

JEFFREY HELZNER

ON THE APPLICATION OF MULTIATTRIBUTE UTILITY THEORY TO MODELS OF CHOICE

ABSTRACT. Ellsberg (*The Quarterly Journal of Economics* 75, 643–669 (1961); *Risk, Ambiguity and Decision*, Garland Publishing (2001)) argued that uncertainty is not reducible to risk. At the center of Ellsberg’s argument lies a thought experiment that has come to be known as the three-color example. It has been observed that a significant number of sophisticated decision makers violate the requirements of subjective expected utility theory when they are confronted with Ellsberg’s three-color example. More generally, such decision makers are in conflict with either the ordering assumption or the independence assumption of subjective expected utility theory. While a clear majority of the theoretical responses to these violations have advocated maintaining ordering while relaxing independence, a persistent minority has advocated abandoning the ordering assumption. The purpose of this paper is to consider a similar dilemma that exists within the context of multiattribute models, where it arises by considering indeterminacy in the weighting of attributes rather than indeterminacy in the determination of probabilities as in Ellsberg’s example.

KEY WORDS: multiattribute, revealed preference, descriptive, uncertainty, methodology

JEL CLASSIFICATIONS: D12, D81

1. INTRODUCTION

Luce and Raiffa (1989) make what is now a well-known distinction between “decision making under risk” and “decision making under uncertainty.” The essence of this distinction is that in decision making under risk the decision maker has access to an objective probability distribution on the relevant state space, while in decision making under uncertainty such access is lacking. For those who endorse subjective

expected utility theory, the distinction between risk and uncertainty is of little significance since it is a basic tenet of subjective expected utility theory that the decision maker has a subjective probability distribution even in cases where objective probabilities fail to be salient. Following a tradition that dates back at least as far as Keynes (1921) and Knight (1921), Ellsberg (1961; 2001) argued against the conflation of risk and uncertainty that is suggested by subjective expected utility theory. Central to Ellsberg's argument that there exist "uncertainties that are not risks" is the following example:

EXAMPLE 0.1 (*Ellsberg*) A ball is to be selected at random from an urn containing 90 balls. 30 of these balls are red. Each of the remaining 60 balls in the urn is either black or white, although the exact ratio of black balls to white balls is unknown. Consider the following two decision problems.

Problem 1	Red	Black	White
<i>e</i>	\$100	\$0	\$0
<i>f</i>	\$0	\$100	\$0

Problem 2	Red	Black	White
<i>g</i>	\$100	\$0	\$100
<i>h</i>	\$0	\$100	\$100

It is well known that a significant number of decision makers, including many who are sophisticated in their knowledge of decision theory, regard *e* as uniquely admissible in Problem 1, while they regard *h* as uniquely admissible in Problem 2. Of course this *e-h* choice pattern is incompatible with subjective expected utility theory since – regardless of the decision maker's utility function and subjective probability distribution on the state space {Red, Black, White} – if the expected utility of *e* is greater than that of *f*, then so must the expected utility of *g* be greater than that of *h*. More generally, the *e-h* choice pattern is incompatible with all theories that maintain the following conditions:

Ordering: The decision maker's preferences constitute a weak order on the set of alternatives. That is, the following conditions hold for all alternatives x , y , and z : (*Transitivity*) if x is at least as good as y and y is at least as good as z , then x is at least as good as z , (*Completeness*) either x is at least as good as y or y is at least as good as x . Furthermore, x is admissible in X , where X is a set of alternatives, just in case x is in X and, for all y in X , if y is at least as good as x , then x is at least as good as y .

Independence: If w , x , y , and z are alternatives on a state space Φ – that is, they are functions from Φ to a set of outcomes – and the following conditions are satisfied for some nonempty subset ψ of Φ : (1) w is at least as good as x , (2) $w(s) = x(s)$ and $y(s) = z(s)$ for all s in ψ , (3) $w(s) = y(s)$ and $x(s) = z(s)$ for all s not in ψ , then y is at least as good as z .

It is clear that no theory endorsing these conditions will be able to accommodate the *e-h* choice pattern: Assume that the conditions hold. By the ordering assumption, either e is at least as good as f or f is at least as good as e . Suppose that the decision maker regards e as uniquely admissible in Problem 1. It must be the case that if f is at least as good as e , then e is at least as good as f . Hence, either way, e is at least as good as f . Since $e(s) = f(s)$ and $g(s) = h(s)$ for s in {White} while $e(s) = g(s)$ and $f(s) = h(s)$ for s in {Red, Black}, it follows that g is at least as good as h . Hence, g is admissible in Problem 2, which is inconsistent with the *e-h* choice pattern.

Attempts to accommodate the indicated choice pattern have tended to focus on dropping independence while maintaining ordering. Indeed, the alternatives to subjective expected utility that are endorsed in Ellsberg (1961) and Ellsberg (2001) relax independence while preserving the ordering requirement. Levi (1974) has advanced what is perhaps the most notable decision theory that drops ordering. We now turn to a brief discussion of Levi's decision theory.

Let P be a set of probability distributions on a finite state space Φ . If p is in P and x is an alternative on Φ , i.e. a

function from Φ to the set of outcomes, then we will write $E_p(x)$ for the expected utility of x against p – we will assume that a real-valued utility function on the set of outcomes has been specified. We will write $S_P(x)$ for the greatest lower bound of $\{E_p(x) \mid p \in P\}$. The relevant parts of Levi's decision theory concern the following notions of admissibility: (*E-admissibility*) x is E-admissible in X if and only if x is in X and there is a probability distribution p in P such that $E_p(x) \geq E_p(y)$ for all y in X . (*S-admissibility*) x is S-admissible in X if and only if x is in X and $S_P(x) \geq S_P(y)$ for all y in X .¹ (*Admissibility*) x is admissible in X just in case x is S-admissible in the set of all y that are E-admissible in X .

As an illustration of Levi's decision theory consider the following application to Ellsberg's three-color urn example: Let P be the set of all distributions p on the state space {Red, Black, White} such that $p(\text{Red}) = \frac{1}{3}$ and let the utility function on rewards coincide with monetary value, i.e. the utility of $\$k$ is k . In Problem 1, e is E-admissible in $\{e, f\}$ since $E_p(e) \geq E_p(f)$ for all p in P such that $p(\text{Black}) \leq \frac{1}{3}$, while f is E-admissible in $\{e, f\}$ since $E_p(f) \geq E_p(e)$ for all p in P such that $p(\text{Black}) \geq \frac{1}{3}$. However, since $S_P(e) = \frac{100}{3}$ while $S_P(f) = 0$, it follows that e is uniquely admissible in $\{e, f\}$. Turning to Problem 2, g is E-admissible in $\{g, h\}$ since $E_p(g) \geq E_p(h)$ for all p in P such that $p(\text{White}) \geq \frac{1}{3}$, while h is E-admissible in $\{g, h\}$ since $E_p(h) \geq E_p(g)$ for all p in P such that $p(\text{White}) \leq \frac{1}{3}$. In this case, since $S_P(g) = \frac{100}{3}$ while $S_P(h) = \frac{200}{3}$, it follows that h is uniquely admissible in $\{g, h\}$.

The following condition goes by various names, e.g. *Sen's α* (Sen, 1971; Kreps, 1988) and the *independence of irrelevant alternatives* (Ray, 1973), within the literature on rational choice: If x is admissible in X and x is an element of Y , where Y is a subset of X , then x is admissible in Y . It is well-known that this condition is necessary for rationalization by a weak order on the set of alternatives (Sen, 1971). Suppose that the ordering condition holds. If x is admissible in X , then it follows that x is at least as good as every y in X . If Y is a subset of X and x is in Y , then it follows that x is

at least as good as every y in Y . Hence, by the second part of the ordering condition, x is admissible in Y .

As noted (Levi, 1974, 1986), Levi's theory allows for violations of Sen's α condition. Consider the following example, a variation of the ones in Levi (1974) and Levi (1986), under the background assumptions of Ellsberg's three-color problem:

	Red	Black	White
t	\$300	\$0	\$300
u	\$300	\$300	\$0
v	\$150	\$150	\$150

Observe that t is E-admissible in $\{t, u, v\}$ since $E_p(t) \geq E_p(u)$ and $E_p(t) \geq E_p(v)$ when $p(\text{White}) \geq \frac{1}{3}$. By symmetry, it is clear that u is E-admissible in $\{t, u, v\}$ since $E_p(u) \geq E_p(t)$ and $E_p(u) \geq E_p(v)$ when $p(\text{Black}) \geq \frac{1}{3}$. However, v is not E-admissible in $\{t, u, v\}$ since $E_p(t) > E_p(v)$ when $p(\text{White}) \geq \frac{1}{3}$ while $E_p(u) > E_p(v)$ when $p(\text{White}) \leq \frac{1}{3}$. Since $S_p(t) = S_p(u)$, it follows that t and v are the admissible alternatives in $\{t, u, v\}$. However, suppose that we remove u and consider admissibility in $\{t, v\}$. From the previous calculations it is clear that t is E-admissible in $\{t, v\}$, but this time v is E-admissible since $E_p(v) \geq E_p(t)$ when $p(\text{White}) \leq \frac{1}{6}$. While both t and v are E-admissible in $\{t, v\}$, v is uniquely admissible in $\{t, v\}$ since $S_p(v) = 150 > 100 = S_p(t)$. Thus, Sen's α condition is violated.

2. FROM EXPECTED UTILITY TO MULTIATTRIBUTE UTILITY

In the previous section we reviewed how Ellsberg's three-color problem invites us to consider the prospects for relaxing either the ordering assumption or the independence assumption of expected utility theory. Seidenfeld (1988) has argued that, at least normatively, the prospects are not good for those theories that maintain ordering while relaxing independence. However, Levi's theory avoids these unfavorable prospects,

since Levi's theory allows for violations of conditions such as Sen's α that are necessary in maintaining the ordering assumption. In what follows we will argue that issues similar to what has been discussed in the context of expected utility theory are in fact relevant within the context of multiattribute utility theory and that the relevance of these issues is such that the traditional status of the ordering assumption within multiattribute utility theory ought to be reconsidered.

We now turn our attention from normative models of decision making under uncertainty to descriptive accounts of choice among multiattribute alternatives. As before we will be concerned with alternatives that admit a certain functional representation. However, whereas previously we were concerned with the interpretation of alternative $f: \Omega \rightarrow V$ as a function from a set of states to a set of outcomes, so that $f(s)$ was understood as the outcome of selecting alternative f if state s should obtain, our present concern is the interpretation of $f: \Omega \rightarrow V$ as a function from the set of attributes to a set of descriptions according to which $f(s)$ specifies the status of alternative f with respect to attribute s .

Use of the multiattribute paradigm is widespread throughout the decision sciences, e.g. from applications to decision analysis (Keeney and Raiffa, 1993) to models of consumer behavior (Wilkie and Pessemier, 1973). Among the most important multiattribute models are those that assume an additive representation of the decision maker's preferences so that there exists a family of real-valued $\{u_i\}_{i \in \Omega}$ functions on the set of alternatives such that following conditions hold for all alternatives f and g : (1) if $f(s) = g(s)$, then $u_s(f) = u_s(g)$ and (2) the decision maker judges f to be at least as good as g if and only if $\sum_{s \in \Omega} u_s(f) \geq \sum_{s \in \Omega} u_s(g)$.

The ordering assumption is implicit in the usual applications involving these models. In other words it is assumed that, when given an opportunity to select from a set X of alternatives, those among X that are deemed admissible are precisely those that maximize the additive index (i.e. those f in X such that $\sum_{s \in \Omega} u_s(f) \geq \sum_{s \in \Omega} u_s(g)$ for all g in X). Furthermore, the independence condition, suitably reinterpreted

by taking Ω to be a set of attributes rather than a set of states, is a necessary condition for the existence of such an additive representation.

Now suppose that an analyst is charged with the task of modeling choice behavior over the following space of alternatives: S_1 , S_2 , and S_3 are stocks and an alternative is specified as a triple $[x_1, x_2, x_3]$, where x_i is the dollar amount allocated to S_i . It will be assumed that the choosing agent does not provide any of his own money for distribution between these three stocks; that is, the principal is supplied by some other party, e.g. the agent's employer. Furthermore, let us imagine that, prior to being confronted with any nontrivial menu of such alternatives, the choosing agent is told that (1) the 1-year return rate for S_1 is known to be $\frac{1}{3}$, (2) the 1-year return rates for both S_2 and S_3 are positive and, although the exact rates for these stocks are unknown, it is known that the sum of the return rates for S_2 and S_3 is $\frac{2}{3}$, and (3) if $[x_1, x_2, x_3]$ is selected, then the choosing agent will, in exactly 1 year from the time of selection, surrender that alternative in exchange for the cash proceeds that were generated by $[x_1, x_2, x_3]$ in the year following the time of selection.² So, for example, a selection of $[x_1, x_2, x_3]$ would result in a payout of $r_1x_1 + r_2x_2 + r_3x_3$ dollars to the choosing agent at the end of the 1-year period, where r_i is the 1-year return rate for stock S_i . Thus, while the choosing agent does not supply the principal for such an alternative, this agent is provided with the proceeds that are generated by the alternative that he chooses.

Given that these alternatives have no value other than the proceeds that they generate, it seems reasonable for our analyst to suppose that each subject's conditional preferences are additive in the sense that his preferences, conditional on the assumption that λ_0 is the 1-year return rate for S_2 , are such that $[x_1, x_2, x_3]$ is at least as good as $[y_1, y_2, y_3]$ iff $\frac{1}{3}x_1 + \lambda_0x_2 + \left[\left(\frac{2}{3}\right) - \lambda_0\right]x_3 \geq \frac{1}{3}y_1 + \lambda_0y_2 + \left[\left(\frac{2}{3}\right) - \lambda_0\right]y_3$. With this assumption, our analyst might hypothesize that a model of such a subject's choice behavior over the given space of

alternatives will in some way depend on the subject's personal probability distribution over the possible values of the 1-year return rate for S_2 . A model of this sort presupposes that the subject has such a numerically precise distribution over these possible values. This, of course, is just the sort of assumption that is called into question by the previously cited work of Keynes, Knight, Ellsberg, and Levi. Additional arguments against this assumption of numerical precision have been advanced by Kyburg (1968), Gardenfors and Sahlin (1982) and in an explicitly descriptive context by Arlo-Costa and Helzner (2005).

Another obvious difficulty with such an approach is that it seems as though no reasonable candidate distributions are available, since, as the reader will recall, if p is a finitely additive probability on some uncountable set S , e.g. open interval $(0, \frac{2}{3})$, then $p(s) = 0$ for some s in S : Suppose that this is not the case. That is, suppose that $p(s) > 0$ for all s in S . For each positive integer n , let $S_n = \left\{ s \in S \mid p(s) > \frac{1}{n} \right\}$. It follows that $S = \bigcup_{i=1}^{\infty} S_i$. Hence, for some positive integer k , S_k has infinitely many members; if not, then S is a countable union of finite sets, which is impossible since S is assumed to be uncountable. Let s_1, \dots, s_{k+1} be $k+1$ distinct elements in S_k . By assumption, $p(s_i) > \frac{1}{k}$ for $i = 1, \dots, k+1$. Hence, by finite additivity, $p(\{s_1, \dots, s_{k+1}\}) = \sum_{i=1}^{k+1} p(s_i) > 1$, which contradicts the assumption that p , as probability measure, takes its values in the closed interval $[0, 1]$.³

Of course there are measures that assign zero probability to some of the possible return rates for S_2 , but there does not seem to be anything in the description of the problem that suggests one such assignment to the exclusion of others. Difficulties of the sort just described might be ignored on the grounds that practically, if not theoretically, the possible return rates for S_2 constitute a finite set. The following sort of difficulty does not depend on having a continuum of possible return rates for S_2 . Indeed, it will suffice to assume that the possible return rates are constrained to a fine, but finite

approximation of $(0, \frac{2}{3})$, e.g. the set of all λ such that $\lambda = (\frac{i}{10^{10}})(\frac{2}{3})$ for some integer i such that $1 < i < 10^{10}$.

Suppose that our analyst presents the following two decision problems to subjects:

Problem 1	S_1	S_2	S_3
p	\$100,000	\$0	\$0
q	\$0	\$100,000	\$0

Problem 2	S_1	S_2	S_3
r	\$100,000	\$0	\$100,000
s	\$0	\$100,000	\$100,000

This example may be thought of as a multiattribute analogue of Ellsberg’s three-color problem where the attributes correspond to the stocks, i.e. so that a given alternative’s position with respect to the i th attribute is given by the dollar amount this alternative has committed to S_i . With respect to this multiattribute example, we do not think that it would be unreasonable to regard p – the alternative with \$100,000 in S_1 and \$0 in the other two stocks – as uniquely admissible in the Problem 1, while regarding s as uniquely admissible in Problem 2. Indeed, since we are now dealing with guaranteed amounts, e.g. p is *guaranteed* to pay \$33,333.33 ..., it seems to us that this pattern is at least as reasonable as the analogous pattern in Ellsberg’s three-color problem.

What is our analyst to make of such a pattern? If our analyst follows tradition and holds ordering sacrosanct, then he must conclude that independence is violated. However, as we know, this is not the only way to interpret the data. Our analyst could consider the possibility that ordering fails and that a more complicated theory is needed to account for the choices in question. Such an explanation could also have the virtue of being able to accommodate the very reasonable possibility that the subject’s conditional preferences are additive in the manner considered previously. Assuming that

these conditional preferences are additive, and thus satisfy independence, the analyst who refuses to consider violations of ordering is in the uncomfortable situation of endorsing a model that accepts the independence of each of the decision maker's conditional preferences but denies independence in the case of the decision maker's unconditional preferences. Such an analyst should at least need to supply some explanation of the apparent violation of the following monotonicity principle: the number of perceived interactions between attributes is weakly monotonic with respect to increases in the perceiver's information state. In the present example, the maximal information states coincide with the possible values of the return rate for S_2 . The above assumption concerning the additivity (hence, independence) of the subject's conditional preferences suggests that the subject would not perceive any interactions between attributes in any of the maximal information states, but, according to the revealed preference of p over q and s over r , perceives some sort of independence-violating interaction between attributes while in the weaker information state that is assumed in the example, i.e. the state in which the agent knows simply that the return rate for S_2 is in the open interval $(0, \frac{2}{3})$.

On the other hand, if our analyst is willing to consider relaxing the ordering assumption, then it is possible to accommodate the observed choices while avoiding the tension that has just been described. If the agent's unconditional preferences are given as the intersection of the agent's conditional preferences, i.e. $[x_1, x_2, x_3]$ is (unconditionally) at least as good as $[y_1, y_2, y_3]$ iff $\frac{1}{3}x_1 + \lambda x_2 + \left[\left(\frac{2}{3}\right) - \lambda\right]x_3 \geq \frac{1}{3}y_1 + \lambda y_2 + \left[\left(\frac{2}{3}\right) - \lambda\right]y_3$ for all possible values of λ , then the agent's unconditional preference relation, what Levi calls categorical weak preference, is incomplete but satisfies independence since each of the conditional preference relations do. Of course the relation between unconditional preference and admissibility must be different from what is supposed in the traditional revealed preference methodology. An analogue of Levi's theory offers one example of how these things might be

related: $[x_1, x_2, x_3]$ is first-tier admissible in X if and only if $[x_1, x_2, x_3]$ is in X and there is some λ such that $\frac{1}{3}x_1 + \lambda x_2 + \left[\left(\frac{2}{3}\right) - \lambda\right]x_3 \geq \frac{1}{3}y_1 + \lambda y_2 + \left[\left(\frac{2}{3}\right) - \lambda\right]y_3$ for all $[y_1, y_2, y_3]$ in X . In other words, $[x_1, x_2, x_3]$ is first-tier admissible in X if it is available and there is some condition under which $[x_1, x_2, x_3]$ is at least as good as every other alternative in X . Similarly, by analogy with S-admissibility, we may formulate a notion of second-tier admissibility. Admissibility is then given by combining the first and second-tier notions: $[x_1, x_2, x_3]$ is admissible in X if and only if it is second-tier admissible in the set of alternatives that are first-tier admissible in X .

3. CONCLUDING REMARKS

Indeterminacy in the context of multiattribute models has received some attention in applications to areas such as problems involving the selection of a candidate to fill a particular job (Weber, 1985) and problems concerning the modeling of consumer behavior (Kahn and Meyer, 1991). However, in keeping with the received view concerning analogous issues in the context of expected utility theory, work in this area has insisted on maintaining the ordering assumption.

It is not difficult to find examples where indeterminacy in the weighting of attributes is present, or at least ought to be suspected. For example, in 2007 the New York City Department of Education decided that, for students applying to pre-kindergarten through second grade, admission to New York's "Gifted and Talented" programs will be determined by using a combination of two standardized tests: The *Otis-Lennon School Ability Test* (OLSAT) and the *Bracken School Readiness Assessment* (BSRA) (Klein, 2007). Specifically, each student's overall Gifted and Talented score will be computed as a weighted average of their OLSAT score and BSRA score, where the weight on the OLSAT score is .75 and the weight on the BSRA is .25. Admission to the Gifted and Talented

programs will be determined according to these overall scores as computed according to this weighted average.

Now, even if we accept that some particular weighted average of scores from these two tests does, in some sense, provide the right model of giftedness, it seems unlikely that the correct value of this weighting parameter could be established with such precision. Are we to believe that several independent attempts at measurement were made and a value of .75 on the OLSAT (.25 on the BSRA) was obtained on each such attempt? Would that better understood procedures of physical measurement had such little need for a theory of error! If, more plausibly, a variety of values were recorded as a result of these various attempts, then we might imagine that the reported value of .75 on the OLSAT (.25 on the BSRA) was informed by the variety of values that were recorded. Perhaps the most familiar way of doing this is to report a value that is derived as a weighted average of the recorded values. That is, if O_1, \dots, O_n are the values for the weight on the OLSAT as obtained from the $m \geq n$ independent attempts at measurement, then one might be tempted to derive a “correct” value for the OLSAT weight by taking a weighted average of O_1, \dots, O_n , but which of the continuum many weighted averages should be taken? If the background assumptions are not sufficient to justify a particular weighting, then this is the reality of situation and there is no reason to expect that a series of theoretical moves will lead to a resolution of the indeterminacy in question. However, principled judgments of admissibility can be made in the face of persistent (or fundamental) indeterminacy, the kind of indeterminacy that precludes a complete ordering of the set of alternatives. As considered in the previous sections, Levi’s work offers one theoretical option in this direction. Of course, other theoretical options that relax ordering might be considered, e.g. Sen’s notion of maximality (Sen, 1970).⁴ Regardless of the particular theory that connects unconditional preference and admissibility, the essential point is that indeterminacy involving the weighting of attributes might require us to rethink the status of the ordering assumption within the context of

multiattribute utility theory. Moreover, recognition of the possibility of indeterminacy in the weighting of attributes should motivate us to adopt approaches for testing the descriptive adequacy of multiattribute theories in a way that does not assume the very ordering assumption that is at issue. Thus, for example, where indeterminacy in the combining of attributes is a serious concern, we ought to eschew the use of questionnaires that require subjects to rank the available alternatives. Similarly, in light of the possibility of violating conditions, such as Sen's α , an open-minded approach to the testing of such theories might necessitate the use of questionnaires that incorporate choice problems that involve more than two alternatives. The use of larger choice sets could place greater demands on the subjects, which in turn could complicate analysis of the resulting data. Nonetheless, the possibility of such difficulties is not a reason to adopt approaches that are blind to violations of ordering.

ACKNOWLEDGMENTS

I would like to thank Teddy Seidenfeld, Horacio Arlo-Costa, and Isaac Levi for many helpful discussions on these topics over the years.

NOTES

1. Other "security" rules might also be considered in the specification of S-admissibility, e.g. *maximin* or *minimax risk*. Levi has maintained an open mind on this topic: "I do not think there is any way of deciding which method of determining security levels is preferable or rational. That question ought to be left up to the decision maker and is to be regarded as a value commitment on the agent's part." (Levi, 1986)
2. With respect to the relationship between their respective return rates, one might imagine S_2 and S_3 as two compa-

- nies that are in direct competition with each other, i.e. S_2 's loss (gain) is S_3 's gain (loss).
3. We recognize that one can consider a concept of possibility that admits possibilities with zero probability. However, such an account must provide some other conceptual role for the resulting notion of possibility. An important example of such an account is given by Levi's notion of 'serious possibility' (Levi, 1980). On Levi's account, serious possibilities can be assigned zero probability, but they have a status that is conceptually distinct from probability in that they emerge from the agent's commitments to a state of full belief.
 4. Sen (1970) allows an incomplete preference relation to serve as a basis for judgments of admissibility. Let R be a binary relation on a set X of alternatives. No assumption is made regarding the completeness of R on X . If P is the "asymmetric part" of R , i.e. xPy iff xRy but not yRx , then the "maximal" alternatives of $Y \subseteq X$ are those alternatives y in Y for which there is no x in Y such that xPy .

REFERENCES

- Arlo-Costa, H. and Helzner, J. (2005), Comparative ignorance and the ellsberg phenomenon, in Cozman, F., Nau, R. and Seidenfeld, T. (eds.), *Proceedings of the Fourth International Symposium on Imprecise Probabilities and their Applications*, pp. 21–30.
- Ellsberg, D. (1961), Risk, ambiguity, and the savage axioms, *The Quarterly Journal of Economics* 75, 643–669.
- Ellsberg, D. (2001), *Risk, Ambiguity and Decision*, Garland Publishing.
- Gardenfors, P. and Sahlin, N. (1982), Unreliable probabilities, risk taking, and decision making, *Synthese* 53, 361–386.
- Kahn, B. and Meyer, R. (1991), Consumer multiattribute judgments under attribute-weight uncertainty, *Journal of Consumer Research* 17, 508–522.
- Keeney, R. and Raiffa, H. (1993), *Decisions with Multiple Objectives*, Cambridge University. The Cambridge edition is a republication of the 1976 work.
- Keynes, J. (1921), *A Treatise on Probability*, MacMillan.
- Klein, J. (2007), New york city gifted and talented program test information 2007-2008: For students in pre-kindergarten, kindergarten, 1st

- grade and 2nd grade, Technical Report, New York City Department of Education.
- Knight, F. (1921), *Risk, Uncertainty and Profit*, Houghton-Mifflin.
- Kreps, D. (1988), *Notes on the Theory of Choice*, Westview.
- Kyburg, H. (1968), Bets and beliefs, *American Philosophical Quarterly* 5, 63–78.
- Levi, I. (1974), On indeterminate probabilities, *Journal of Philosophy* 71, 391–418.
- Levi, I. (1980), *The Enterprise of Knowledge*, MIT Press.
- Levi, I. (1986), The paradoxes of allais and ellisberg, *Economics and Philosophy* 2, 23–53.
- Luce, R. and Raiffa, H. (1989), *Games and Decisions: Introduction and Critical Survey*, Dover Publications. The Dover edition is a republication of the 1957 work.
- Ray, P. (1973), Independence of irrelevant alternatives, *Econometrica* 32, 987–991.
- Seidenfeld, T. (1988), Decision theory without ‘independence’ or without ‘ordering’: What is the difference?, *Economics and Philosophy* 4, 267–290.
- Sen, A. (1970), *Collective Choice and Social Welfare*, Holden Day.
- Sen, A. (1971), Choice functions and revealed preference, *Review of Economic Studies* 38(3), 307–317.
- Weber, M. (1985), A method of multiattribute decision making with incomplete information, *Management Science* 31(11), 1365–1371.
- Wilkie, W. and Pessemier, E. (1973), Issues in marketing’s use of multiattribute models, *Journal of Marketing Research* 10(4), 428–441.

Address for correspondence: Jeffrey Helzner, Department of Philosophy, Columbia University, New York, NY, USA.
E-mail: jh2239@columbia.edu