Brian Skyrms

The Value of Knowledge

Why is it better to know something rather than nothing? Perhaps because knowledge is an end in itself. But we also seek knowledge in order to make informed decisions. Informed decisions, we believe, are better than uninformed ones. What is the relevant sense of "better" and what basis is there for the belief? There is a Bayesian answer to these questions. It is that coherence *requires* you to believe in the value of knowledge.

It is a theorem in the context of the decision theory of L. J. Savage (1954) and a standard treatment of learning by conditionalization that the prior expected utility of making an informed decision is always at least as great as that of making an uninformed decision, and is strictly greater if there is any probability that the information may affect the decision. The proof is stated in a characteristically trenchant note by I. J. Good (1967).

This paper will investigate the extent to which this account can be generalized to various models of the acquisition of information and to variant theories of expected utility, its import for the theory of deliberation, and its connection with payoff in the long run.

I. Good Thinking

Suppose that one can either choose between n acts, $A_1, \ldots, A_n$ now or perform a cost-free experiment, E, with possible results $[e_k]$ and then decide. The value of choosing now is the expected value of the act with the greatest expected utility (the prior Bayes act):

$$\begin{aligned} & MAX_j \, \Sigma_i \, Pr(K_i) \, U(A_j \,\&\, K_i) \\ = \; & MAX_j \, \Sigma_k \, \Sigma_i \, Pr(K_i) \, Pr(e_k | K_i) \, U(A_j \,\&\, K_i). \end{aligned}$$

The value of making an informed decision conditional on experimental result e is the expected utility conditional on e of the act that has the highest expected utility after assimilating the information e (the posterior Bayes act associated with e):

$$MAX_j \, \Sigma_i \, Pr(K_i | e) \, U(A_j \,\&\, K_i).$$

The present value of making an informed decision is the expectation:

$$\begin{aligned} & \Sigma_k \Pr(e_k) \text{ MAX}_j\ \Sigma_i \Pr(K_i \mid e_k)\ U(A_j \ \&\ K_i) \\ = \quad & \Sigma_k \Pr(e_k) \text{ MAX}_j\ \Sigma_i \Pr(e_k \mid K_i) \Pr(K_i)/\Pr(e_k)\ U(A_j \ \&\ K_i) \\ & \text{(by Bayes's Theorem)} \\ = \quad & \Sigma_k \text{ MAX}_j\ \Sigma_i \Pr(e_k \mid K_i) \Pr(K_i)\ U(A_j \ \&\ K_i). \end{aligned}$$

The formulas for the present value of making an informed decision and an uninformed decision differ only in the order of the first two operations. But it is true on general mathematical grounds that Σ_k MAX$_j$ g(k,j) is greater than or equal to MAX$_j$ Σ_k g(k,j), with strict inequality if MAX$_j$ g(k,j) is not the same for all k, q.e.d.

The situation is easy to visualize when there are only two possible experimental results: e_1,e_2. Suppose that there are three possible acts: A_1,A_2,A_3, whose expected utility is graphed as a function of Pr(e_2) in figure 1.

If the experiment is performed and e_2 is the result, then Pr(e_2) will equal 1, and A_1 will have the greatest expected utility. If the experiment is performed and e_1 is the result, then Pr(e_2) will equal zero and A_2 will be the optimal act. If prior to the experiment Pr(e_2) = 1/3, then A_3 will be the prior Bayes act. The expected utility of the prior Bayes act is thus the convex function indicated by the bold line in figure 2:

The expected utility of an informed decision is plotted as the dotted line connecting the expected utility of A_1 at Pr(e_2) = 1 with the expected utility of A_2 at Pr(e_2) = 0. The vertical difference between the bold and dotted lines is the net gain in prior expected utility resulting from the determination to make an informed decision. This is zero only if both experimental outcomes lead to the same posterior Bayes act or if one is already certain of the experimental outcome; otherwise it is positive.

I want to take explicit notice here of some of the features of the system within which the proof proceeds. In the first place the expected utility theory assumed is statistical decision theory as found in Savage (1954), where a distinction is made between states of the world, acts, and consequences; where states of the world together with acts determine the consequences; and where the relevant expectation is unconditional expected utility:

$$\Sigma_i \Pr(K_i)\ U(A \ \&\ K_i).$$

We will see in section V that in the decision theory of Jeffrey (1965), where the value of an act goes by its conditional expected utility, Good's theorem fails.

Secondly, by using the same notation for acts and states pre- and postexperiment, we are assuming that performance of the experiment itself cannot affect the state in any relevant way and that, so far as they affect consequences, the generic acts available postexperiment are equivalent to those available preexperiment. This assumption is violated, for example, in an Abbott and Costello movie where

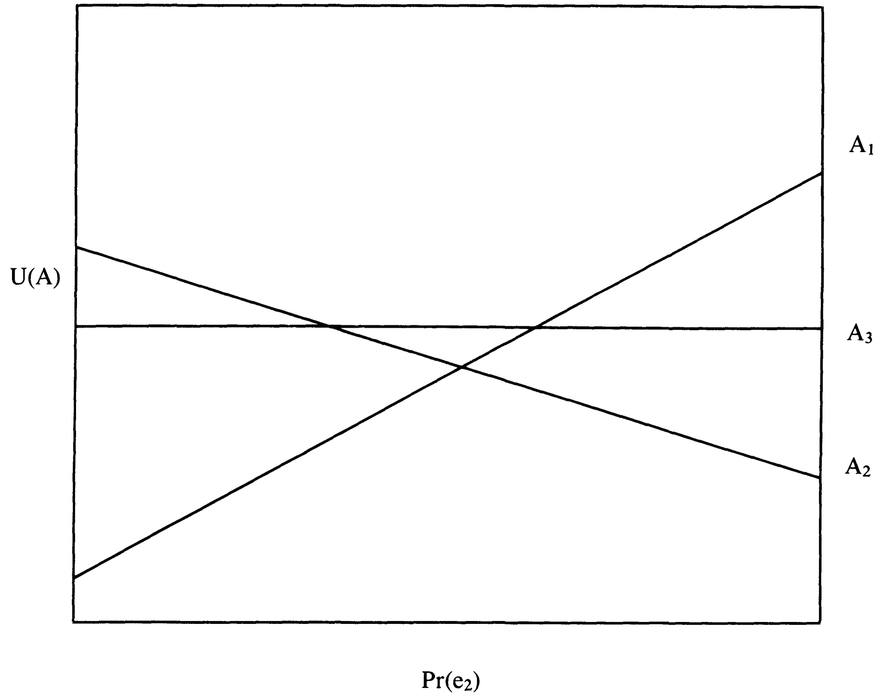

Figure 1. Expected Utility as a Function of Pr (e_2).

Costello is charged with keeping a match dry until a crucial time. Abbott keeps asking Costello if the match is dry and Costello keeps replying "Yes." Finally it is time to light the match. "Are you sure it's dry?" Abbott asks for the final time. "Yes, I lit it a little while ago to make sure," Costello replies.

In the third place, the proof implicitly assumes not only that the decision maker is a Bayesian but also that he knows that he will act as one. The decision maker believes with probability one that if he performs the experiment he will (i) update by conditionalization and (ii) choose the posterior Bayes act. For an example where (i) fails, consider the case of the wrath of Khan. You are a prisoner of Khan, who (you firmly believe) is about to insert a mindworm into your brain, which will cause you to update by conditionalization on the denial of the experimental result. For an example where (ii) fails, consider the compulsive loser who conditionalizes correctly on the experimental result, but then proceeds to choose the act with minimum posterior expected utility.

I am not taking issue here with any of the assumptions, but merely making them explicit. As it stands, the theorem is a beautifully lucid answer to a fun-

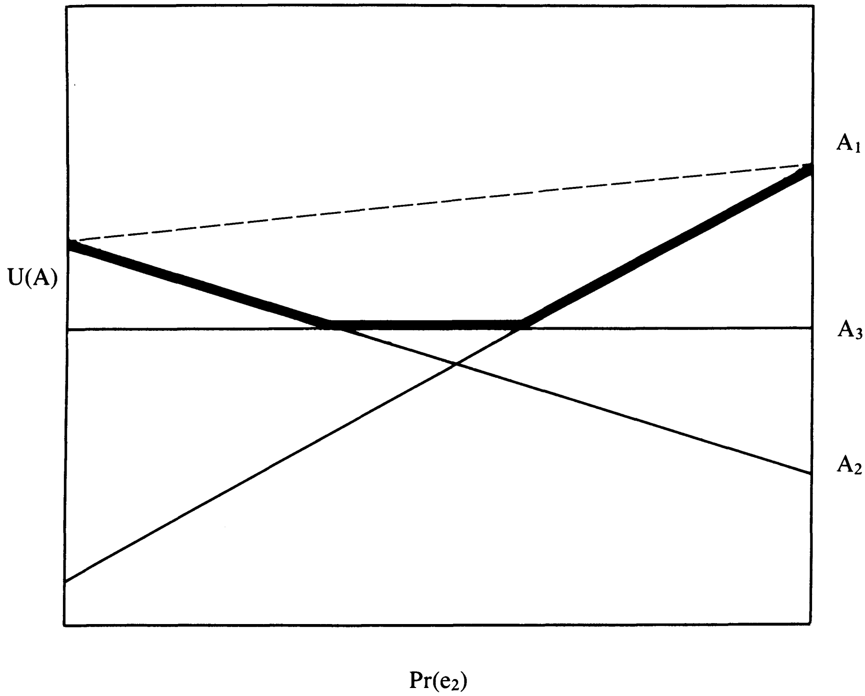

Figure 2. The Value of Information.

damental epistemological question. It will be of some interest, however, to see to what extent the assumptions can be relaxed and the theorem retained.

II. Probable Knowledge

Does knowledge always come nicely packaged as a proposition in one's subjective probability space? To attempt an affirmative answer would be either to defend a form of "the myth of the given" of unprecedented strength, or to relapse into skepticism. But the standard Bayesian theory of learning from experience by conditionalization and, in particular, the analysis of the last section appear to tacitly make just this assumption. This is not because Bayesians have been ignorant of the problem, but rather because it is much easier to raise the difficulty than to suggest any constructive treatment.

One positive suggestion put forward by Richard Jeffrey (1965, 1968) is to generalize Bayes's rule of conditionalization. Suppose that an observational interaction falls short of making proposition p certain, but makes it highly probable. Suppose also that the only effect on the observer's subjective probability space is through the effect on p; i.e., the probabilities conditional on p and on its negation remain constant. Then we have what Jeffrey calls belief change by probability kinematics on the partition [p,-p]. Conditionalization on p and on its negation are extreme cases. Jeffrey extends the notion to any partition, all of whose members have positive probability in the obvious way. We have belief change by probability kinematics on that partition, just in case posterior probabilities conditional on members of the partition (where defined) remain unchanged.

Does the analysis of the last section extend to the case of probable knowledge? Paul Graves has recently shown that it does (Graves, forthcoming). Here is a sketch of Graves's analysis.

Suppose that you can either choose now among n acts, or perform a cost-free experiment whose result will be a belief change by probability kinematics on partition Γ. There may be no proposition in your language capturing just the phenomonological "feel" of the possible observational inputs. But are there entertainable propositions that capture the possible effects on your probability space of the observational interaction? Well, you could just *describe* the possible final probability measures that could come about. There is no reason why you could not think about these possible outcomes now, expanding your probability space to allow final probability, pr_f,to enter as a random variable.

You believe now that your observational interaction is a legitimate way of acquiring information, and so you have now:

$$\text{(M)} \quad PR(q \mid pr_f = pr^*) = pr^*(q).$$

You believe that your belief change will be by probability kinematics on partition Γ, so for any final probability pr^*, any proposition q that is "first order" (i.e., does not involve pr_f), and any member γ of Γ, you have:

$$\text{(PK)} \quad pr^*(q \mid \gamma) = PR(q \mid \gamma),$$

from the definition of probability kinematics on Γ. By (M) this is equivalent to:

$$\text{(S)} \quad PR(q \mid pr_f = pr^* \ \& \ \gamma) = PR(q \mid \gamma).$$

Since we are sure that the belief change is by probability kinematics on Γ, it is sufficient, to specify the possible final probabilities, that we specify just the final probabilities of members of the partition thus:

$$\wedge_i \, pr_f(\gamma_i) = \alpha_i,$$

since only one final probability can meet this specification and come from the initial probability by probability kinematics on Γ. So (S) becomes:

$$\text{(M}'\text{)} \quad PR(q \mid \wedge_i \, pr_f(\gamma_i) = \alpha_i \ \& \ \gamma) = PR(q \mid \gamma).$$

Now the foregoing is all done in terms of what your present probabilities are about the way that your final probabilities will be after the ineffable observational interaction, without speculation as to the nature of that interaction. Nevertheless, it implies that your probabilities are structured *as if* your experimental result consisted in learning $\wedge_i \, pr_f(\gamma_i) = \alpha_i$, and conditionalizing on the result. Then Good's theorem goes through just as in the last section, with these sentences in place of the e_ks.

Graves's treatment assumes for simplicity that the prior probabilities are concentrated on a finite number of possible combinations of posterior probabilities of members of the partition, but the analysis generalizes in a straightforward way to the continuous case. Consider belief change by probability kinematics on the partition $[e_1,e_2]$ where the prior for prf(e_1) is given by a continuous probability density. Here we need to strengthen (M) to:

$$\text{(M+)} \quad PR[e_1 \mid pr_f(e_1) \varepsilon I] \; \varepsilon I,$$

where I is any closed interval such that $PR[pr_f(e_1) \varepsilon I] > 0$. An immediate consequence of (M+) is the expectation principle:

$$\text{(E)} \quad PR\,(e_1) = E\,[pr_f\,(e_1)].$$

This together with the fact that the the expected utility of the Bayes act is a convex function of pr(e_1) leads immediately to the Good theorem. Let $\varnothing$ be the expected utility of the Bayes act. From the convexity of $\varnothing$ it follows on general mathematical grounds (Jensen's inequality: Royden 1968, 110) that:

$$E\,[\varnothing\,\{pr_f(e_1)\}] \geq \varnothing\,[E\,\{pr_f(e_1)\}],$$

so by (E):

$$E\,[\varnothing\,\{pr_f(e_1)] \geq \varnothing[\,PR(e_1)];$$

but by (M+) and the theorem on total probability, the prior expectation of the expected utility of the posterior Bayes act is equal to the prior expected utility of the decision to make an informed decision, q.e.d.

III. Ramsey's Anticipation

Good makes no great claims of originality. He cites treatments of the value of evidence by Lindley (1965) and Raiffa and Schlaifer (1961) as partial anticipations. To this list one must surely add Savage himself, who discusses the value of observation and indeed proves a form of the Good theorem in Chapter 7 (and appendix 2) of *The Foundations of Statistics* (1954). It would not be suprising if one could trace the basic idea back a little further. But it is worth reporting that Frank Ramsey had it back in the 1920s.

In two manuscript pages on "Weight or the Value of Knowledge," Ramsey (unpublished) sketches a version of the theorem. He proceeds in a rich setting where your act consists in choosing the values of a list, $x_1,x_2, \ldots$; of "control vari-

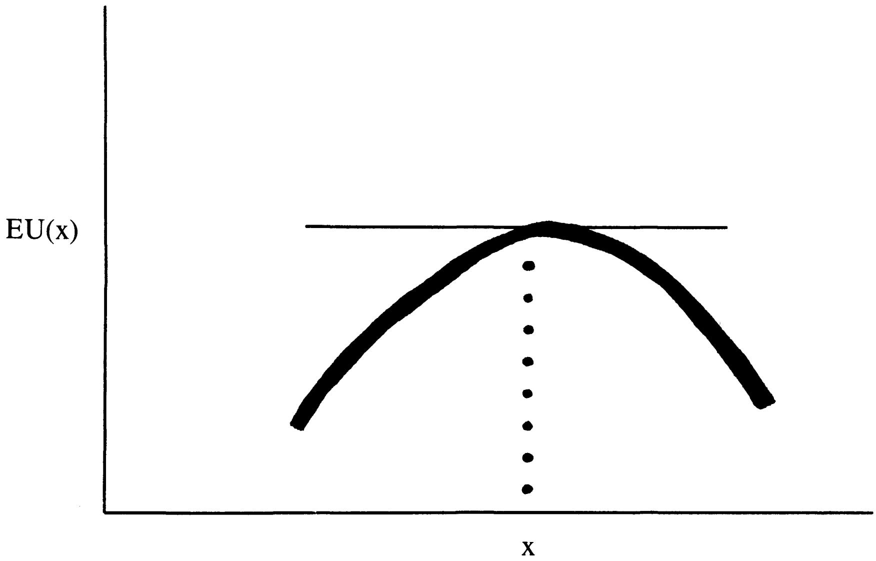

Figure 3. EU(x) Assuming a Maximum.

ables." He supposes that there is an unknown proposition, a, such that the expected utility of x_i considered as a function of x_i [for a fixed value of pr(a)] is continuous and twice differentiable and assumes its maximum at a nonextreme value of x_i; so that at the maximum $\delta EU(x_i \mid a)/\delta x_i = 0$ with the second derivative negative. This situation is illustrated in figure 3. It is also assumed that different acts are optimal for $pr(a)=1$ and $pr(a)=0$.

In this context Ramsey considers a function, $\varnothing(p)$, which he calls the "expectation of advantage in regard to a if I expect it with probability p." This is what we called in section 1 "the expected utility of the prior Bayes act." He argues that the second derivative of this function must be everywhere positive; i.e., that the function must be strictly convex. I reproduce as figure 4 Ramsey's own illustration of the situation.[1] It can be compared with the case of a finite number of acts shown in figure 2 of section 1.

It follows immediately that the expected value of an experiment whose possible results are a, not-a is positive. If Ramsey had simply noted this, it would have been enough for the theorem. Instead, however, he did something much more interesting. He considered the case in which the experimental result was not the truth value of a, but the truth value of another proposition, k, whose only effect on expected utility of the acts is in its alteration of the probability of the

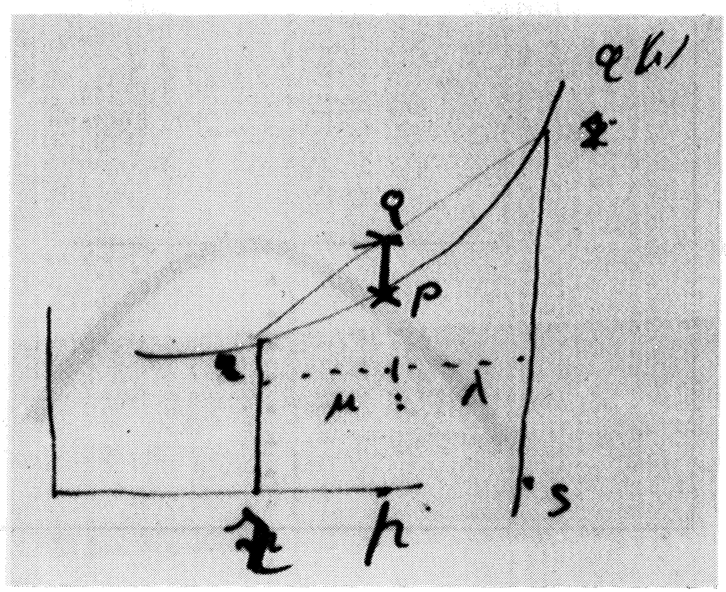

Figure 4. Ramsey's Diagram. (Courtesy of the University of Pittsburgh).

experiment. (In other words, probability of a conditional on the experimental result considered as a random variable is a sufficient statistic for the experiment.) In Ramsey's illustration, r is the probability of a conditional on k (together with background knowledge) and s is the probability of a conditional on the denial of k. As Ramsey notes, it is evident (from the strict convexity of ∅) that the gain in expected utility associated with undertaking the experiment (the length of the line segment PQ) must be positive unless k is irrelevant to the probability of a, and $p=r=s$. This can be seen as a special case of the problem discussed in the last section, where we have belief change by probability kinematics on a, not-a with only two possible results.

This is of special interest because we have other evidence that probability kinematical ideas were not foreign to Ramsey. There is a partial anticipation of probability kinematics in an 1851 paper by W.F. Donkin. Ramsey took a page of notes from this article. The manuscript is in the Pittsburgh archives. I quote Ramsey quoting Donkin:

> If there be any number of mutually exclusive hypotheses h_1, h_2, h_3 -; of which the probabilities relative to a particular state of information are p_1, p_2, p_3 -, and if new information is gained which changes the probabilities of some of them suppose of $h_{m=1}$ and all that follow, *without having otherwise any reference to the rest*, then the probabilities of these latter have the *same ratios* to

one another after the new information that they had *before*— . . . (Emphasis is Ramsey's.)

IV. Dynamic Probability and Other Forms of Generalized Learning

Suppose that you are in an even more amorphous learning situation than the kind that motivates Jeffrey's ideas. There is no nontrivial partition that you expect with probability one to be sufficient for your belief change. Perhaps you are in a novel situation where you expect the unexpected observational input. Perhaps there is to be no external observational input, and you are in the realm of what Good calls "dynamic probability." You are just going to think about some subject matter and the input, if you are lucky enough to have one, will be the "aha erlebniss." How unstructured can the setting for belief change be while the Good theorem is retained?

Reflection on section 2 suggests that the theorem does not really depend on any restrictions imposed by probability kinematics. After all, in the case of a discrete space where all the atoms have positive probability, *any* belief change is by probability kinematics on the partition of unit sets of the atoms. And the theorem did not depend on the belief change being by probability kinematics on any special partition. We do, however, need some form of principle M to assure you that you regard the upcoming belief-changing process as a learning experience.

Consider the case in which there are a finite number of relevant states of the world and a finite number of acts. Let the states of the world be $K_1, \ldots K_n$. The set of probability measures over these states can be represented as a k-1 dimensional simplex in k dimensional Euclidian space. The expected utility of the Bayes act is a function of these probabilities. Now suppose that the you are about to undergo what you expect will be a learning experience in that your prior probabilities satisfy an appropriate form of principle M:

$$(M++)\text{: } \inf \text{PROB } (q) \leq \text{PR } [q \mid pr_f \; \varepsilon \text{ PROB}] \leq \sup\text{PROB } (q),$$

for any q, where PROB is any rectangle (hyperinterval) of probability measures, such that PR[$pr_f \varepsilon$PROB] is positive. This is all we assume about the learning experience. M++ guarantees that:

$$(E)\text{: } \quad PR = E\,[pr_f].$$

$\varnothing$ is still convex. By the appropriate form of Jensen's inequality (Loeve 1963, 159):

$$E\,[\varnothing(pr_f)] \geq \varnothing\,[E(pr_f)],$$

and by (E):

$$E\,[\varnothing(pr_f) \geq \varnothing\,[PR].$$

The argument can be generalized under appropriate regularity conditions to cover infinite state spaces and infinite numbers of acts, but this involves some mathematical complications.[2]

Here is Good thinking reduced to its bare essentials. Nothing at all about the nature of the event that is to occasion your belief change has been specified, excepting your belief in the epistemological legitimacy of the impending belief change as embodied in (M). Bayesians have found it difficult to say anything informative about belief revision situations with so little structure, but in the presence of higher-order probabilities and condition (M) Good's theorem emerges from the mists unscathed.

V. Conflicting Expectations

I noted at the onset that Good's theorem is proved in the context of a Savage-type expected utility theory. The remark was not idle. In the variant form of expected utility theory best known to philosophers—that developed in Richard Jeffrey's *Logic of Decision* (1st ed.)—Good's theorem fails (Adams and Rosenkrantz 1980; Skyrms 1982). And for a number of forms of "causal decision theory" introduced by philosophers as alternatives to Jeffrey's theory, the question has not been adequately discussed (Skyrms 1982).

In Jeffrey's (1965) theory, there is no distinction between states, acts, and consequences. Probabilities and utilities are defined on a Boolean σ-algebra of propositions, and what we might intuitively take to be states, acts, and consequences are represented as propositions. The expected value of a proposition, A, is a conditional expected utility:

$$\text{(Jeffrey) } V(A) = \Sigma_i \Pr(A \mid B_i)\, V(A \,\&\, B_i),$$

for *any* partition $[B_i]$. The probability in question is subjective degree-of-belief. Note, in particular, that A may have positive probabilistic relevance to B because A is *evidence* for B, rather than a causal factor favorable to B. For this reason we call the theory, which holds that the choice-worthiness of an act is measured by this conditional epistemic expected value, "evidential decision theory." Although Jeffrey (1965) is, perhaps, the most thorough development of evidential decision theory, it is by no means the only endorsement of this type of theory. Evidential decision theory, in various forms, has a number of advocates.

Jeffrey's theory has as a stated aim the elimination of the causal concepts implicit in Ramsey's notion of a gamble and in Savage's distinction between acts, states, and consequences from the logic of decision. This gives rise to anomalous results in certain situations in which the act is symptomatic rather than causative.

Newcomb's problem (see Nosick [1969]) is such a case. You have just taken a psychological examination that predicts fairly well how people behave in the impending decision situation. There are two boxes, one transparent and one opaque. Your choice is either to take the opaque box and get whatever is under it, or to take both boxes (yes, both!) and get everything under the opaque box, together with \$1,000, which is visible under the transparent box. The experimenter has put \$1,000,000 under the opaque box if her test predicted that you will take only it; nothing otherwise. The money is either there already or not, and you are convinced that there will be no cheating on the part of the experimenter. Your subjective probabilities that the experimenter will make the right prediction given that you take one box and given that you take both are both greater than .6. Taking one box is evidence that the million is there, but can in no way be a cause of its being there. A straightforward application of Jeffrey's theory leads to the recommendation that you take only the opaque box.

In order to avoid such results, Robert Stalnaker (1972) suggested that the Jeffrey expectation be modified by replacing the conditional probability with the probability of a subjunctive conditional. This idea was systematically developed by Gibbard and Harper (1980), who define expected utility of an act, A, as:

$$\text{(SGH)}\ \Sigma_i \Pr [\text{If I were to do A, then } O_i]\, D(O_i),$$

where $[O_i]$ is a partition of ultimate consequences and D is desirability.

The Stalnaker-Gibbard-Harper theory is, like that of Savage, made for conditions of determinism, for it is assumed that the subjunctive conditionals involved themselves form a partition. David Lewis (1980) suggested an extension of Stalnaker's approach to indeterministic situations by replacing the probability of the conditional with the expectation of the value of chance in a conditional with chance consequent:

$$\text{(Lewis)}\ \Sigma_{ij}\, x_j \Pr [\text{If I were to do A, CHANCE } (O_i) = x_j]\, V(O_i).$$

I sought the same generality in an extension of the Savage approach to conditions of indeterminism. My proposal was to take expected utility as:

$$\text{(Skyrms)}\ U(A) = \Sigma_{ij} \Pr(K_i) \Pr(C_j \mid A \,\&\, K_i)\, U(C_j \,\&\, K_i \,\&\, A),$$

where $[K_i]$ is a partition of states interpreted as causal preconditions of the decision, $[C_j]$ is a partition of consequences, and A is the act. The foregoing three proposals go by the (possibly misleading) name "causal decision theory."

Now Good's theorem fails for evidential decision theory. We can use a variation on Newcomb's problem to make the point. Suppose the experimenter offers you a peek under the opaque box before you make your decision, but cautions you that the accuracy of prediction holds up for subjects who are offered this option. An evidential decision theorist will presumably refuse the offer, reasoning that whatever he sees, he will subsequently take both boxes (on evidential theoretic grounds), which will be bad news (i.e., evidence that there is nothing under the opaque box).

What is the status of Good's theorem in causal decision theory? The answer is simplest with respect to the Stalnaker-Gibbard-Harper theory, for that theory can be viewed as a reformulation of the theory of Savage. In Savage's theory, the act together with the state determine the consequence, so that Savage can represent acts as random variables on the space of states. We can use the Savage framework to give a semantics for the subjunctive conditionals that occur in the Gibbard-Harper account. The conditional: "If I were to do A, consequence C would follow," is true in a state just in case the act A maps that state onto consequence C. The set of states in which such a subjunctive conditional is true is the inverse image of the consequence C under the function A. Conversely, given the Gibbard-Harper account, one can reconstruct Savage-type states as appropriate bundles of subjunctive conditionals. These remarks fall somewhat short of showing that the theories are fully intertranslatable, but are sufficient for our purposes. The Stalnaker-Gibbard-Harper theory automatically inherits Good's theorem from its Savage counterpart.

What about the more general indeterministic forms of causal decision proposed by Lewis and myself? Here again we have one question rather than two, for a relation holds between our respective proposals analogous to that between Savage and Stalnaker-Gibbard-Harper. A state of the world, K, makes the subjunctive conditional "If I were to do A, CHANCE (C) = x," true just in case Pr $[C | A \& K] = x$.

It will suffice then, to discuss one of the proposals. I choose my own. The expected utility of A therein:

$$\Sigma_{ij} \Pr(K_i) \Pr(C_j | A \& K_i) U(C_j \& A \& K_i),$$

can be thought of as an expansion of Savage's:

$$\Sigma_i \Pr(K_i) U(A \& K_i),$$

wherein Savage's $U(A \& K_i)$ is analyzed as:

$$\Sigma_j \Pr(C_j | A \& K_i) U(C_j \& A \& K_i).$$

Accordingly, Good's theorem extends to this case provided that this quantity is independent of each experimental result. A sufficient condition for this is that we have both the following:

(I) $U(C_j \& A \& K_i \& e_k) = U(C_j \& A \& K_i)$ for all i,j,k.
(II) $\Pr(C_j | A \& K_i \& e_k) = \Pr(C_j | A \& K_i)$ for all i,j,k.

(I) is one precise way of saying in this context that information is really free. (II) may be thought of as saying that all the experimental results do is give you information about the state of the world. They do not affect your belief about the conditional chance of a consequence on an act obtaining in a given state of the world. Given the intended interpretation where $[K_i]$ is a sufficient partition for conditional chance of consequence on act, and thus state together with act determine chance of consequence, condition (II) will be fulfilled if the theory has been properly applied. Good's theorem holds for the most general form of causal decision theory.

VI. Deliberational Equilibrium and Rational Decision

The dominant rationality concept in Bayesian decision theory is maximization of expected utility. The dominant rationality concept in the theory of games is equilibrium. The equilibrium concept has been absent from most discussions of the philosophical foundations of decision theory because those discussions have neglected aspects of the *process of deliberation*.

In the simplest cases, the dynamics of deliberation may be trivial. One calculates expected utilities and maximizes. But there are more complex cases in which the very process of deliberation can generate new information relevant to the evaluation of the expected utilities at issue. In such cases, an act may have highest expected utility in terms of the probabilities available at the onset of deliberation, but look worse than other alternatives if we feed in the information that it is the act about to be chosen.

Such cases have moved toward center stage in the philosophical discussions of the respective merits of evidential and causal decision theory. To deal with them, apologists for both theories have suggested that some sort of equilibrium condition be added to the principle of maximum expected utility in the account of individual rational decision.

Richard Jeffrey and Ellery Eells explore the possibility that the class of equilibrium decisions under evidential decision theory may not include the anomalous cases that embarrass the theory. The idea is that the extra information generated by deliberation leading to a decision would screen off the spurious correlation between act and state of the world. And William Harper, responding to problems raised by Reed Richter, suggests that causal decision theory needs to be supplemented by an equilibrium requirement. Harper, following Jeffrey, calls this a requirement of "ratifiability."

The foregoing discussion of the value of knowledge in causal and evidential decision theory suggests the following two propositions:

(I) The addition of an equilibrium requirement is *inconsistent* with the expected utility principle in evidential decision theory.

(II) The equilibrium condition, in so far as it is legitimate, *follows* from the expected utility principle in causal decision theory.

I will discuss them in order.

If a decision maker finds himself about to make an equilibrium decision, there is no problem of consistency between the equilibrium condition and the expected utility hypothesis. But suppose that the act he has calculated to have maximum expected utility will not retain maximum expected utility if the information that he is about do it is used to update his probabilities. The "ratifiability" defense of evidential decision claims that the evidential decision theorist who has just selected "take one box" as the Bayes act is in this position; that were he to update and recalculate at the moment of truth he would regret not taking two boxes. Let us suppose for the sake of argument that a case can be made out for the dubious claim that the additional information will so render act and state independent. The question still remains as to whether the decision maker should update and recalculate. This is a question of the value of knowledge, and we have seen that even in the ideal case in which the evidential decision maker is offered a free knowledge of the state of nature the theory may lead him to reject that offer. That is, the evidential decision theorist may associate negative expected utility with a policy of moving to an equilibrium decision.

In causal decision theory, on the other hand, the general validity of Good's theorem suggests that in the absence of substansial processing costs, one should assimilate whatever information is generated by the process of deliberation. Deliberation then becomes a dynamic process of moving towards the apparent optimal act under conditions of informational feedback. In such a decision process, a decision maker cannot choose a nonequilibrium act. As he gets close to choosing the act, informational feedback will make it appear less than maximally attractive, and will point his deliberation in a different direction. In causal decision theory, a policy of selecting an equilibrium decision is almost a consequence of the expected utility principle.

Why almost? Well, in the first place there are real-world costs of computation that might be significant. But even if we idealize away these costs, there is a kind of fallacy of composition involved in assuming that one can justify a strategy of informational feedback by induction on the application of the Good theorem. Consider an artificial paradoxical situation in which there is always free information available, and the decision as to which act to perform could be postponed for any finite time. Under a strategy of always assimilating free information, the decision maker would never act! When the strategy as a whole is evaluated against "Choose the prior Bayes act," we have seen that the conditions of the theorem are violated. There is no choice of a posterior Bayes act subsequent to the deliberational process. Such a deliberational strategy is not Good thinking.

A correct Bayesian treatment of deliberation must evaluate deliberational strategies as wholes. The question of optimum deliberational strategies for problems of the kind we have been discussing is a large and complex question, which I will not address in any detail here. A few qualitative points can, however, be made on the basis of the foregoing discussion.

Suppose that you are confronted with two opaque boxes, A and B, and are to choose one and receive its contents. A mean demon, who you believe is a good predictor of your behavior conditional on either choice, has put money under the box she predicted you won't choose. Let's say \$1,000 if it is A and & 1,100 if B, to make it interesting. Suppose you initially incline towards B, but when about to take B, you recalculate and find that A is more appealing; but when about to take B, you find A more appealing; etc. You find your deliberations oscillating, with convergence—if any—towards indecision rather than decision. You are suffering from Richter's complaint (Richter 1984). Richter's theory of its etiology is that there is something wrong with causal decision theory. Sharvy (1983) and Harper (1986) reply that you should consider mixed decisions, which do indeed include an equilibrium decision. Richter turns the screw. Randomization costs something. In fact, let's say that the reason that this is a mean demon is that she will boil you in oil if she catches you using a nondegenerate mixed strategy.

In this case I would say that the cause of Richter's complaint is not a defect in causal decision theory, but rather a misapplication of Good's theorem. The initial choice is not between "Decide now between A and B," and "Deliberate, generate free information, and then face essentially the same choice situation." Rather, if we take the story at face value, it is between the former and "Deliberate, choose A, B, or a mixed strategy; or fail to make any choice." If this is the only other sort of deliberational strategy open to you, and you assign any positive probability to deliberation's getting you into the kind of trouble described, causal decision theory can recommend "Decide now between A and B," in accord with Richter's intuition.

The sort of pathology to which Richter calls our attention is not simply the lack of an equilibrium decision, as is shown in the following example (Skyrms 1984, 83):

> You are to choose one of three shells [A_1,A_2,A_3], and will receive what is under it. No mixed acts are allowed. (If you attempt to randomize, even mentally, the attempt will be detected and you will be shot.) A very good predictor has predicted your choice. If he predicted A_1, he put 10 cents under shell one and nothing under the others. If he predicted A_2, he put \$10 under shell two and \$100 under shell three. If he predicted A_3, he put \$20 under shell three and \$200 under shell two.

Suppose you start deliberation with a very small probability of choosing A_1 and equal probabilities of choosing A_2 and A_3. If you deliberate continuously, moving towards maximum expected utility with informational feedback, you will suffer from Richter's complaint, hanging up between A_2 and A_3. Supposing that if you can't come to a decision, you get nothing, the expected utility of choosing the prior Bayes act is higher than that of the deliberational strategy that cannot lead to a decision.

The example was designed to show something else as well. Suppose that you also have another deliberational strategy under consideration: "Choose the equilibrium decision with highest prior expected utility" (Harper 1984). There is a unique equilibrium decision here, but the conditions for the Good theorem fail as before.

What can we say about the conditions under which the Good theorem *will* apply to a deliberational strategy? In the first place, the process as a whole must be cost free. If it is possible that the process does not lead to any decision, then there is an associated cost equal to the difference in payoff of the prior Bayes act and of no decision. For the process as a whole to be cost free, the expected costs of no decision must be zero. (This may happen because a deliberational strategy has prior probability zero of leading to no decision, even though it may possibly lead to no decision.) Otherwise the costs of deliberation by that strategy must be balanced against its benefits. (One way of guaranteeing a decision is to adopt a strategy with a time limit, such that if no pure act has been selected by the time limit, the mixed act corresponding to the decision maker's state of indecision about pure acts will be selected as a default.)

In the second place, the deliberational strategy as a whole must satisfy condition (M). Thus, for any fixed act, its initial expected utility must be equal to the initial expectation of its final expected utility. This can be thought of as a condition that the deliberational strategy be unbiased. One cannot, for instance, get away with a deliberational strategy designed to generate good news. And for deliberation to be nontrivial, one must be uncertain about where deliberation will lead. Since initial probability is the initial expectation of final probability, we have:

$$\text{If } Pr_i[pr_f(p)=\alpha] = 1, \text{ then } Pr_i(p) = \alpha.$$

If you know where you're going, you're already there. Thus to the extent that deliberation generates information by computation, the results of computation must be initially uncertain.

Condition (M) can also give us some guidance regarding the internal structure of a deliberational strategy. Consider the deliberational strategy that computes expected utility; assigns probability one to the act with maximum expected utility (provided there is a unique one); revises probabilities of states of nature accordingly; and recomputes, etc. In the mean-demon case, this strategy oscillates between assigning probability one to box A and probability one to box B. From the point of view of automatic control, some damping in the system would be desirable. Condition (M) provides a less *ad hoc* justification. Applying it now to stages of deliberation we have:

$$\text{If } Pr_i[p] = 1 \text{ then } Pr_i[pr_f(p) < 1] = 0$$

or, by contraposition, *if you think that you may change your mind, you're not certain*. A strategy that expects informational feedback that may with some positive probability alter the Bayes act, and that proceeds stagewise in accordance with (M), will not leap immediately to the assignment of probability one to the prior

Bayes act, but will rather move some distance in the direction of attractive options. Very simple models of this kind show suprisingly nice properties of convergence to equilibrium decisions. (E.g., the Nash dynamics discussed in Skyrms [1986] "Deliberational Equilibria.")

Questions of deliberational dynamics become more interesting in game-theoretic contexts in which Bayesian players who know each other's initial starting points deliberate, while attempting to second-guess each other's deliberations. In such contexts, informational feedback may be more interesting than in examples like the mean demon, and individual deliberational equilibriums may be related to game theoretic equilibriums. Rational deliberation may lead Bayesian players from an initial position to a unique game-theoretic equilibrium solution in non-zero sum games with multiple equilibriums. In this way, deliberational dynamics is related to the "tracing procedure" of Harsanyi and Selten (Harsanyi 1975). Detailed discussion of these questions is beyond the scope of this paper, but perhaps enough has been said to show that considerations of the value of information must play a foundational role in the theory of deliberation.

VII. Condition M

The analysis of generalized Good thinking shows that condition M is of central importance. Fulfillment of condition M shows that you regard any impending belief change as a generalized learning experience. This applies even when the experience cannot be neatly summarized as being given a proposition in your probability space. Consider the example of Ulysses and the sirens. Prior to the encounter with the sirens, Ulysses has initial probabilities, Pr, about what his probabilities will be after hearing the siren song. The possible sensory siren inputs cannot be easily summarized as propositions in his probability space. Ulysses believes that the sirens have the power and predilection to cloud men's minds so that they cease to believe that the rocks are dangerous (R). Condition M fails:

$$\Pr(R \mid pr_f(R) < .1) > .1,$$

and Ulysses believes that this sort of input can get him in deep trouble. Prudently, he makes arrangements (i) to prevent the input, and (ii) to prevent himself from acting effectively if (i) fails.

In the setting of the original Good theorem, condition M holds automatically. Since the learning experience consists of conditionalizing on the experimental result, e_k, final probability after the learning experience as a random variable in the initial probability space is identical to probability conditional on the partition of possible experimental results:

$$pr_f(\bullet) = \Pr[\bullet \| \{e_k\}],$$

which is a point function that takes as its value the ordinary conditional probability Pr[• | e] for the member of the partition, e, that includes the point. But condition M holds by definition for probability conditional on a partition. Conversely, in the "dynamic probability" setting, with a finite number of possible pr_fs, we expand the initial small probability space by taking its product with space generated by taking the possible pr_fs as atoms. In the larger space, we have the partition of possible final probabilities (each of which, we are assuming, has positive prior probability). Condition M is necessary and sufficient for the final probability as a random variable to be equal to probability conditional on *this* partition. M is the principle that relates final probability to conditional probability.

M is a principle of dynamic *coherence*. An agent's degrees of belief are statically coherent at a time in the weak sense if it is impossible to make a "Dutch book" against him in a finite number of bets that he regards at that time as fair or favorable, such that he suffers a net loss in every possible outcome. An agent's degrees of belief are *statically coherent in the strong sense* if it is impossible to make a Dutch book against him using bets all of which he considers favorable (or, equivalently, assuming payoffs are real valued, using a countable number of bets that he considers fair or favorable). Static coherence in the strong sense requires the agent's degrees of belief to be a countably additive probability measure (Adams 1962), and we make this strong assumption in this paper. An agent is *dynamically coherent* (in strong and weak senses) if a Dutch book cannot be made against him by betting at different times, the bettor at the time of betting knowing no more than the agent does. (Goldstein 1983; Skyrms forthcoming a,b; Van Fraassen 1984).

An agent has probabilities now about her probabilities tomorrow. Consider the simple case in which there are only a finite number of probabilities that she may have tomorrow, each of which today has finite positive probability. We assume that tomorrow the agent will *know her own mind*: $Pr_f\ (pr_f = Pr_f) = 1$, so controversial features of higher order probability do not come into play. If we give the epistemic situation only this much structure, what can we say about dynamic coherence?

A necessary and sufficient condition for dynamic coherence is that the agent's initial probability, Pr_i, satisfy:

$$\text{(M)}\ PR(q \mid pr_f = pr^* = pr^*(q).$$

Necessity is established directly by describing a betting strategy that, by a finite number of bets today and tomorrow, can make a Dutch book against any agent who violates M (as in Goldstein 1983; or Van Fraassen 1984). Sufficiency is essential because the random variable pr_f can be construed as probability conditional on a partition. Any payoff function that the bettor can attain by betting at t_1 today and then tomorrow, if pr_f in some set S, can be attained by betting only at t_1 utilizing bets conditional on S (where S is a disjunction of a finite number of pr_2s). (Something similar can again be said with regard to more general versions of M. See Goldstein [1983]; Gaifman [1986]; Skyrms [forthcoming a]).

M is a condition for convergence.[3] Consider an infinite sequence of learning situations at time $t_1, t_2, \ldots$ and a corresponding sequence of probabilities $pr_1, pr_2, \ldots$ that indicate respectively the upshot of these learning situations. The situations might, for instance, be observations of outcomes of flips of a coin with unknown bias with updating by conditionalization. Prior to the sequence, we suppose that you have a big probability space on which $pr_1, pr_2, \ldots$ are random variables. Suppose also that you have the appropriate version of M to assure the expectation principle:

$$E\,[pr_{n+1} \,\|\, pr_1, \ldots pr_n] = pr_n.$$

(In our example, the random variable consisting of the expected value of probability of heads after observing n+1 tosses, conditional on the probabalities after observing the first n tosses, is identical to the probability after observing n tosses considered as a random variable.) Then the sequence of upcoming revised probabilities as random variables, $pr_1, pr_2, \ldots$ forms a nonnegative *martingale*. Then with probability one, the sequence $pr_1, pr_2, \ldots$ converges to a random variable pr_{ff} with:

$$E\,[pr_{ff}] = E\,[pr_1] = PR$$

The random variable, pr_{ff}, is final probability in the light of the whole learning sequence. In our example, pr_{ff} is a reasonable facsimile of "the true chance that the coin comes up heads." Your initial probability of heads, PR(H), is equal to your initial expectation of the true chance of heads, E $[pr_{ff}(H)]$. The example generalizes (see Dynkin 1978; Diaconis and Freedman 1981). Almost everything we know about convergence to a limiting relative frequency—the strong law of large numbers, de Finetti's theorem, and generalizations of de Finetti's theorem for various versions of partial exchangeability—are special cases of this martingale convergence argument.

Returning to the theme of this paper, what can we say about the value of knowledge in the long run? Suppose that learning experiences were really free in terms of time and opportunity costs, so that one could undertake an infinite number of learning experiences in a finite time, and then make a decision. By the martingale convergence theorem, the same argument used in section 4 for pr_f, goes through here for pr_{ff}, establishing in a general way the value of making a most informed decision.[4]

Notes

1. I would like to thank the Ramsey family and the Archives for Logical Positivism at the University of Pittsburgh for permission to reproduce this figure.

2. The general case is treated in my "On the Principle of Total Evidence with and without Observa-

tion and Sentences," in *Logic, Philosophy of Science and Epistemology*, Proceedings of the 11th International Wittgenstein Symposium (Holder-Pichler-Tempsky: Vienna, 1987).

3. The following argument assumes countable additivity.

4. Research partially supported by NSF grant SES-8605122.

References

Adams, E., and Rosenkrantz, R. 1980. Applying the Jeffrey Decision Model to Rational Betting and Information Acquisition. *Theory and Decision* 12: 1–20.

Adams, E. 1962. On Rational Betting Systems. *Archive fur Mathematische Logik und Grundlagenforschung* 6: 7–18; 112–28.

Armendt, B. 1980. Is there a Dutch Book theorem for Probability Kinematics? *Philosophy of Science* 47: 583–88.

——. 1986. Foundations of Causal Decision Theory. *Topoi* 1.

de Finetti, B. 1977. "Probabilities of Probabilities: A Real Problem or a Misunderstanding?" In *New Developments in the Applications of Bayesian Methods*, ed. Aykac and Brumat 1–10. Amsterdam: North-Holland.

Diaconis, P., and Freedman, D. 1981. "Partial Exchangeability and Sufficiency." In *Statistics: Applications and New Directions*, 205–36. Calcutta: Indian Statistical Institute.

Diaconis, P., and Zabell, S. 1982. Updating Subjective Probability. *Journal of the Americal Statistical Association* 77: 822–30.

Donkin, W. F. 1851. On Certain Questions Relating to the Theory of Probabilities. *Philosophical Magazine*.

Dynkin, E. B. 1978. Sufficient Statistics and Extreme Points. *Annals of Probability* 6: 705–30.

Eells, E., and Sober, E. 1986. Common Causes and Decision Theory. *Philosophy of Science* 53: 223–45.

Eells, E., 1982. *Rational Decision and Causality*. Cambridge: Cambridge University Press.

——. 1984. Metatickles and the Dynamics of Deliberation. *Theory and Decision* 17: 71–95.

——. 1985. Weirich on Decision Instability. *Australasian Journal of Philosophy* 63: 473–78.

——. 1984. Causal Decision Theory. In *PSA 1984* East Lansing: Philosophy of Science Assn.

Gaifman, H. 1988. "A Theory of Higher Order Probabilities." In *Causation, Chance and Credence*, ed. B. Skyrms and W. Harper. Dordrecht: Reidel.

Gibbard, A., and Harper, W. 1980. "Counterfactuals and Two Kinds of Expected Utility." In *IFS*, ed. Harper et al., 153–90. Dordrecht: Reidel.

Goldstein, M. 1983. The Prevision of a Prevision. *Journal of the American Statistical Association* 78: 817–19.

Good, I. J. 1950. *Probability and the Weighing of Evidence*. New York: Hafner.

——. 1967. On the Principle of Total Evidence. *British Journal for the Philosophy of Science* 17: 319–21.

——. 1974. A Little Learning Can Be Dangerous. *British Journal for the Philosophy of Science* 25: 340–42.

——. 1983. *Good Thinking: The Foundations of Probability and Its Applications*. Minneapolis: University Of Minnesota Press.

——. 1981. The Weight of Evidence Provided by Uncertain Testimony or from an Uncertain Event. *Journal of Statistical Computation and Simulation* 13: 56–60.

Graves, P. Forthcoming. "The Total Evidence Principle for Probability Kinematics." In *Philosophy of Science*.

Harper, W. 1986. Mixed Strategies and Ratifiability in Causal Decision Theory. *Erkentniss* 24: 25–36.

——. 1984. Ratifiability and Causal Decision Theory: Comments on Eells and Seidenfeld. In *PSA 1984*. East Lansing: Philosophy of Science Association.

——. Forthcoming. "Causal Decision Theory and Game Theory." In *Causation in Decision, Belief Change and Statistics*, ed. Harper and Skyrms. Dordrecht: Reidel.

Harsanyi, J. C. 1967. Games with Incomplete Information Played by Bayesian Players, parts I,II,III. *Management Science* 14: 159–83, 320–34, 486–502.

——. 1973. Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points. *International Journal of Game Theory* 2: 1–23.

——. 1975. The Tracing Proceedure: A Bayesian Approach to Defining a Solution Concept for N-Person Non-cooperative Games. *International Journal of Game Theory* 4: 61–94.

Horwich, P. 1985. Decision Theory in the Light of Newcomb's Problem. *Philosophy of Science* 52: 431–50.

Jeffrey, R. 1965. *The Logic of Decision*. N.Y.: McGraw Hill; 2d rev. ed., Chicago: University of Chicago Press 1983).

——. 1968. "Probable Knowledge." In *The Problem of Inductive Logic*, ed. Lakatos. Amsterdam: North-Holland.

——. 1981. The Logic of Decision Defended. *Synthese* 48: 473–92.

Kyburg, H. 1980. Acts and Conditional Probabilities. *Theory and Decision* 12: 149–71.

——. Forthcoming. "Powers." In *Causation in Decision, Belief Change and Statistics*, ed. W. Harper and B. Skyrms. Dordrecht: Reidel.

Levi, I. 1975. Newcomb's Many Problems. *Theory and Decision* 6: 161–75.

——. 1982. A Note on Newcombmania. *The Journal of Philosophy* 79: 337–42.

——. 1983. The Wrong Box. *The Journal of Philosophy* 80: 534–42.

Lewis, D. 1979. Prisoner's Dilemma Is a Newcomb Problem. *Philosophy and Public Affairs* 8: 235–40.

——. 1980. Causal Decision Theory. *The Australasian Journal of Philosophy* 59: 5–30.

Lindley, D. V. 1965. *Introduction to Probability and Statistics*, pt. 2. Cambridge: Cambridge University Press.

Loeve, M. 1963. *Probability Theory*, 3d ed. New York: van Nostrand.

Marshak, J. 1975. Do Personal Probabilities of Probabilities have an Operational Meaning? *Theory and Decision* 6: 127–32.

Nash, J. 1951. Non-Cooperative Games. *Annals of Mathematics* 54: 286–95.

Nozick, R. 1969. "Newcomb's Problem and Two Principles of Choice." In *Essays in Honor of Carl G. Hempel*, ed. N. Rescher. Dordrecht: Reidel.

Raiffa, H., and Schlaiffer, R. 1961. *Applied Statistical Decision Theory* Boston: Harvard Graduate School of Business Administration.

Ramsey, F. P. 1931. *The Foundations of Mathematics and Other Essays*, ed. R. B. Braithwaite. New York: Harcourt Brace.

——. Unpublished. Manuscripts 005–20–01, 005–20–03, 003- 13–01 in the Archives for Scientific Philosophy in the Twentieth Century at the Hillman Library of the University of Pittsburgh.

Richter, R. 1984. Rationality Revisited. *The Australasian Journal of Philosophy* 62: 392–403.

——. 1986. Further Comments on Decision Instability. *The Australasian Journal of Philosophy* 64: 345–49.

Royden, H. I. 1968. *Real Analysis*, 2d ed. London: Macmillan.

Savage, L. J. 1954. *The Foundations of Statistics*. N.Y.: Wiley.

——. 1967. Difficulties in the Theory of Personal Probability. *Philosophy of Science* 34: 305–10.

Seidenfeld, T. 1984. Comments on Eells. In *PSA 1984*. East Lansing: Philosophy of Science Association.

Sharvy, R. 1983. Richter Destroyed. Circulated typescript.

Skyrms, B. 1980a. *Causal Necessity*. New Haven: Yale University Press.
——. 1980b. "Higher Order Degrees of Belief." In *Prospects for Pragmatism*, ed. D. H. Mellor. Cambridge: Cambridge University Press.
——. 1982. Causal Decision Theory. *The Journal of Philosophy* 79: 695–711.
——. 1986. Deliberational Equilibria. *Topoi 1*.
——. 1985. Ultimate and Proximate Consequences in Causal Decision Theory. *Philosophy of Science* 52: 608–11.
——. 1984. *Pragmatics and Empiricism*. New Haven: Yale University Press.
——. 1987a. "Coherence." In *Scientific Inquiry in Philosophical Perspective*, ed. N. Rescher, 225–42. Pittsburgh: University of Pittsburgh Press.
——. 1987b. "Dynamic Coherence." In *Foundations of Statistical Inference*, ed. I. B. MacNeill and G. J. Umphrey. Dordrecht: Reidel.
——. 1987c. Dynamic Coherence and Probability Kinematics. *Philosophy of Science* 54: 1–20.
——. 1987d. On the Principle of Total Evidence with and without Observation Sentences. In *Proceedings of the 11th International Wittgenstein Symposium*, 187–95. Vienna: Holder-Pichler-Tempsky.
——. 1987e. Updating, Supposing and MAXENT. *Theory and Decision* 22: 225–46.
——. 1988. "Conditional Chance." In *Probability and Causality: Essays in Honor of Wesley Salmon*, ed. J. Fetzer, 161–78. Dordrecht: Reidel.
——. Forthcoming a. "Deliberational Dynamics and the Foundations of Bayesian Game Theory." In *Epistemology* [Philosophical Perspectives vol. 2], ed. J. E. Tomberlin. Northridge, Calif.: Ridgeview.
——. Forthcoming b. Probability and Causation. *Journal of Econometrics*.
Sobel, J. H. 1983. Expected Utilities and Rational Actions and Choices. *Theoria* 49.
——. 1985. Circumstances and Dominance Arguments in a Causal Decision Theory. *Synthese* 52.
——. 1986. Notes on Decision Theory: Old Wine in New Bottles. *Australasian Journal of Philosophy* 64: 407–37.
Stalnaker, R. 1972. "Letter to David Lewis," In *Ifs* (1980), ed. Harper et al. Dordrecht: Reidel.
Teller, P. 1973. Conditionalization and Observation. *Synthese* 26: 218–58.
——. 1976. "Conditionalization,Observation, and Change of Preference." In *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, ed. W. Harper and C. Hooker. Dordrecht: Reidel.
Uchii, S. 1973. Higher Order Probabilities and Coherence. *Philosophy of Science* 40: 373–81.
Van Fraassen, B. 1980. Rational Belief and Probability Kinematics. *Philosophy of Science* 47: 165–87.
——. 1984. Belief and the Will. *Journal of Philosophy* 81: 235–56.
Vickers, J. 1965. Some Remarks on Coherence and Subjective Probability. *Philosophy of Science* 32: 32–38.
Weirich, P. 1985. Decision Instability. *The Australasian Journal of Philosophy* 63: 465–72.
——. 1986. Decisions in Dynamic Settings. In *PSA 1986*. East Lansing: Philosophy of Science Assn.