

Medical Decision Making

<http://mdm.sagepub.com/>

Why Expected Utility Theory Is Normative, but Not Prescriptive

Jonathan Baron

Med Decis Making 1996 16: 7

DOI: 10.1177/0272989X9601600103

The online version of this article can be found at:

<http://mdm.sagepub.com/content/16/1/7.2>

Published by:



<http://www.sagepublications.com>

On behalf of:



[Society for Medical Decision Making](#)

Additional services and information for *Medical Decision Making* can be found at:

Email Alerts: <http://mdm.sagepub.com/cgi/alerts>

Subscriptions: <http://mdm.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://mdm.sagepub.com/content/16/1/7.2.refs.html>

>> [Version of Record](#) - Feb 1, 1996

[What is This?](#)

The Normative Status of Expected Utility Theory

A Word from the Editor

In this issue, we present a new feature. Several comments, representing a variety of views, were solicited as responses to a paper that took what is, for this journal, an unusual direction: Brian Cohen's paper is neither a report of empirical research on clinical decision making nor the working through of a mathematical model of a clinical problem. It is a philosophical analysis of some of the theoretical foundations of these more frequently appearing contributions. The commentators all saw this paper prior to publication, and Cohen responds to their views.

Members of the Society for Medical Decision Making and many contributors to the literature on clinical applications of decision analysis have recurrently lamented the failure of the approach to decision making at issue—expected utility theory—to gain wide ac-

ceptance among clinicians. Cohen raises the question of whether decision analyses, by virtue of their foundation in expected utility theory, implicitly commit clinicians to maximizing collective, as opposed to individual, good. The builders of the modern version of expected utility theory attempted to sever all links between it and nineteenth-century British utilitarianism. Cohen argues that they did not succeed. The continuing debate about how to balance individual and societal interests in medical decisions suggests that we still have a way to go before the issues are resolved, either in theory or in implementing public policy.

This forum is intended to provoke dialogue and discussion. Readers' comments and reactions are most welcome.—A.S.E.

Why Expected Utility Theory Is Normative, but Not Prescriptive

JONATHAN BARON, PhD

Brian Cohen¹ argues, and I agree, that expected utility (EU) theory needs more justification than is typically provided, that the usual "intuitive-appeal" defenses of independence and transitivity are unconvincing, that an appropriate defense of EU must be based on goal achievement, and that many applications of EU to medical decision analysis are questionable. But he dismisses the axiomatic argument too quickly.* I here summarize my defense of the independence principle—apparently the most controversial axiom—in terms of goal achievement (Baron², pp 297–298; later stated better³, pp 322–332, 4, ch. 4) and comment on the implica-

tions of accepting EU for medical decision analysis.

It is important to point out that a normative theory is an idealized abstraction, an analytic framework that we impose on the disorderly flow of existence for the sake of defining a standard. The theory need not be applicable directly. It may instead allow us to evaluate the rules we use in practice. Rules that are better than others according to the standard are called "prescriptive."⁵ A normative theory of decision making may be based on maximizing the achievement of goals (or the satisfaction of values). The analytic framework assumed by EU involves thinking about decisions in terms of uncertain states of the world, about which we have beliefs; outcomes, for which we have goals or values; and options, which determine outcomes given states. The point is to evaluate options in terms of goals for outcomes and beliefs about states. This evaluation should depend on the world itself, not on which of several equivalent descriptions we choose.

Address correspondence to Dr. Baron: Department of Psychology, University of Pennsylvania, 3815 Walnut Street, Philadelphia, PA 19104-6196; or (e-mail) baron@psych.upenn.edu.

*Brevity required him to ignore many other issues, and I shall follow him here. However, his discussion of my views neglects, among other things, my defense of the axiomatic argument.

Maximization of goal achievement implies independence. To show this, I consider a different form of the independence principle, which says: *If there is some state of the world that leads to the same outcome no matter what choice you make, then your choice should not depend on what that outcome is.* Consider the following pair of choices:

	S_1		S_0
	$p = 0.05$	$p = 0.20$	$p = 0.75$
A vs B	\$3,000 \$0	\$3,000 \$4,000	\$0 \$0
C vs D	\$3,000 \$0	\$3,000 \$4,000	\$3,000 \$3,000

Your choice of A vs B is that of Cohen's lotteries II vs IV. In C and D, the \$0 outcome is replaced with \$3,000. We can think of these lotteries as arising from two states of the world S_1 and S_0 . S_1 leads to Cohen's lottery I for choice A or C and to Lottery III for choice B or D. S_0 leads to \$0 for A and B, and \$3,000 for C and D, so its outcome is out of your control.

Many people are tempted to choose C in the C–D choice and B in the A–B choice. Independence implies that you should choose either A and C or B and D (or that you should be indifferent everywhere). In other words, the outcome in S_0 doesn't matter. To see this, assume:

1. The relation "is better in terms of your goals" (abbreviated $>_g$) applies to lotteries considered as outcomes. Thus, we may think of "getting a ticket for (or the right to play) a lottery with a 0.8 chance of \$4,000 and a 0.2 chance of \$0" as the outcome of S_1 for options B or D. Likewise, we may think of the outcome "\$3,000" as a lottery. Getting \$3,000, after all, is an abstraction from an infinitude of possible courses of events—getting \$3,000 and subsequently losing it through a bad investment, getting \$3,000 and buying a new computer, getting \$3,000 and hearing simultaneously that the value of one's stocks have declined (or increased) by the same amount, etc. If we were to assume that our values led to a transitive ordering of things like "\$3,000" but not of things like "a 0.8 chance of \$4,000," then the application of the theory would depend on how we describe equivalent outcomes, not on the world itself. Moreover, given that we can describe anything as a lottery, denial of this assumption would amount to denial of the entire enterprise. Neither EU nor any other theory could evaluate options.

2. Outcomes described in the same way (e.g., "\$3,000" or "0.8 chance of \$4,000") are evaluated the same, regardless of the existence of outcomes that could have occurred but did not. This assumption—necessary because it is part of the antecedent condition of the independence principle itself—is typically false for

lotteries. The feeling of winning or losing is part of the outcome, and this feeling clearly depends on such counterfactuals. To aid your imagination in thinking of a case where this assumption is plausibly true, assume that the outcome of the chosen option (e.g., "\$3,000" or a ticket for a 0.8 chance of "\$4,000") will be sent as anonymous donations to your favorite charity, that the recipients will not know anything about how the result was reached or what else it might have been, and that you care only about the benefits to the charity, not your own feelings.

Assumption 2 implies: *events or outcomes that do not happen cannot affect the achievement of goals.* Thus, two identically-described outcomes (including lotteries) are equal in goal achievement no matter how you reach them, once you reach them.[†]

There are two possible states, S_1 and S_0 . If S_1 occurs, then the outcome in S_0 does not matter—does not affect whether $C >_g D$ or $D >_g C$ —because S_0 did not happen (assumption 2). If S_0 occurs, then the outcome in S_0 does not matter either, because the outcome is the same (assumption 2). Therefore, the outcome in S_0 does not matter. Assumption 1 implies that either $A >_g B$ or not. Since the outcome of S_0 does not matter, if $A >_g B$ then $C >_g D$, and if not $A >_g B$ then not $C >_g D$. If you want to maximize achievement of your goals, you must follow the independence principle.

This principle, plus other principles, implies EU. Some of the common violations of EU theory, such as the tendency to choose B and C, violate the independence principle directly, so the principle is not descriptively true, just normatively correct. That is, goal maximization requires that we follow it.

Expected utility theory helps to make an argument for utilitarianism.^{4,9} When several individuals face the same decision under uncertainty, as is typical for medical decisions, it is impossible for each individual to follow EU yet for the overall outcome to be non-utilitarian.

Expected utility is not always prescriptively correct.³ Normative theories provide a standard for evaluating decisions or methods of making them, but the attempt to follow this standard is often self-defeating. We may do better *according to the standard* if we use various heuristics or alternative formal methods instead. Or, it may be that *some* ways of trying to follow the standard are better than others.

One case of this sort may be the use of "standard gambles" (no pun intended) in decision analysis. People choose lottery C because they feel certain of getting \$3,000. This kind of "certainty effect" is one of several factors that make responses to lotteries violate EU theory.¹⁰ The problem seems to be in the way we think

[†]This is the essential idea of the dynamic-consistency argument of Hammond⁶ (defended against the criticisms of McClennan⁷ by Haslam and Baron.⁸

about probabilities themselves. Standard gambles lead to inconsistent estimates of utilities of identical outcomes. Other methods, particularly difference measurement, may be better.^{3, ch 17}

More generally, the goal-achievement definition of utility implies that true utility is not necessarily the same as utility inferred from preferences. Resistance to the goal-achievement view comes from the behaviorist beast, which, having largely moved out of the house of psychology, has taken up residence in economics.

References

1. Cohen B. Is expected utility theory normative for medical decision making? *Med Decis Making*. 1996;16:1–6.
2. Baron J. *Thinking and Deciding*. New York: Cambridge University Press, 1988.
3. Baron J. *Thinking and Deciding*. 2nd ed. New York: Cambridge University Press, 1994.
4. Baron J. *Morality and Rational Choice*. Dordrecht, The Netherlands: Kluwer, 1993.
5. Baron J. *Rationality and Intelligence*. New York: Cambridge University Press, 1985.
6. Hammond PH. Consequentialist foundations for expected utility. *Theory and Decision*. 1988;25:25–78.
7. McClennan EF. *Rationality and Dynamic Choice: Foundational Explorations*. New York: Cambridge University Press, 1990.
8. Haslam N, Baron J. Rationality and resoluteness. Review of *Rationality and Dynamic Choice: Foundational Explorations*, by EF McClennan. *J Math Psychol*. 1993;37:143–53.
9. Broome J. *Weighing Goods: Equality, Uncertainty and Time*. Oxford, England: Basil Blackwell, 1991.
10. Kahneman D, Tversky A. Prospect theory: an analysis of decisions under risk. *Econometrica*. 1979;47:263–91.

The Strengths and Limitations of Expected Utility Theory

GEORGE WU, PhD

Expected utility (EU) theory has long served as the backbone for decision analysis. The theory has been applied to decision problems in management, public policy, and, most importantly for readers of this journal, medicine. The dominant status enjoyed by EU theory reflects a presumed *normative* justification: it is often argued that rational, reflective individuals would want to adhere to the basic principles of EU theory. In this issue, Cohen argues that EU theory is inappropriate for medical decisions because these are one-time decisions and expectation is a long-run property that applies only to repetitive decisions.¹ The question of whether EU is an appropriate normative standard is vitally important for medical decision making. However, the arguments Cohen offers are not convincing enough to abandon EU theory in medical or other contexts. The theory does have limitations, but before we abandon a proven theory we must propose alternative theories and scrutinize them as closely as we have scrutinized EU theory.

The Importance of Axioms

Expected utility is implied and implies a set of axioms: ordering, continuity, independence, and transitivity.

Building a normative theory of decision making from axioms has great merit. Axioms impose a consistency between actual and hypothetical choices. Cohen is correct that “intuitive rules of consistency alone cannot serve as a sufficient basis for rationality.” He seems, however, to imply that proponents of EU theory have taken consistency to be sufficient justification for it. On the contrary, EU theory has been defended in a number of ways. The first and most natural support stems from the intuitive appeal of an axiom, the extent to which it provides a reasonable guide for individual behavior.^{2,3} Axioms decompose a larger theory into smaller pieces that can be separately judged as normative principles. For many students of decision theory, the independence axiom survives as a normative principle despite the substantial psychological pull of EU “counter-examples” such as the Allais paradox.⁴ Second, EU theory supporters have enumerated many implications of violating EU, such as dynamically inconsistent behavior, vulnerability to a money pump, a negative value of information, and “horizon inflexibility” (the optimal choice depends on how far the decision tree is carried out).^{5–8}

Central to Cohen's thesis is the argument that the normative status of EU can be judged only through the lens of a broader vision of rationality, that “the value of rationality is its ability to help people further their own interests.” Goal achievement seems to be indisputable as a principle of rationality. However, it

Address correspondence to Dr. Wu: Graduate School of Business Administration, Harvard University, Boston, MA 02163.