely continue to appear in problem statements described as changes in wealth, as we do here).

A further refinement of practice clarified the "buying" and "selling" price of a lottery. Both purchases and sales obviously change the decision maker's wealth, and so may affect the willingness to accept a lottery and the value imputed to it. The "selling price" was defined as the wealth change corresponding to the "certainty equivalent" that appears in the continuity axiom statement. The buying price for the lottery (*p: a, c*) became that amount b for which the lottery (*p: a-b, c-b*) has a certainty equivalent of 0 (that is, the *status quo* as assessed at one's wealth before acquiring the lottery).

It is easy to show that the buying and selling prices agree in sign for all utility functions, and so agree in value at zero. It is also a standard result that for decreasingly positive risk averse decision makers, the way most people seem to be, the buying price is less than the selling price when both are positive.

It is uncontroversial that the modeling conventions and definitions which accommodate non-constant risk aversion are consistent in the sense that the same constellation of risks and grants will yield the same expected utility values regardless of how the problem is described. The conventions and definitions do not deny Pfanzagl's observation that { (*p: a, c*), *x* } restates (*p: a+x, c+x*), but rather they integrate this fact into decision practice.

A Pfanzagl Exchange Puzzle

A decreasingly positive risk averse decision maker has initial wealth *w* and holds a favorable and acceptable lottery (*p: a, c*) whose selling price is *s* and whose buying price is *b*, $s > b > 0$. If asked, and after consulting the received theory for guidance, the decision maker asserts that he or she will accept no sum less than *s* in exchange for the lottery, as the theory counsels.

The decision maker is offered the combination { (*p: a-b, c-b*), *b* } to replace the lottery he or she owns, and accepts this Pfanzagl exchange for the (slightly) superior situation that it is compared with the original lottery. An opportunity to consume, perhaps to enjoy a meal at an expensive restaurant, then presents itself, which leads the decision maker to spend the *b* to experience the good. The experience ends while the residual lottery, (*p: a-b, c-b*), remains pending.

A lottery dealer then offers to take the residual lottery from the decision maker without compensation to either party. From the dealer's point of view, this may be a good deal, since the actuarial value of the lottery is positive.

After consulting the received theory, the decision maker, whose wealth is once again *w*, calculates the selling price of the lottery as zero. Finding zero to be offered, the decision maker surrenders the lottery to the dealer, as the theory counsels to be fair.

A series of free exchanges thus exists in which the decision maker would sell the original lottery (literally, accept a sum of money and surrender the lottery) for less than *s*. Part of our confidence in the reasonableness of the received theory is the assumption of the transitivity axiom that no such series exists. Had the dealer offered *b* at the outset for (*p: a, c*), the existence of this series would rebut to the decision maker's assertion that the minimum selling price, determined by the riskless sum that appears in the received continuity axiom, was *s*.

An aggressive dealer could resolve the impasse at the outset by writing the combination { (*p: a-b, c-b*), *b* } and making the Pfanzagl exchange, which is a riskless transaction for both parties. The dealer would then produce a restaurant guide, let nature take its course, and accost the decision maker as he or she left the eatery.

This is not a Dutch book problem. Nobody exploits anybody else, and the dealer can tell the decision maker all about the plan. Note also that except for that misunderstanding about what the minimum selling price is, the theory gives the decision maker good advice throughout. The Pfanzagl exchange is harmless, and what the theory counsels at the restaurant door accurately reflects the decision maker's preferences and circumstances at the time.

The story presents no difficulty for the user of a constant risk aversion utility function, for whom the buying and selling prices are equal. Nor would there be a difficulty for the user of any utility function if the buying and selling prices of the original lottery (*p: a, c*) were zero.

The story illustrates that it may be difficult in principle to answer the question "What sum of money is worth the same to you as owning the lottery *L*?" For many people, including some willing to consult the theory for advice, the truth is "It depends." The received continuity axiom assumes a different kind of answer, and requires that kind of answer for its confident and unheeded application.

That is, regardless of any specific "value" that may or may not be associated with a lottery, the lottery might be willingly exchanged for any of several amounts. The bases for that diversity of answers to the question of worth are themselves principles of choice used in building the theory, in concert with the nature of money and its uses.

Allais violations appear to be an occasion when something upon which the truthful answer depends is relevant. Among the possible considerations is the following hypothetical series of exchanges.

Let *L* be the three-way lottery in Allais, and *L** be that lottery with each of its outcomes reduced by the amount *good*. The combination {*L**, *good*} is not inferior to *L*. Suppose one has consumption opportunities (or exisitng family responsibilities, or other pressing resource claims), so that one has immediate uses for some part *c* of *good* besides gaming.

Depending upon the buying price of *L** and the size of *c*, it could be that *good-c* alone is preferred to {*L**, *good-c*}; if so, then you would give away the residual lottery if you could. The state *good-c* and the satisfaction of the competing claims for resources can be directly achieved by choosing *good*, but cannot be directly achieved by choosing *L*.

A strong continuity principle beset by paradox when confronted with the implications of what money is and how it can be used furnishes no grounds for rejecting as irrational violations of the principle in monetary problems.

Attention now turns to the extent of the damage and to strategies for repair.

Consequences for Other Theories

The received axioms' domain of origin is Von Neumann and Morgenstern's game theory. Pfanzagl pointed out that von Neumann and Morgenstern's concept of "strategic equivalence," adopted by them for *n*-person games, implies the same requirement as his own consistency axiom, that utility functions be either linear or exponential.

As we have seen, restricting the form of the utility function to that extent is now widely held to be unsuitable for general-purpose decision theory. On the other hand, it is quite possible that the specific domain of strategic choice may have attributes which justify the restriction. Clearly, von Neumann and Morgenstern thought that such attributes were present, although they presumably did not know how the attributes imply the restriction.

If they are right about the attributes, then the axioms, in the context that von Neumann and Morgenstern deployed them, would be immune to Pfanzagl exchange puzzles of the kind discussed in this paper. While puzzles do reveal a difficulty in generalizing the axioms to other domains, no change in the theory of the specific domain is necessarily supported by the existence of challenges in generalizing one of its elements.

In discussing other theories which rely on expected utility, let us begin by saying that confining the theory's scope in money choices to comparisons among risks and between risks and the *status quo* is *not even remotely* being advocated here. Nevertheless, it is interesting to consider how little difference that would make to expected utility's role as a foundational element of some other theories.

For example, experiment design in statistics, to the extent that it involves money, like the expenses of experiments, rarely offers a literal stipend or fine as an alternative to uncertain and costly experimentation. Insofar as experiment design involves things other than money, such as the inalienable worth of aiding the progress of knowledge, then it is untouched by anything said here.

Prominent among the theoretical uses of the received theory is its contribution to gambling-based semantics and justifications for subjective probability. But the issue there is to show how probabilities in the open unit interval (0 1) can describe degrees of uncertainty. Since the basic expected utility theorem is secure, both as to truth and as to method of demonstration, the arguments based upon it are also secure in (0 1). As for how the bounding integers denoting certainty might be combined with fractions denoting uncertainty, that point was well covered by Boole in the Nineteenth Century.

Confinement of scope of continuity would not entirely moot Allais' critique. He discusses at length the problem of "near certainties." No reasonable theory could recommend any dramatic distinction between the amount b and the risk (.999: b b+1), which must be similar in affect and effect to b itself. The lottery even shares with money a key feature which drives the Pfanzagl exchange puzzle, effective fungibility. If those truly are the odds, then I might well offer to pay the restauranteur (.999: b b+1) for my meal, and we would both call it square.

While this untidiness may be a concern for theorists relying upon expected utility, it is not a new concern. It is well known that judgments about very small probabilities (*i.e.* a common species of near certainty) are fraught (von Winterfeldt and Edwards 1986). Maybe, then, what Allais is telling theorists about near certainty is something that they would discover for themselves soon enough.

Consequences for Decision Craft

There is a case for retaining the continuity axiom unaltered, warts and all. The advice the theory gives is good advice, well worth thinking about. In the Pfanzagl exchange puzzle, "Hold out for a better price" and "Under the new circumstances, settle for less" are both defensible based on the information that would be available to any advisor at the times when advice is sought.

So, too, is the advice in Allais good. Savage recounts his mental journey from champion of the received theory, through his descent into violation, to his re-embrace of orthodoxy. His epiphany came when he constructed one of those series of exchanges the theory guarantees to exist.

He imagined literal lottery tickets, and saw the certainty of *good* as owning three blocks of tickets whose probabilities of winning add up to 100%. Both of Allais' hypotheticals, he then realized, ask whether he would forfeit a "1% chance of *good*" ticket in order to exchange a "10% of *good*" ticket block for a "10% chance of *great*" block.

This is an altogether sensible way to view the situation. There are other ways, too. Allais' original respondents, likely among them European survivors of a then-recent devastating war, might have seen that the certainty of *good*, with the right to leave the encounter with it in hand, buys food and fuel, while the illiquid chance of *good* in the other choice has no current use except to be held or traded.

Nevertheless, that Savage's is not the only way to look at the situation furnishes no argument for not at least considering the theory's perspective. The ability of the received theory to provide its perspective relies on its unmodified axioms.

A rational person might consider a choice from the theory's perspective and then decline its advice in some situations where monetary certainties and near certainties contend with heartier risks. On this view, what the received continuity axiom says about risks and certainties is rationally tenable, but compliance is not rationally obligatory. That has some implications for practical decision modeling.

One practical task is the assessment of a decision maker's utility curve and subjective probabilities, often based on hypothetical betting. As von Winterfeldt and Edwards (1986) report, practitioners have long been wary of posing for curve assessment purposes hypothetical bets which feature choices between risks and amounts for sure. The basis for that wariness may be concern about a psychological "certainty effect," whose empirical foundation is surely in part attributable to Allais. Be that as it may, the precautions already in place seem adequate to prevent miscalibration from anything discussed here.

Later on in a decision modeling episode comes sensitivity analysis, which checks that utility and probability values found to be crucial to the final recommendation have been accurately measured. "Near certainties," for example, might and already do receive some of the attention they need at this point.

The continuity anomaly is not a matter of accurate measurement. It is the situation itself that warrants scrutiny for variably risk-averse decision makers. Thus, existing consensus practice should be enlarged to include re-examination of all choice junctures which offer the possibility of release, or near-release, from monetary uncertainty, whether or not the numbers involved present a "close call" in expected utility value.

Also unlike existing sensitivity analysis, this checking must be done at the problem description level, not on the model. Modelers routinely "rearrange" lotteries leading to choices into choices between lotteries, and *vice versa*. The issue of concern is whether or not the actual decision situation, as opposed to its "equivalent" model, offers the possibility of release.

The *release* that the sensitivity analyst is looking for is the appearance of resources which are able to leave the frame of the problem. One example would be a terminal grant, like the one in Allais' problem (if you choose *good* for sure, then both you and the money leave the problem).

For another instance, in a slightly more complicated version of the earlier Pfanzagl exchange puzzle, suppose the decision maker were offered the opportunity to replace (p: a, b) with either the Pfanzagl exchange or else the amount $(b + s)/2$. That amount is less than s, and the theory would bluntly counsel rejecting it, even as the Pfanzagl exchange is recognized to be harmless.

A shrewd counselor would ask the decision maker "how sure are you that you will hold onto that side payment of b?" The decision maker may in fact be better off taking the $(b + s)/2$, despite its being less than s.

Then again, the decision maker might intend to hold onto the side payment through the resolution of the residual lottery (of positive value to the decision maker if his or her wealth is greater than w, as it now is, and will be so long as the side payment is held). If dealers lurk and bistros beckon, then there is no way to know which is really the better exchange, except to ask.

Conclusions

The received decision theory has difficulty when variably risk averse decision makers compare monetary certainties with monetary risks. Appeal to the theory thus fails to exclude Allais violations, among other things, from rational respectability. Nettlesome though that may be, the consequences of the anomaly are nevertheless mild.

There is scant threat to the use of the expected utility arguments as components of other theories, nor is any strenuous reform of decision analytical routine indicated. The received theory as it stands reliably provides advice worthy of consideration, albeit not of uncritical compliance, by a reasoning agent confronting risk.

Still, however mild, there really is a problem with the usual decision theory when it confronts variable risk aversion and liquidity. That warrants some caution before implementing the theory's prescriptions, or decrying the irrationality of other incompatible advice.

References

Allais, M. 1953, Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axioms de l'Ecole Americaine, *Econometrica* 21: 503-546 (with editor's note).

Ellsberg, D. 1961, Risk, Ambiguity, and the Savage Axioms, *Quarterly Journal of Economics* 75: 643-669.

Pearl, J. 1988, *Probabilistic Reasoning in Intelligent Systems*, San Mateo CA: Morgan Kaufman, chapter 6.

Savage, L. J. 1972. *The Foundations of Statistics*, New York: Dover.

Pfanzagl, J. 1959, A General Theory of Measurement Applications to Utility, *Naval Research Logistics Quarterly* 6: 283-294.

Pratt, J.W. 1964, Risk Aversion in the Small and the Large, *Econometrica* 32: 122-136.

Ramsey, F. P. 1926, Truth and Probability, reprinted in Kyburg, H.E. Jr. and Smokler, H.E. 1964, *Studies in Subjective Probability*, New York: Wiley.

von Neumann, J. and Morgenstern, O. 1944, *Theory of Games and Economic Behavior*, Princeton NJ: Princeton University Press, chapter 3 and Appendix.

von Winterfeldt, D. and Edwards, W. 1986, *Decision Analysis and Behavioral Research*, Cambridge UK: Cambridge University Press, chapters 4, 7, and 11.

Appendix

Expected Utility Theorem. There exists a function U() whose domain is the set of outcomes in the interval [x, y], such that for lotteries A = (p: a, b) and C = (q: c, d) with outcomes in the interval (x, y): A $\geq$ C $\Leftrightarrow$ pU(a) + (1 - p)U(b) $\geq$ qU(c) + (1 - q)U(d).

Sketch Proof. Despite the lack of a specification of how riskless amounts are compared with risky lotteries, the proof of the theorem is very similar to the proofs of the corresponding theorem from the received axioms.

The basic steps are: (1) for z in (x, y) define U(z) as the probability for which (.5: z, x) ~ [U(z): y, x]. The existence and uniqueness of that probability for a given z follows from continuity of risk.

For the $\Rightarrow$ direction:

(2) Observe that A $\geq$ C just when (.5 A, x) $\geq$ (.5 C, x), from independence and the ordering of outcomes. Note that in applying the independence axiom here, lotteries have been compared only with lotteries, and riskless grants compared only with riskless grants.

(3) From compound probability, (.5 A, x) ~ [p: (.5 a, x), (.5 b, x)], and (.5: C, x) ~ [q: (.5: c, x), (.5: d, x)]

(4) By independence, substitute [U(a): y, x] for (.5 a, x), and substitute similarly for the lotteries in b, c, and d on the right sides of the tilde expressions in step 3.

(5) Rearrange the resulting { p: [(U(a): y, x], [U(b): y, x] }by compound probability to [pU(a) + (1-p) U(b): y, x], and proceed similarly for the other expression in c and d.

(6) Since the lotteries from (5) both involve the same prizes, conclude from probability dominance that pU(a) + (1-p) U(b) $\geq$ qU(c) + (1-q)U(d).

The argument for the converse applies steps parallel to (2) through (5) in reverse order. //