Nonstandard numbers for qualitative decision making *

(Extended Abstract)

Daniel Lehmann
Institute of Computer Science,
Hebrew University,
Jerusalem 91904, Israel
lehmann@cs.huji.ac.il

Abstract

The consideration of nonstandard models of the real numbers and the definition of a qualitative ordering on those models provides a generalization of the principle of maximization of expected utility. It enables the decider to assign probabilities of different orders of magnitude to different events or to assign utilities of different orders of magnitude to different outcomes. The properties of this generalized notion of rationality are studied in the frameworks proposed by von Neumann and Morgenstern and later by Anscombe and Aumann. It is characterized by an original weakening of their postulates in two different situations: nonstandard probabilities and standard utilities on one hand and standard probabilities and nonstandard utilities on the other hand. This weakening concerns both Independence and Continuity. It is orthogonal with the weakening proposed by lexicographic orderings.

*This work was partially supported by the Jean and Helene Alfassa fund for research in Artificial Intelligence and by grant 136/94-1 of the Israel Science Foundation on "New Perspectives on Nonmonotonic Reasoning".

1 Infinitesimal probabilities

Suppose you are considering playing dice. You have to choose between betting on six (*b6*) and betting on four (*b4*). The sums won are the same. You experiment with the dice and come to the conclusion that the chances of the dice falling on six are equal to those it falls on four. You conclude that you are indifferent between *b6* and *b4*. You are now offered a third bet (*e6*): you win if the dice falls on six or falls on one of the (twelve) edges. You ponder the chances of the dice falling on an edge and conclude that they are too small to make you prefer the bet *e6* to *b4*. You are indifferent between *b4* and *e6*. Since you believe in the transitivity of indifference, you conclude you are indifferent between *e6* and *b6*.

You now consider two more bets. In *e*, you win a large sum if the dice falls on any one of the edges. In *f*, you win the same sum if the dice falls on the edge that lies between face six and face five. You decide you prefer bet *e* to bet *f*.

According to the theory of rationality proposed by von Neumann and Morgenstern [19], you are irrational. If you are indifferent between *e6* and *b6*, it must be that you give subjective probability zero to the event of the dice falling on one of its edges. In this case, you *must* be indifferent between *e* and *f*. You may certainly be rational and indifferent between e and f, but must any rational decider be such? This work suggests the preferences above can be explained by choosing some number ϵ, positive but infinitesimally close to zero (as in Robinson's [15]), and assigning a subjective probability of $\frac{1-\epsilon}{6}$ to the dice falling on any of its faces and a probability of $\frac{\epsilon}{12}$ to the dice falling on any one of its edges. The numbers $\frac{1-\epsilon}{6}$ and $\frac{1-\epsilon}{6} + \epsilon$ are *qualitatively* (the term will be formally defined below) equivalent. The numbers ϵ and $\frac{\epsilon}{6}$ are *not* qualitatively equivalent.

2 Infinitesimal utilities

Here is a different example. Suppose you have to choose between two lotteries. In the first lottery you may win, with probability p, a week's vacation in Hawaii. With probability $1 - p$ you get nothing. In the second lottery you may win, with the *same* probability p, the same vacation in Hawaii, but, with probability $1 - p$ you get a consolation prize: a free copy of your favorite magazine. Since the free copy is preferred to nothing, von Neumann-

Morgenstern's independence postulate implies that lottery two is preferred to lottery one. But couldn't a rational decision maker be indifferent between the two lotteries? One, often, I think, buys a lottery ticket in a frame of mind focused on the big prize and not on the consolation prize. This behavior is by no means general, as attested by the fact lotteries often offer consolation prizes, but should a decision maker indifferent between the two lotteries be considered *irrational* in all situations? I think not.

A variation on this example considers also a third lottery, in which one wins a trip to Paris with the same probability p as above, and nothing with probability $1 - p$. Suppose you try to compare lotteries one and three. You ponder at length the advantages and disadvantages of the two vacation spots, and decide you are indifferent between the trip to Hawaii and the trip to Paris, all relevant considerations taken into account. You conclude that you are also indifferent between the lotteries one and three. The Independence axiom of von Neumann and Morgenstern implies that lottery two is preferred to lottery three. But it has been argued that it is quite unreasonable to expect the very slight improvement that lottery two presents over lottery one to overcome the lengthy and delicate deliberation that made you conclude that the trips to Hawaii and Paris are equivalent for you.

Similar examples have been put forward to argue that the indifference relation is not always transitive: lottery three is equivalent to lottery one and to lottery two, but lottery two is preferred to lottery one. The system presented in this paper endorses the transitivity of indifference, but allows a decision maker to be indifferent between lotteries one and two.

This work proposes to consider that the utility you attach to the trips to Paris and Hawaii are equal, say 1. The utility you attach to the free copy of your favorite magazine is ϵ, some positive number, infinitesimally close to zero. The utility of lottery one and of lottery three is p, that of lottery two is $p + (1 - p)\epsilon$, which is qualitatively equivalent to p. Nevertheless, the free copy of the magazine has utility ϵ, that is *not* qualitatively equivalent to zero.

3 Infinite utilities

Let us consider yet another situation. A patient has to choose between two options:

1. (option p) do nothing and die in a matter of weeks,

2. (option q) undergo surgery, the result of which, depending on some objective probabilities, may be a long and happy life (denoted by l) or an immediate death on the operation table (denoted by d). We shall denote by λ the probability of the surgery being successful, i.e., the probability of l. The probability of death on the operation table is, therefore, $1 - \lambda$.

In other terms, one has to compare *mixtures* p and $\lambda\, l + (1-\lambda)\, d$. Assuming one prefers l to p and prefers p to d, von Neumann and Morgenstern's postulates imply what we shall call property $\mathbf{P}$: there is a unique $\lambda \in]0,1[$ for which one is indifferent between p and $\lambda\, l + (1-\lambda)\, d$. For any $\mu > \lambda$, one prefers $\mu\, l + (1-\mu)\, d$ to p, and for all $\mu < \lambda$, one prefers p to $\mu\, l + (1-\mu)\, d$. If one thinks that a long and happy life is overwhelmingly preferable to p so as to make the distinction between p and d insignificant, i.e., that $\mu\, l + (1-\mu)\, d$ is preferable to p for any $\mu \in]0,1[$, property $\mathbf{P}$ fails and one deviates from von Neumann-Morgenstern's rationality postulates.

But would we really dismiss as *irrational* such a behavior, or such a preference? The consideration of an infinite utility for l explains the preference. Von Neumann-Morgenstern's point of view is perfectly acceptable and I have no criticism for someone who adheres to it and decides there is indeed some μ, very close to zero, perhaps, such that p is equivalent to $\mu\, l + (1-\mu)\, d$. My only claim is that someone who thinks $\mu\, l + (1-\mu)\, d$ is preferable to p for any $\mu \in]0,1[$ cannot be considered *irrational* outright.

An argument, very similar to the one just presented, for prefering a mixture $\mu w + (1-\mu) d$ to some p, for any $\mu \in]0,1[$, even though p is preferred to d, appears in Pascal's [14]. There w denotes eternal bliss (the reward of the believer if God exists), d denotes a life spent in error by a believer in a God that does not exist, and p denotes a life spent by a non-believer. This argument, known as Pascal's wager, is very well-known and the reader may find in a detailed discussion in [13].

A number of papers [5, 18, 2] discussed, in the setting of the St. Petersburg's paradox, the existence of unbounded utilities. In the last of these papers, Aumann argues very convincingly that utilities should be bounded. At first sight, one may think that infinite utilities imply unbounded utilities, and therefore Aumann argues also against infinite utilities, but this is not the case. His argument may be summarized in the following way: if utilities were unbounded, for any $\lambda \in]0,1[$ there would be a consequence (c) such that a lottery $\lambda\, c + (1-\lambda)\, d$ is prefered to a long and happy life (l). But this seems very unreasonable to Aumann. His argument is directed against *unbounded* utilities, but ineffective against *infinite* utilities. Certainly, no consequence is infinitely preferable to l and therefore, if there are infinite utilities, the utility of l is infinite. But there is absolutely no problem if one assumes that the utility of l is infinite and maybe even maximal (nothing is prefered to l). In this case, Aumann's argument disappears. The fact that l is infinitely prefered to some other consequence, d for example, or a sum of money, will influence the preferences of a decider between lotteries involving l and consequences such as d.

4 Background

Utility theory is discussed in the framework of [19, Chapter 3], see also [7]. Let $\mathcal{H}$ be a boolean algebra of subsets of X, and P a convex set of probability measures on $\mathcal{H}$.We assume that P is finitely generated. Convex means here that:

$$\forall p, q \in P,\ \forall \lambda \in]0,1[,\ \ \lambda p + (1-\lambda) q \in P.$$

Here $]0,1[$ denotes the open real interval in some model, maybe nonstandard, of the real numbers.

Von Neumann and Morgenstern have characterized the binary relations $>$ on P that can be defined by a linear functional u on P, when $]0,1[$ is the standard interval, in the following way:

$$\forall p, q \in P,\ \ p > q \ \Leftrightarrow\ u(p) > u(q). \tag{1}$$

In Equation 1, the functional u is a function from P to the standard set of real numbers $\mathbf{R}$ and the relation $>$ in the right hand side is the usual strict ordering on $\mathbf{R}$.

Their characterization is equivalent to the following, due to Jensen [8] (see [7, p. 1408]) three conditions, for all $p, q, r \in P$ and all $\lambda \in]0,1[$ (a weak order is an asymmetric and negatively transitive binary relation):

A1 $>$ on P is a weak order,

A2 $p > q \ \Rightarrow\ \lambda p + (1-\lambda) r > \lambda q + (1-\lambda) r$,

A3 $(p > q, q > r) \ \Rightarrow\ \exists \alpha, \beta \in]0,1[,$ such that

$$\alpha p + (1-\alpha)r > q > \beta p + (1-\beta)r.$$

The three conditions above are not the original postulates of von Neumann and Morgenstern, they are equivalent to them. They will be referred to, nevertheless, in this work, as von Neumann and Morgenstern's postulates. The purpose of this work is to generalize von Neumann and Morgenstern's characterization to deal with *qualitative* probabilities or with *qualitative* utilities. In the sequel, $p \geq q$ will denote $q \not> p$ and $p \sim q$ will denote the conjunction of $p \geq q$ and $q \geq p$.

5 Qualitative Decision Theory

Qualitative decision theory has been developed mostly in opposition to quantitative decision theory, stressing decision methods that do not satisfy von Neumann-Morgenstern's or Savage's [16] postulates, the postulates generally accepted for quantitative decision theory. The focus in qualitative decision theory has always been on methods and algorithms, more than on an axiomatic treatment ([4] is an exception).

A different approach is proposed here: qualitative and quantitative decision theories can be viewed as special cases of a unified general theory of decision that contains both. This unified theory is a generalization of the quantitative theory. The power of the generalization lies in the consideration of nonstandard models of the set of real numbers for utilities and a definition of indifference that neglects infinitesimally small differences. In this paper probabilities will always be standard. Some first results, for nonstandard probabilities and standard utilities have been presented in [11]. Preliminary ideas appeared in [12].

A well-established tradition in Decision Theory considers Expected Utility Maximization as the only rational policy. Following this view, an act f is strictly preferred to an act g iff the utility expected from f is strictly larger than that expected for g. Since expected utilities are real numbers, *strictly larger* has its usual, *quantitative* meaning. The main claim of this paper is that the qualitative point of view may be subsumed by a slightly different definition of *strictly larger*. Suppose we consider any model elementarily equivalent to the real numbers, more precisely, any (standard or nonstandard) model of the real numbers, $\mathcal{R}$ (for the standard model, we shall use $\mathbf{R}$). Let x and y be elements of $\mathcal{R}$. To make matters simpler, suppose that both x and

y are positive. The number x is quantitatively larger than y iff $x - y > 0$. What could be a reasonable definition of *qualitatively larger*? Clearly, if x is qualitatively larger than y then it must be quantitatively larger: in a sense *qualitatively larger* means *definitely larger*. A first idea that may be considered is to use a notion that proved fundamental for nonstandard analysis (the monads of [15], or see [10]): the notion of two numbers being *infinitely close*, and consider that a number x is *qualitatively* larger than a number y iff x is larger than y and *not infinitely close* to y, i.e., iff $x - y$ is strictly larger than some positive standard number. At first this idea looks appealing: if ϵ is strictly positive and infinitesimally close to zero, and x is a standard, strictly positive, real, we do *not* want $x + \epsilon$ to be qualitatively larger than x. At a second look, one realizes that the size of $x - y$ should not be judged absolutely, but relatively to the size of x: for example $\epsilon^2 + \epsilon$ should be qualitatively larger than ϵ^2. Therefore I propose the following definition:

Definition 1 *Let x and y be positive. We shall say that x is qualitatively larger than y and write $x \succ y$ iff $\frac{x-y}{x}$ is strictly positive and* not *infinitesimally close to zero: in other terms, iff there is a strictly positive* standard *number r such that $\frac{x-y}{x} \geq r$.*

The definition may be extended to arbitrary numbers in an obvious way:

1. if $x \geq 0$ and $0 > y$, then $x \succ y$, and

2. $x \succ y$ iff $-y \succ -x$,

$x \preceq y$ shall denote $x \not\succ y$ and $x \sim y$ shall denote that $x \preceq y$ and $y \preceq x$.

Notice that, if we choose, for $\mathcal{R}$, the standard model of the reals, $\mathbf{R}$, then $x \succ y$ iff $x > y$. Therefore our treatment would include the classical approach, if we allowed also negative utilities. As said above, in this paper, we concentrate on the case all utilities are positive. Is our framework, with positive nonstandard utilities, a generalization of the classical theory, with standard positive and negative utilities? Since, in the classical setting, utilities are defined only up to an additive constant, bounded utilities may always be considered to be positive, by adding a positive large enough constant. In view of Aumann's [2] critique of unbounded utilities, we feel that the present framework encompasses the most important part of classical theory.

Notice also that Definition 1 relies on the notion of a *nonstandard* number, and that notion is not first-order definable.

Expected Qualitative Utility Maximization, the paradigm of rationality proposed here is the version of Expected Utility Maximization that obtains when, for probabilities and utilities,

- the models chosen for the real numbers *may* be nonstandard, and
- real numbers are compared *qualitatively*, i.e., by $\succ$.

At this stage, I do not know of an axiomatic characterization of Expected Qualitative Utility Maximization in its most general form: nonstandard probabilities and utilities. But two orthogonal special cases have been characterized in full: first, the case in which probabilities may be nonstandard but utilities are standard and secondly, the case in which probabilities are standard but utilities may be nonstandard. In the first case, we want to characterize the binary relations $>$ on P that can be defined by a functional $u : P \rightarrow \mathbf{R}$ into the standard real numbers in the following way:

$$\forall p, q \in P, p > q \Leftrightarrow u(p) > u(q), \tag{2}$$

and u is pseudo-linear, i.e.:

$$\forall p, q \in P, \forall \lambda \in]0,1[, u(\lambda p + (1-\lambda)q) \sim \lambda u(p) + (1-\lambda)u(q). \tag{3}$$

In Equation 3, the interval $]0,1[$ may be non-standard and therefore the right-hand side of the equivalence may be nonstandard. This case is treated in Section 7.

In the second case, we characterize the binary relations $>$ on P that can be defined by a *linear* functional $u : P \rightarrow \mathcal{R}_+$ in the following way:

$$\forall p, q \in P, p > q \Leftrightarrow u(p) \succ u(q). \tag{4}$$

Here the interval $]0,1[$ is the standard one. This case in treated in Section 8.

One should immediately notice that, if $c > 0$, the utility function $c\, u(p)$ defines the same ordering as $u(p)$, and is linear or pseudo-linear iff u is. But, contrary to what happens in the classical setting, if $d \in \mathcal{R}$ the function $d + u(x)$ does not, in general, define the same ordering as $u(p)$. Such an instability under an additive constant, and in particular an asymmetry between gains and losses has been found in the behavior of decision makers in many instances [6, 9, 17]. The question of whether Expected Qualitative Utility Maximization is a realistic model for explaining such behavior cannot be discussed in this work.

6 Maximin as Expected Qualitative Utility Maximization

Considering nonstandard utilities enables us to obtain decision criteria that do not satisfy von Neumann-Morgenstern's postulates and were so far considered as part of the realm of *qualitative* decision theory.

As noticed above, Expected Qualitative Utility Maximization generalizes Expected Utility Maximization: if one chooses the standard model for real numbers then Expected Qualitative Utility Maximization boils down exactly to Expected Utility Maximization, at least when utilities are bounded. We shall show now that considering nonstandard utilities enables us to obtain decision criteria that do not satisfy von Neumann-Morgenstern's postulates and were so far considered as part of the realm of *qualitative* decision theory. A version of the Maximin criterion will be presented. The Maximin criterion has been proposed by A. Wald [20], in a different framework. The criterion to be presented is a variation on this theme.

Assume the set X is finite and $\mathcal{H}$ contains all subsets of X. Let the elements of X be $x_0, \ldots, x_{n-1}$. Let ϵ be a number that is positive and infinitesimally close to zero and let our utility function u be the linear function defined by: $u(x_i) = \epsilon^{n-i-1}$, for $i = 0, \ldots, n-1$. Notice that $x_i < x_j$ iff $i < j$. The utility of a mixture $\lambda x_i + (1-\lambda) x_j$ is $\lambda \epsilon^{n-i-1}$ if $x_i < x_j$ and ϵ^{n-i-1} if $x_i \sim x_j$.

Suppose $i < j$ and $i' < j'$. Then, $\lambda x_i + (1-\lambda)x_j < \mu x_{i'} + (1-\mu)x_{j'}$ iff $x_i < x_{i'}$ or $x_i \sim x_{i'}$ and $\lambda > \mu$. The decision maker therefore compares different mixtures by comparing the worst possible outcomes and, if they are the same, their respective probabilities. This is some form of Maximin criterion and does not satisfy **A2** or **A3**, but it has been considered a rational way of deciding by many authors, and it is amenable to Expected Qualitative Utility Maximization.

7 Postulates for nonstandard probabilities and standard utilities

The postulates that characterize this first case are **A1**, **A3** and the following **B2**.

Definition 2 $\lambda \in]0,1[$ *is* negligible *iff, for any* $p, q \in P$, $\lambda\, p + (1-\lambda)\, q \sim q$.

The intuitive meaning of *negligible* is infinitesimally close to zero.

B2 $p > q, \lambda$ not negligible $\Rightarrow \lambda p + (1-\lambda)r > \lambda q + (1-\lambda)r$.

8 Postulates for standard probabilities and nonstandard utilities

The postulates that characterize this second case are **A1** and the following.

To formulate our independence property, it is best to set the following definition.

Definition 3 *We shall say that p overrides q and write $p \gg q$ iff $p > q$ and for any q' such that $q > q'$ and for any $\lambda \in]0,1[$, $\lambda q + (1-\lambda)p \sim \lambda q' + (1-\lambda)p$.*

The intuitive meaning of $p \gg q$ is that p is so much preferred to q that, in any lottery in which p and q are the prizes, if one does not win p, one does not even care to cash q, but would as well get any lesser prize q'. Notice that, since $>$ is asymmetric, the relation $\gg$ is also asymmetric and therefore irreflexive.

Our independence property may now be formulated as:

A′2 $p > q, r \not\gg p \Rightarrow \forall \lambda \in]0,1[\ \lambda p + (1-\lambda)r > \lambda q + (1-\lambda)r$.

The intuitive meaning of **A'2** is that any lottery is sensitive to both its prizes, unless one of the prizes overrides the other one.

A′3 $p > q > r \Rightarrow \exists \alpha \in]0,1[$ such that $\alpha p + (1-\alpha)r > q$.

A″3 $p > q > r, p \not\gg q \Rightarrow \exists \beta \in]0,1[$ such that $q > \beta p + (1-\beta)r$.

9 Comparison with previous work

Numerous works during the fifties and the sixties considered weakenings of the von Neumann-Morgenstern's postulates. Nonstandard analysis [15] appeared late on the scene. This work proposes an original weakening based on nonstandard analysis.

Our postulates are very close to the original postulates of von Neumann and Morgenstern. In particular the ordering $<$ is modular (weak total) and the indifference relation $\sim$ is transitive. In the case of nonstandard utilities, we weaken both **A2** and **A3**, in a closely linked manner. Notice that **A2**, **A'3** and **A"3** together imply **A3**, since **A2** says that $p \gg q$ implies that for any w, $w \geq q$. The lexicographic orderings of [3] provide one of the best known weakenings of von Neumann and Morgenstern's postulates. The weakening they present is essentially orthogonal to ours. Indeed, the lexicographic orderings define a preference relation that satisfies **A2**. Any ordering that satisfies our postulates and those of von Neumann-Morgenstern postulates. To explain better the difference between lexicographic orderings and our qualitative ordering (in the case probabilities are standard), assume P is the real plane $\mathbf{R}^2$ ordered by the lexicographic ordering: $(x, y) < (x', y')$ iff either $x < x'$ or $x = x'$ and $y < y'$. For $\lambda \in [0, 1]$, define $\lambda(x, y) + (1 - \lambda)(x', y')$ to be $(\lambda x + (1 - \lambda)x', \lambda y + (1 - \lambda)y')$. Notice that $(0, 0) < (1, 10) < (2, 0)$, but there is no $\lambda \in]0, 1[$ such that $(1, 10) = \lambda(0, 0) + (1 - \lambda)(2, 0)$ since $\lambda(0, 0) + (1 - \lambda)(2, 0) = (2(1 - \lambda), 0)$. Both lexicographic and qualitative orderings imply the failure of property **P** of section 3. But lexicographic orderings also implies the failure of the following property that holds in the case probabilities are standard: if $p > q > r$ and there exists some $\beta \in]0, 1[$ such that $q > \beta p + (1 - \beta)r$, then there exists some $\gamma \in]0, 1[$ such that $q \sim \gamma p + (1 - \gamma)r$. Indeed there exists some $\beta \in]0, 1[$ such that $(1, 10) > \beta(2, 0) + (1 - \beta)(0, 0) = \beta(2, 0)$: for example $\beta = 0.4$, and nevertheless there is no λ as above. Lexicographic and qualitative orderings stem from different concerns and have very different characteristics.

10 Subjective probability

10.1 Anscombe-Aumann's framework

In [1], Anscombe and Aumann consider a finite set S (of states) and the set $\mathbf{F}$ (of acts) of mappings: $S \mapsto P$. They show that a single postulate , added to **A1**, **A2** and **A3** is enough to characterize the orderings obtainable from subjective probabilities on states and linear utilities: for any $a, b \in \mathbf{F}$,

A4 If $\forall s \in S, s \neq s_0 \Rightarrow a(s) = b(s)$, then $a > b \;\Rightarrow\; a(s_0) > b(s_0)$.

In the last part of **A4**, $a(s_0)$ and $b(s_0)$ stand for the corresponding constant functions.

For the case of standard utilities (and nonstandard probabilities), the same single added postulate is enough to guarantee the corresponding result: existence of nonstandard subjective probabilities and a pseudo-linear utility function into the standard reals.

For the case of standard probabilities (and nonstandard utilities), in addition to **A1**, **A'2**, **A'3**, **A"3** and **A4**, one needs an additional postulate to ensure the subjective probabilities are standard. This postulate deals with Savage-null states.

$$\mathbf{A'5} \quad t \in S,\ a \in \mathbf{F},\ a(t) \gg a \ \Rightarrow\ t \text{ is null.}$$

where null is defined below and $\gg$ is defined in Definition 3.

Definition 4 *Let $t \in S$. The element t is said to be* null *iff for any $a, b \in \mathbf{F}$ we have $a \sim b$ if a and b agree everywhere except possibly on t, i.e., for any $s \neq t$, $a(s) = b(s)$.*

11 Conclusion

Nonstandard models of the real numbers provide for a natural notion of qualitative equivalence and a principle of Qualitative Utility Maximization. This work characterizes in full the situation in which one allows nonstandard utilities but insists on standard probabilities. In this framework one may consider consequences that are infinitely preferable to others and criteria of the Maximin family. The study of games with nonstandard utilities seems appealing. The dual case of nonstandard probabilities and standard utilities and the most general case of nonstandard probabilities and utilities are yet to be characterized. They will include consideration of subjective probabilities infinitesimally close to zero.

Acknowledgments

I want to thank Israel Aumann for his encouragements and suggestions to improve the presentation of this work.

References

[1] F. J. Anscombe and R. J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205, 1963.

[2] Robert J. Aumann. The St. Petersburg paradox: A discussion of some recent comments. *Journal of Economic Theory*, 14(2):443–445, April 1977.

[3] Lawrence Blume, Adam Brandenburger, and Eddie Dekel. Lexicographic probabilities and choice under uncertainty. *Econometrica*, 59(1):61–79, January 1991.

[4] Ronen I. Brafman and Moshe Tennenholtz. On the foundations of qualitative decision theory. In *Proceedings of the National Conference on AI (AAAI'96)*, pages 1291–1296, 1996.

[5] D. L. Brito. Becker's theory of the allocation of time and the St. Petersburg paradox. *Journal of Economic Theory*, 10:123–126, 1977.

[6] P.C. Fishburn and G.A. Kochenberger. Two-piece von neumann-morgenstern utility functions. *Decision Sciences*, 10:503–518, 1979.

[7] Peter C. Fishburn. Utility and subjective probability. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 2 of *Handbooks in Economics 11*, chapter 39, pages 1397–1435. North-Holland, 1994.

[8] N. E. Jensen. An introduction to bernoullian utility theory, I: Utility functions. *Swedish Journal of Economics*, 69:163–183, 1967.

[9] D. Kahneman and A. Tversky. Prospect theory: an analysis of decision under risk. *Econometrica*, 47:263–291, 1979.

[10] H. J. Keisler. *Foundations of Infinitesimal Calculus.* Prindle, Weber&Schmidt Inc., Boston, 1976.

[11] Daniel Lehmann. Generalized qualitative probability: Savage revisited. In Eric Horvitz and Finn Jensen, editors, *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*, pages 381–388, Portland, Oregon, August 1996. Morgan Kaufmann.

[12] Daniel Lehmann. What is qualitative: A framework for qantitative and qualitative decision theory. In *Qualitative Preferences in Deliberation and Practical Reasoning*, Stanford, CA, March 1997. 1997 AAAI Spring Symposium.

[13] Daniel Lehmann. Nonstandard utilities. Technical Report TR-98-3, Leibniz Center for Computer Science, Computer Science Institute, Hebrew University, Jerusalem, Israel, January 1998. revised version.

[14] Blaise Pascal. *Pensées*, volume 28 of *Les Cents Chefs d'Œuvre de la Langue Française*. Robert Laffont, Paris, France, 1960. first published in 1670.

[15] Abraham Robinson. *Non-standard Analysis*. North-Holland, Amsterdam, 1966.

[16] Leonard J. Savage. *The Foundations of Statistics*. John Wiley, 1954.

[17] P.J.H. Schoemaker. *Experiments On Decision Under Risk*. Martinus Nijhoff, Boston, 1980.

[18] L.S. Shapley. The St. Petersburg paradox: A con game? *Journal of Economic Theory*, 14:439–442, 1977.

[19] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, second edition edition, 1947.

[20] Abraham Wald. *Statistical Decision Functions*. John Wiley & Sons, New York, 1950.