



Dynamic Coherence and Probability Kinematics

Brian Skyrms

Philosophy of Science, Vol. 54, No. 1. (Mar., 1987), pp. 1-20.

Stable URL:

<http://links.jstor.org/sici?sici=0031-8248%28198703%2954%3A1%3C1%3ADCAPK%3E2.0.CO%3B2-Y>

Philosophy of Science is currently published by The University of Chicago Press.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/ucpress.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Philosophy of Science

March, 1987

DYNAMIC COHERENCE AND PROBABILITY KINEMATICS*

BRIAN SKYRMS†

*Department of Philosophy
University of California at Irvine*

The question of coherence of rules for changing degrees of belief in the light of new evidence is studied, with special attention being given to cases in which evidence is uncertain. Belief change by the rule of conditionalization on an appropriate proposition and belief change by “probability kinematics” on an appropriate partition are shown to have like status.

1. Dynamic Coherence. Frank Ramsey introduces the notion of static coherence of degrees of belief—coherence at a given time—as a kind of pragmatic consistency property:

These are the laws of probability, which we have proved to be necessarily true of any consistent set of degrees of belief. . . . If anyone’s mental condition violated these laws, his choice would depend on the precise form in which the options were offered him, which would be absurd. He could have a book made against him by a cunning bettor and would then stand to lose in any event. (Ramsey 1931, p. 182)

Ramsey rightly treats the possibility of a Dutch book as a symptom of deeper pathology. The bettor who violates the laws of the probability

*Received May 1985; revised November 1985.

†Research partially supported by National Science Foundation Grant SES-84-09898. Some of these results were presented at the Stanford University Statistics Colloquium in January, 1985, and to the Ohio State University Conference on Decision Theory in March, 1985. I would like to thank Ernest Adams, Brad Armendt, Nancy Cartwright, Persi Diaconis, Richard Jeffrey, Paul Teller, and Bas van Fraassen for comments on an earlier version of this paper.

Philosophy of Science, 54 (1987) pp. 1–20.
Copyright © 1987 by the Philosophy of Science Association.

calculus leaves himself open to having book made against him because he will consider two different sets of odds as fair for an option depending on how that option is described; the equivalence of the descriptions following from the underlying Boolean logic.

Ramsey (and following him, de Finetti and Savage) argued for the representation of coherent degrees of belief at a time as a finitely additive probability measure on a Boolean algebra of propositions. However, the argument can be extended in a natural way to justify countable additivity of the probability measure on a Boolean sigma-algebra of propositions (Adams 1962). Protection against a Dutch book resulting from a finite number of bets perceived as fair requires finite additivity; protection against one resulting from a countable number of bets perceived as fair requires countable additivity. We will assume countable additivity.

Later in the same essay, Ramsey discusses *changes* in degrees of belief:

Since an observation changes (in degree at least) my opinion about the fact observed, some of my degrees of belief after the observation are necessarily inconsistent with those I had before. We have, therefore, to explain exactly how the observation should modify my degrees of belief; obviously if p is the fact observed, my degree of belief in q after the observation should be equal to my degree of belief in p before, or by the multiplication law to the quotient of my degree of belief in pq by my degree of belief in p . When my degrees of belief change in this way we can say that they have been changed consistently by my observation. (Ramsey 1931, p. 192)

Ramsey is describing the process of belief change by *conditionalization* on p , “the fact observed”:

$$Pr_{\text{new}}(q) = Pr_{\text{old}}(q/p) = Pr_{\text{old}}(p \& q) / Pr_{\text{old}}(p)$$

The relevant conditional probabilities are often calculated by means of Bayes’ theorem, in which case the rule of updating by conditionalization is known as *Bayes’ rule*. Ramsey’s use of the word “consistently” in the foregoing passage might suggest that he has in mind a coherence argument to justify this rule parallel to the argument he used to justify the laws of the probability calculus as rules for consistent degrees of belief at a fixed time. But no such argument is given, either in this essay or elsewhere in his published writings. We are left to speculate as to whether he considered the argument too obvious to put down or whether he had no such argument.

There is a coherence argument, given explicitly by de Finetti, for the ratio definition of conditional probability:

$$Pr_{\text{old}}(q/p) = Pr_{\text{old}}(p \& q) / Pr_{\text{old}}(p).$$

This involves conditional bets. A bet on q conditional on p is won if p and q are both true, lost if p is true and q false, and called off if p is false. Conditional bets can be compounded from unconditional bets. A bet on $p \& q$ that wins c if $p \& q$ and loses d otherwise and a bet on not- p that wins d if not- p and loses f otherwise taken together give net payoff of a bet on q conditional on p that has zero payoff if the condition p is not realized, wins $c - f$ if $p \& q$, and loses $d + f$ if $p \& \text{not-}q$. Coherence requires that conditional probabilities as fair betting ratios for conditional bets mesh with fair betting ratios for unconditional bets, as in the standard definition of conditional probability. None of this in itself, however, gives us a coherence argument for the rule of conditionalization for *changing* degrees of belief.

In an essay examining Savage's system, Ian Hacking points to the rule of conditionalization as an implicit assumption. He calls it the *dynamic assumption* of personalism:

The idea of the model of learning is that $Prob(h/e)$ represents one's personal probability after one learns e . But formally the conditional probability represents no such thing. If, as in all of Savage's work, conditional probability is a defined notion, then $Prob(h/e)$ stands merely for the quotient of two probabilities. It in no way represents what I have learned after I have taken e as a new datum point. It is only when we make the dynamic assumption that we can conclude anything about learning from experience. (Hacking 1967, p. 315)

Hacking argues that no dynamic Dutch book argument is possible:

. . . a man knowing e would be incoherent if the rates offered on h unconditionally differed from his rates on h conditional on e . But no incoherence obtains when we shift from the point before e is known to the point after it is known. . . . since the man announces his post- e rates only after e is discovered, and simultaneously cancels his pre- e rates, there is no system for betting with him which is guaranteed success in the sense of a Dutch book. (1967, p. 315)

Given the way Hacking structures the problem, with the bettor simply betting at two different times with someone whose degrees of belief at each of those times are statically coherent, it is clear why he thinks that no Dutch book is possible. But the problem can be given more structure.

Indeed, any plausible treatment of the case must give it more structure, for there must be some way of indicating in the statement of the problem that the man in question learns only e , rather than e together with some extra information. Furthermore, it is of some importance that we are discussing coherence not of separate degree-of-belief states but of a *rule* or *strategy* for changing such states upon receiving a proposition as input.

Let us consider a perfect model situation for the application of such rules. Suppose that our bookie has at time 1 a prior probability assignment, Pr_1 , over a probability space, and there is a countable partition of that space, $\{e_i\}$, such that each member of that partition has positive prior probability. At time 2, the true member of the partition is announced and the bookie moves to a posterior probability, pr_2 , according to a strategy that treats the announced member of the partition as total input. Such a strategy is a *function*, STRAT, which maps members of $\{e_i\}$ onto posterior probability distributions. The strategy of conditionalization maps e_i onto the posterior probability, Pr_2 , such that $Pr_2(q) = Pr_1(q/e_i)$ for all q . There are many other possible strategies. We will allow a cunning bettor to bet with our bookie at each time. He must make bets that the bookie considers fair or favorable (non-negative expected utility) at the time. He is allowed to know the bookie's probabilities at the times of the bets and he is allowed to know the bookie's *strategy*. Formally, the bettor's strategy consists of a pair of functions. The first maps the pair consisting on the bookie's (Pr_1, STRAT) onto bets to be made at t_1 that for the bookie have non-negative expected utility at t_1 according to his pr_1 ; the second maps $(Pr_1, \text{STRAT}, e_i)$ onto bets to be made at t_2 that for the bookie have non-negative expected utility at t_2 according to his $pr_2 = \text{STRAT}(e_i)$. Let us say that the bettor makes a *dynamic Dutch book* against the bookie, if no matter what the true member, e_i , of the partition, the bettor's strategy leaves him at t_2 with a finite number of bets whose net payoff is positive for every point in e_i . And we will say that the bookie's strategy is *dynamically coherent* if no bettor's strategy makes a dynamic Dutch book against it.

Now there is an argument, due to David Lewis,¹ which shows that in the type of situation under consideration a dynamically coherent strategy must proceed by conditionalization:

THEOREM I (LEWIS). *If the bookie's strategy does not proceed by conditionalization on the evidence, he leaves himself open to a dynamic Dutch book.*

Proof. Suppose that for some proposition (measurable set), q , in the bookie's probability space, and some member of his evidential partition, e , $Pr_1(q/e)$ unequal to $Pr_e(q)$ where Pr_e is the Pr_2 onto which STRAT maps e . Then either (1) $Pr_1(q/e) < Pr_e(q)$ or (2) $Pr_1(q/e) > Pr_e(q)$. If the first, let $d = Pr_e(q) - Pr_1(q/e)$. The bettor can now proceed as follows: At t_1 he proposes (i) the conditional bet to the

¹Reported in Teller (1973, 1976). Anyone who doubts the value of this argument because of qualms about the existence of a suitable prior should consider the arguments in Freedman and Purves (1969).

effect that if $e \& q$ the bookie pays him $\$ Pr_1(\text{not-}q/e)$; if $e \& \text{not-}q$ he pays the bookie $\$ Pr_1(q/e)$; if $\text{not-}e$ the bet is called off; and (ii) the sidebet on $\text{not-}e$ such that if e , he pays the bookie $\$ d Pr_1(\text{not-}e)$; if $\text{not-}e$, the bookie pays him $\$ d Pr_1(e)$. The bookie judges these bets as fair at t_1 . At t_2 if a member of the evidential partition other than e was revealed, the conditional bet is canceled and the cunning bettor wins the sidebet and gains $\$ d Pr_1(e)$. If e was the member of the evidential partition revealed at t_2 , then the bettor proposes an additional bet (iii) on q such that if q , he pays the bookie $\$ Pr_e(\text{not-}q)$; if $\text{not-}q$, the bookie pay him $\$ Pr_e(q)$. The bookie regards these as fair, according to Pr_e . The net result of (i) and (iii) is then that the bettor wins $\$ d$ no matter whether q or $\text{not-}q$. He has lost $\$ d Pr_1(\text{not-}e)$ on the sidebet (ii), giving him again a net gain of $\$ Pr_1(e)$ on (i), (ii), and (iii). Case (2) is similar.

In the dynamic case, as in the static one, the possibility of a Dutch book is the result of an underlying pragmatic inconsistency. If the bookie has a rule for changing degrees of belief of the kind under discussion, and is in a situation such as we have described where the rule is applicable, then the bettor can achieve the effect of a bet on q conditional on a member, e , of the evidential partition in one of two ways. He can make the conditional bet at t_1 either directly or by making a finite number of unconditional bets that achieve the same effect. Or he can simply resolve to wait until t_2 and make the bet on q at the revised rates if and only if e is the member of the partition announced. The second possibility is foreseeable by the bookie; it is based on his own rule for revising his degrees of belief. For the bookie to evaluate conditional bets consistently, his strategy must update his degrees of belief by conditionalization on the evidence.

2. Probability Kinematics. Can we build an adequate theory of evidence in which conditionalization is the only rule for updating degrees of belief? There is a tradition in epistemology that goes back to the Stoics that is favorable to an affirmative answer. According to this general view, there is a kind of evidence that is the foundation for all knowledge. When in possession of such evidence one is unmistakably led to certainty in the appropriate evidential proposition; that proposition bears “the mark of truth” that compels unqualified assent. All other propositions are evaluated in relation to these foundational ones. Various versions of this view have surfaced since Hellenistic times, with the mark of truth being conferred by the senses or intuition or some combination; the most recent version being promulgated by one wing of the logical positivist movement. On such a view, one simply updates by conditionalization as new

evidence comes in. If one's ultimate prior were also somehow given by intuition or reason, all of epistemology would be grounded in certainty. Such was, in its essentials, the program of Carnap's inductive logic, at least in the beginning.

But even leaving aside the question of the grand prior, it is hard to swallow the view that evidence comes as neatly packaged as the Stoics claim it does. Their doctrine was forcefully opposed in their own time by the academic skeptics who held that all knowledge is "probable" rather than certain, and in our own time by epistemologists as diverse as J. L. Austin, Sir Karl Popper, and Wilfrid Sellars. According to the strict skeptical view, conditionalization is *never* justified. One need not be a complete skeptic, however, to doubt that conditionalization is always the appropriate model. Even a positivist who believed in the possibility of an adequate language of sense data might not have that language in hand yet. Everyone has an interest in how belief revision in the light of evidence that does not render any proposition in the agent's language certain should be handled.

Addressing this problem, Richard Jeffrey (1965) suggested a generalization of the rule of conditionalization as a way to deal with some cases of uncertain evidence. Suppose $\{e_j\}$ is a partition all of whose members have positive initial probability, Pr_i . A subsequent probability, Pr_j , is said to *come from* Pr_i by *probability kinematics on* $\{e_j\}$ iff:

$$Pr_i(q/e_j) = Pr_j(q/e_j) \text{ for all } j \text{ and all } q.$$

Jeffrey uses the name "probability kinematics" to suggest the absence of informational forces that might deform the probabilities conditional on members of the partition. In statisticians' language the partition $\{e_j\}$ is a *sufficient partition* for the class of probability distributions that can come from Pr_i by probability kinematics on $\{e_j\}$, and a measurable function whose set of inverse images is the partition is a *sufficient statistic* for that class of probability distributions. (For a nice discussion, see Diaconis and Zabell 1982.)

Belief revision by probability kinematics is a natural generalization of conditionalization. The partition must, of course, be one appropriate to the evidence; the evidential event should give us information salient only to the relative probabilities of members of the partition. If the information gives one member of the partition final probability of one, we have belief change by conditionalization on that member of the partition. The set of probabilities that can come from a prior that gives every member of $\{e_j\}$ positive probability, by probability kinematics on $\{e_j\}$, is a convex set of which the probability measures that come by conditionalization on a member of $\{e_j\}$ are the extreme points.

There is an obvious question to ask: "Is there a dynamic coherence

argument for probability kinematics?” The answer is not quite obvious, because the whole point of probability kinematics is to deal with the sort of situation where there is no proposition in the agent’s language that represents the epistemic input of the evidential experience. How then, do we represent his strategy for changing degrees of belief? Let us consider these questions in the context of an example of the sort of situation that Jeffrey had in mind.

3. The Observation Game. Player A (the bookie) is shown a jellybean under dim light, and on the basis of this observational event may revise his prior probabilities of its color. Subsequently, he is told its true color by the gamemaster. Player A has three salient probability distributions: Pr_1 (before the observation); Pr_2 (after the observation); Pr_3 (after the gamemaster announces the true color); over a discrete probability space whose points represent $\langle \text{color, flavor} \rangle$ pairs. Sets of points in this space can be thought of as representing propositions about the bean in question. (I will use the ampersand, ‘&’, for set intersection and the dash, ‘-’ for set complement.) Player B (the cunning bettor) can make bets with player A at any of the corresponding times, t_1, t_2, t_3 , regarding the color and flavor of the bean. Player B doesn’t get to make any observations, but at each time he is allowed to know player A’s probability measure over this space at that time. He also gets to know player A’s *strategy* for changing degrees of belief.

3.1. Strategies. Allow me to begin with some heuristic considerations that motivate the definition of strategies. Player A comes into the game with an initial probability, Pr_1 , which is modified in response to the observation to yield Pr_2 . If the information that the observational event supplied were just that the true situation were in some set in his probability space, then we could require that his strategy specify his Pr_2 as a function of that given set (as in the strategy of conditionalization). But we are here interested in the case in which information conveyed by observation cannot be captured in this way. We suppose that the light is too dim, the probability space too crude, to allow for this possibility. Lacking such an observational proposition, we require at this point only that player A’s strategy specify a class of possible Pr_2 ’s that he takes to be permissible.

At t_3 player A learns just the color of the jellybean, and here he does have a set, COLOR, which captures what he learns. His strategy must specify his output, Pr_3 , as a function of his $\langle Pr_2, \text{COLOR} \rangle$. We need, however, some way of building into the specification of his strategy that his observation at t_2 is salient only to the partition of colors; of ruling out that he cheats by perhaps sniffing the bean when he is supposed to be

only observing its color. If the change from Pr_1 to Pr_2 reflects only the acquisition of information about color, then when at t_3 player A is told the true color of the jellybean, this should supersede whatever imperfect information about color he obtained by the act of observation in dim light. We build in the prohibition against extra illicit information by requiring that Pr_3 be a function of COLOR alone.

We will assume for simplicity that player A's Pr_1 and Pr_2 must give each atom of the probability space positive probability and that his Pr_3 gives each flavor positive probability. Player A must believe the gamemaster: $Pr_3(C) = 1$ for the color, C , that the gamemaster announced. Player A's strategy must address the possibility that any color may be the one announced whatever his Pr_2 . [(iv) below.]

Formally, *Player A's strategy* may be taken to be a set, STRAT, of quadruples: $\langle Pr_1, Pr_2, \text{COLOR}, Pr_3 \rangle$ such that: (o) Pr_1 , the prior that player A brings to the situation, is the same in each quadruple and gives each atom positive probability. (i) In each quadruple Pr_2 gives each atom positive probability and Pr_3 gives each flavor positive probability. (ii) The Pr_3 of a quadruple gives the color, which is the third component of that quadruple, probability 1. (iii) If two quadruples in STRAT agree on COLOR, they agree on Pr_3 . (iv) For every Pr_2 that occurs as the second coordinate of a quadruple in STRAT and every COLOR, there is a quadruple in STRAT whose second and third coordinates are, respectively, that Pr_2 and that COLOR.

Player B specifies a strategy that tells him what to bet at t_1, t_2, t_3 given player A's strategy and player A's probability measures up to the appropriate time. At any time, he can make a finite number of bets that player A considers fair or favorable at the time; or he may refrain from betting.

Formally, *Player B's Strategy* is an ordered triple of partial functions (partial because player B may decide not to bet), $\langle F_1, F_2, F_3 \rangle$, where F_1 where defined maps player A's STRAT onto a finite number of bets with non-negative expectation for A according to A's Pr_1 ; F_2 where defined maps $\langle \text{STRAT}, Pr_2 \rangle$ pairs onto a finite number of bets with non-negative expectation according to Pr_2 ; F_3 where defined maps $\langle \text{STRAT}, Pr_2, \text{COLOR}, Pr_3 \rangle$ quadruplets onto a finite number of bets that have non-negative expectation according to Pr_3 . A *Sequence of play* is a septuplet, $\langle Pr_1, B_1, Pr_2, B_2, \text{COLOR}, Pr_3, B_3 \rangle$ where $\langle Pr_1, Pr_2, \text{COLOR}, Pr_3 \rangle$ is in player A's strategy, and B_1, B_2 , and B_3 , respectively, are the bets or absence of bets that this quadruple elicits in the obvious way from player B's strategy. Let us say that player B's strategy *scores against* that of player A if (1) there is a sequence of play in which player B ends up at t_3 with bets whose net result is positive for him for all flavors (a Dutch book against A), and (2) for every sequence of play, player B's winnings are non-negative for every flavor. We will say that player A's strategy is *bulletproof* if no strategy for player B will score against it.

3.2. *Bulletproof Strategies and Sufficient Partitions.* We will say that the partition of colors is *Sufficient for A's Strategy*, or in Jeffrey's terminology that *A's strategy proceeds by probability kinematics on the partition of colors* just in case for each $\langle Pr_1, Pr_2, C, Pr_3 \rangle$ in A's STRAT, each color, C , and flavor, F :

$$Pr_1(F/C) = Pr_2(F/C) = (\text{where defined}) Pr_3(F/C).$$

THEOREM II. *It is a necessary condition for A's strategy in the Observation Game to be bulletproof that A's strategy proceed by probability kinematics on the partition of colors.*

Proof (1). Suppose that A's strategy set has as a member a quadruple, $\langle Pr_1, Pr_2, C, Pr_3 \rangle$ such that there is a flavor, F , such that $Pr_1(F/C)$ is unequal to $Pr_3(F/C)$. Then either $Pr_1(F/C)$ is less than $Pr_3(F/C)$ or $Pr_1(F/C)$ is greater than $Pr_3(F/C)$. Consider the first case, and let $e = Pr_3(F/C) - Pr_1(F/C)$. Player B can then end up at time t_3 with a Dutch book against player A. At t_1 he makes a complex betting arrangement with player A: He proposes that if $C \& F$, player A pay him \$ $Pr_1(F/C)$; if $C \& \neg F$, he pays player A \$ $Pr_1(F/C)$; if $\neg C$, then the bet is called off. (ii) He proposes in addition that if C , he pays player A \$ $ePr_1(\neg C)$; if $\neg C$, player A pays him \$ $ePr_1(C)$. Player A regards both bets (i) and (ii) as fair according to his Pr_1 . (The conditional bet, (i), can be constructed from a finite number of unconditional bets that he considers fair, as pointed out in section 2.) Then player B waits until time t_3 . If the gamemaster announces that the color is something other than C , (i) requires no payment and player B receives \$ $ePr_1(C)$ as a consequence of winning (ii). If, on the other hand, the color C is announced, he then makes a further arrangement with player A: (iii). If $C \& F$, he pays player A \$ $Pr_3(\neg F/C)$; if $C \& \neg F$, player A pays him \$ $Pr_3(F/C)$. Player A, by the lights of his current Pr_3 , will regard (iii) as fair because his current Pr_3 must be the same as the one in the quadruple we started with by the requirement that Pr_3 be a function of color alone. The net result of (i) together with (iii) is that player B wins \$ e no matter whether F or $\neg F$. He has lost \$ $ePr_1(\neg C)$ on (ii) reducing his winnings to \$ $ePr_1(C)$ as before. The argument for case (ii) is similar, with player B taking the opposite ends of bets (i) and (iii).

Proof (2). Suppose that player A's STRAT contains a quadruple, $\langle Pr_1, Pr_2, C, Pr_3 \rangle$ such that there is a flavor, F , such that $Pr_2(F/C)$ is unequal to $Pr_3(F/C)$. Player B then waits until t_2 . If the Pr_2 of the quadruple is not present at t_2 , player B does not bet. However, if it is present, he can end up at t_3 with a Dutch book against player A by betting just as in Proof (1) with Pr_2 substituted for Pr_1 and t_2 substituted for t_1 everywhere.

Proof (3). Suppose that player A's STRAT contains a quadruple $\langle Pr_1, Pr_2, C, Pr_3 \rangle$ such that there is a flavor, F , such that $Pr_1(F/C)$ is unequal to $Pr_2(F/C)$. Then either (i) $Pr_1(F/C)$ is unequal to $Pr_3(F/C)$ or (ii) $Pr_2(F/C)$ is unequal to $Pr_3(F/C)$ or (iii) $Pr_3(F/C)$ is undefined. In cases (i) and (ii), we have seen how the cunning bettor should bet under proofs (1) and (2), respectively. In case (iii), player A's STRAT must contain another quadruple with the same Pr_1 and Pr_2 in which $Pr_3(F/C)$ is well defined by the definition of a STRAT. This quadruple brings us back under case (i) or (ii).

Player B's strategy is this: He searches player A's STRAT for a quadruple of the kind described under proof (1). If he finds one, he takes the first one that he finds and bets relative to it as described under (1). If not, he searches for a quadruple as described under (2). If he finds one, he takes the first one that he finds and bets as described under (2). If not, he does not bet at all. If player A's strategy does not proceed by probability kinematics on the partition of colors, player B's strategy will score against him. Player B will either find a quadruple of the kind described under proof (1) or one of the kind described under (2). If the first, he surely ends up with a Dutch book against player A at t_3 ; if the second, he either ends up with a Dutch book against at t_3 or makes no bet at all.

Inspection of the proof of the foregoing theorem shows that player B knows the magnitude, e , of the discrepancy between player A's strategy and probability kinematics before he makes his first bet; at t_1 in case 1 and at t_2 in case 2. He could, then, modify the strategy given by inflating the stakes by a factor of $\$ N/e Pr_1(C)$ to fit the e and C involved, where N is an arbitrarily large positive real. Such a modified strategy for player B yields him a payoff of $\$ N$ in each course of play in which he bets at all. In case 1, he is assured of such a payoff. In case 2, he is assured of such a payoff if the offending Pr_2 assignment appears, and does not bet otherwise.

3.3. Bulletproof Strategies and Potential Centering. Because we allow player A to pick his set of possible Pr_2 's in the specification of his strategy, there is another way in which his strategy can fail to be bulletproof. He might, for example, choose his possible Pr_2 's such that for each one, $Pr_2(C\&F)$ is less than $Pr_1(C\&F)$, for some particular color-flavor pair. In that case, player A's probability for $C\&F$ would have to move in a foreseeable direction, a fact that player B could exploit by betting against $C\&F$ at t_1 and buying back the bet at a profit at t_2 .

If player A's strategy is such that for some proposition, Q , in the color-flavor space, $Pr_1(Q)$ greater than $Pr_2(Q)$ for all Pr_2 's allowed by the strategy, we will say that player A's strategy is *OUT*. If his strategy is such

that for some Q in the color-flavor space, there is a positive e such that $Pr_1(C\&F) - Pr_2(C\&F)$ not less than e for all Pr_2 's, we will say that his strategy is *DISTANT*. If his strategy is such that for some such Q , for some Pr_2 , $Pr_1(Q)$ is greater than $Pr_2(Q)$ and for no Pr_2 is Pr_1 less than $Pr_2(Q)$, we will say that his strategy is *NOT-IN*.

LEMMA. *If player A's strategy is DISTANT, player B can always end up with a Dutch book at t_2 which assures him of whatever payoff he chooses in advance at t_1 . If his strategy is OUT player B can always end up at t_2 with a Dutch book against him. If his strategy is NOT-IN, player B's strategy can score against his.*

Proof is obvious.

Being *not-IN*, *OUT*, and *DISTANT* are increasingly serious defects for player A's strategy, the least serious of which still prevents his strategy from being bulletproof. If the pr_1 of player A's strategy can be represented as a mixture of the pr_2 's of his strategy, his strategy cannot be *OUT* or *DISTANT*. If his strategy has a finite number of pr_2 's, and the pr_1 of his strategy can be represented as a mixture of his pr_2 's in which each pr_2 has positive weight, then his strategy isn't *not-IN*, (because if it were *not-IN*, the hypothesized mixture would give Q probability less than pr_1). The question of centering will be raised again in a context where it is possible to bet on the pr_2 's.

4. Higher-Order Probabilities and Absolute Dutch Books. The foregoing discussion of the observation game focused on the concept of a *bulletproof strategy*. The state of being bulletproof is a strong coherence property, that is, a guarantee against a certain kind of conditional Dutch book. A *conditional Dutch book* is a set of bets such that they result in a net loss to the bookie (player A) if the condition is realized, and result in zero net transaction otherwise. A conditional Dutch book on a condition of positive probability can always be turned into an unconditional Dutch book by making the appropriate sidebet against the condition; but in the observation game, the conditions in question include the specification of an observational probability distribution, pr_2 . If player A has well-defined probabilities over courses of play in the game, the conditions of the conditional Dutch books that constitute *scores* against his strategy may well for him have probability zero. I did not assume in the foregoing discussion that the observer in the observation game had any such probabilities. The question arises as to what more can be said if he does. [The subsequent discussion owes much to the important work of Armendt (1980).]

4.1. Probability Kinematics. If for every sequence of play, player B's strategy results in bets at t_3 whose net result is positive for him for all

flavors, we will say that player B's strategy constitutes an *unconditional dynamic Dutch book* against that of player A. Under reasonable conditions, we can show that a strategy of belief change by probability kinematics on the observational partition almost everywhere (with respect to Pr_1) is a necessary condition for avoiding an unconditional Dutch book.

Let us modify the observation game to get a version with higher-order probabilities as follows: Player A's three probability measures, Pr_1 , Pr_2 , and Pr_3 , are over an enlarged probability space. This space is the product of (a) the original discrete space of color-flavor pairs of "The Observation Game" and (b) the space of probability measures over space (a) with Lebesgue measurable sets in the appropriate n -space serving as the measurable sets. (These probability measures are to be interpreted as player A's Pr_2 about which player A is uncertain at t_1 .) We assume that player A has a fixed initial probability, Pr_1 . Since Pr_3 is a function of COLOR, we can take a set specified by a given Pr_2 and color as tantamount to a specification of a quadruple, $\langle Pr_1, Pr_2, \text{COLOR}, Pr_3 \rangle$, in player A's strategy. We will speak loosely of probabilities of sets of quadruples in this sense.

Let us call a quadruple, $\langle Pr_1, Pr_2, \text{COLOR}, Pr_3 \rangle$, in Player A's strategy VULNERABLE if there is a color, C , and flavor, F , such that $Pr_i(F/C)$ is defined for $i = 1, 2, 3$ and it is not the case that it takes on the same value for $i = 1, 2, 3$.

LEMMA 1. *If the set of VULNERABLE quadruples in player A's strategy has positive probability, player B can make an unconditional Dutch book against player A.*

Proof. We showed that if in the Observation Game a VULNERABLE quadruple is played out by player A, player B ends up with a Dutch book at t_3 that guarantees him fixed positive winnings, $\$K$, no matter what the flavor turns out to be. If the set of VULNERABLE quadruples has positive measure, then player B could guarantee himself an unconditional Dutch book by making a sidebet against the set of vulnerable quadruples, V . He offers to pay player A $\$K Pr_1(-V)$ if V , if player A will pay him $\$K Pr_1(V)$ if $-V$, guaranteeing himself winnings of $\$K Pr_1(V)$.

A quadruple disagrees with probability kinematics—is UNKINEMATIC—if for some color, C , and flavor, F , $Pr_1(F/C)$ is unequal to $Pr_2(F/C)$, or $Pr_2(F/C)$ is unequal to $Pr_3(F/C)$, or both. The UNKINEMATIC quadruples are the VULNERABLE ones together with ones in which $Pr_1(F/C)$ is unequal to $Pr_2(F/C)$ but $Pr_3(F/C)$ is undefined.² We

²Because some other color is the third coordinate of the quadruplet.

must deal with the case in which player A gives his set of UNKINEMATIC quadruples positive probability and gives his set of VULNERABLE quadruples zero Pr_1 . Note that in this case, the agent has the optimistic belief ($Pr_1 = 1$) that nature will render him invulnerable to the consequences of any lack of kinematicity ($Pr_1(F/C)$ unequal to $Pr_2(F/C)$) by conveniently canceling the conditional bets by producing a color other than C at t_3 . He cannot square this with the requirement of the game that each Pr_2 in player A's strategy give each Pr_2 positive probability.

LEMMA 2. *If player A gives his set of UNKINEMATIC quadruples positive Pr_1 , but gives his set of VULNERABLE quadruples zero Pr_1 , then player B can make an unconditional dynamic Dutch book against him.*

Proof. Since there are only a finite number of colors, if the set of unkinematic quadruples has positive measure and the set of vulnerable quadruples has zero measure, then there must be some particular color, C , such that the set of unkinematic quadruples with $Pr_1(F/C)$ unequal to $Pr_2(F/C)$ for some F has positive measure and such that the subset of quadruples that contain C has zero measure. Call these sets UNK_C and VUL_C , respectively. If UNK_C has positive measure, then there is some positive number, e , such that the set of quadruples that contain $Pr_2(C)$ at least as great as e has positive measure,³ since the observation game requires that for each color, each Pr_2 gives it positive measure. Call this set UNK'_C and the subset in which C does come up, VUL'_C . Now consider the following fair-bet strategy: At t_1 B proposes (1) that player A pay him \$1 if VUL'_C and that he pay player A nothing if not. Player A considers this fair since he gives VUL'_C zero Pr_1 . If at t_2 player A is in UNK'_C player B offers (2) to sell back bet 1 for $\$e$. Player A now considers this fair or favorable. B's strategy so far constitutes a Dutch book conditional on UNK'_C . But since UNK'_C has positive Pr_1 , this can be turned into an unconditional Dutch book by a suitable sidebet (3) against UNK'_C at t_1 as before.

We will say that player A's strategy *proceeds by probability kinematics almost everywhere in Pr_1* if his Pr_1 gives his set of UNKINEMATIC quadruples probability zero. Lemmas 1 and 2 do not quite establish that this property is a necessary condition for avoiding an unconditional dynamic Dutch book because nothing has been said that requires these sets to be measurable. However, it is most reasonable to add the requirement that these sets be measurable if we regard player A's strategy not just as

³We use countable additivity here.

a set theoretical entity, but as a strategy that is specifiable in an effective manner.

THEOREM III. *If player A's sets of VULNERABLE and UNKINEMATIC quadruples are measurable and his strategy does not proceed by probability kinematics almost everywhere in Pr_1 , player B can make an unconditional dynamic Dutch book against him.*

4.2. Centering. In the present context we can say more about centering. Goldstein (1983) and van Fraassen (1984) both argue that coherence requires that pr_1 be equal to the prior expectation of pr_2 considered as a random variable. That is, pr_1 must be the "center of mass" of the possible pr_2 's where pr_1 is taken as the measure of mass. I will take some liberties in adapting their arguments to the present setting.

Consider first the case in which there are a finite number of possible pr_2 's, each with positive pr_1 :

THEOREM IVA (VAN FRAASSEN). *The bookie's prior, pr_1 , must be such that for any proposition, Q , of the color-flavor space, and any number, a , such that $pr_1[pr_2(Q) = a]$ is positive:*

$$pr_1[Q/pr_2(Q) = a] = a$$

or the bettor can make an unconditional dynamic Dutch book against him.

Proof. Suppose that $pr_1[Q/pr_2(Q) = a] = a - e$ for some positive e . Then at t_1 , (i) the bettor makes a conditional bet with the bookie to the effect that if $pr_2(Q) = a \& Q$, the bookie pays him $\$1 - (a - e)$; if $pr_2(Q) = a \& \neg Q$, he pays the bookie $\$a - e$; if it is not the case that $pr_2(Q) = a$, the bet is called off. (ii) In addition, he makes a sidebet to the effect that if $pr_2(Q) = a$, he will pay the bookie $\$e pr_1[\neg pr_2(Q) = a]$; if $\neg pr_2(Q) = a$, the bookie will pay him $\$e pr_1[pr_2(Q) = a]$. Then he waits until t_2 . If it is not the case that $pr_2(Q) = a$, the conditional bet (i) is canceled and he wins the sidebet (ii). If it is the case that $pr_2(Q) = a$, then (iii) he bets on Q such that if Q , he pays the bookie $\$1 - a$; if $\neg Q$ the bookie pays him $\$a$. Then the net effect of (i) and (iii) is that he wins $\$e$ no matter whether Q or not, which is reduced by his loss of the sidebet (ii). In any event, his net gain from (i), (ii), (iii) is $\$e pr_1[pr_2(Q) = a]$. If $pr_1[Q/pr_2(Q) = a] = a + e$ for some positive e , the bettor takes the other end of bets (i) and (iii). It follows immediately that $pr_1(Q) = E_1[pr_2(Q)]$; that the prior probability of any proposition in the color-flavor space is the prior expectation of its posterior probability. To get the same result in a more general setting, we need to bet on pr_2 falling within an interval.

THEOREM IVB (GOLDSTEIN). *The bookie's prior, pr_1 , must be such that for any proposition, Q , of the color-flavor space, and any closed interval, I , such that $pr_1[pr_2(Q) \text{ in } I]$ is positive:*

$pr_1[Q/pr_2(Q) \text{ in } I]$ in I

or the bettor can make an unconditional dynamic Dutch book against him.

Proof. As in IVA.

The principles that figure in these theorems are sometimes referred to in the philosophical literature as “Miller” principles. One or both are discussed in Goldstein (1983), Skyrms (1980a, 1980b), and van Fraassen (1984).

5. Converse Dynamic Dutch Book Arguments. A dynamic Dutch book argument shows that if you have a rule for updating probabilities in a certain type of situation, and your rule does not meet certain standards, some kind of Dutch book can be made against you. David Lewis's Dutch book argument for conditionalization and my generalization of that argument to probability kinematics are examples. Converse Dutch book arguments show that if the standards are met, no such Dutch book can be made. A Dutch book theorem has little force if its correlative converse is not true. For the *dynamic* case, converse Dutch book arguments have not been pursued in the philosophical literature.

5.1. Conditionalization. Consider the conditions for Lewis's dynamic Dutch book argument for conditionalization. The bookie has a prior at time 1 such that every member of a finite partition, p_i , has positive prior probability. At time 2, he learns the true member of the partition and changes his probability assignment. His rule for change is a function that maps the pair (prior, member of the partition learned) to his posterior. A bettor gets to know the bookie's strategy, gets to know the true member of the partition, when the bookie does, and gets to make a finite number of bets at each time that the bookie considers fair at that time. The bettor's strategy is a pair of functions, one from the bookie's probability assignment at time 1 to a finite set of bets that the bookie considers fair at that time, and a second from the revealed member of the partition and the bookie's probability at time 2 to a finite set of bets that he considers fair at that time. The bookie's strategy together with that of the bettor determine a *payoff function*, which gives the net payoff of all bets for each state of the world. Lewis's argument shows that if the bookie's strategy is not to update by conditionalization, the bettor can choose a strategy such that the two yield a payoff function for the bettor that is positive everywhere.

THEOREM V. *If the bookie's strategy is to update by conditionalization, then there is no bettor's strategy that constitutes a Dutch book against him.*

Proof. Any payoff function that the bettor can achieve against the conditionalizing strategy by betting at t_1 and t_2 can be achieved by an alternative strategy that relies only on a finite number of bets all made at time 1. For every bet that the bettor's original strategy makes at t_1 , the modified strategy will make at t_1 . For every bet on Q that the bettor's original strategy makes at t_2 , if P_i is the true member of the partition, the modified strategy substitutes a bet on Q conditional on P_i made at t_1 , (which can be attained by a finite number of unconditional bets made at t_1). Given that the bookie is a conditionalizer and that the true member of the partition is announced at t_2 , the payoff must be the same. A dynamic Dutch book can therefore be made against the conditionalizer only if a static Dutch book can be made against him at t_1 by a finite number of bets that he considers fair. This we know to be impossible given that he respects the probability calculus at t_1 since the expectation of the sum of a finite number of random variables, each with zero expectation, is zero; while the expectation of a betting arrangement which constitutes a Dutch book must be negative.

5.2. Probability Kinematics. In the discussion of the Observation Game, it is shown that the bookie's strategy must proceed by probability kinematics on the partition, to be "bulletproof." This is somewhat weaker than an unconditional Dutch book. The converse would then be somewhat stronger than the converse to a Dutch book; that is, that belief change by probability kinematics is "bulletproof."

Consider first the case of "coarsegrained observation by candlelight" where the bookie in the observation game has only a finite number of pr_2 's possible. We will argue that if the bookie has a strategy of belief change by probability kinematics on the partition of colors, and if his strategy meets an INTERIOR condition, we can embed the observation game in a bigger Lewis game such that the bookie's strategy in the original game fails to be bulletproof only if a strategy of conditionalization in the Lewis game can have a dynamic Dutch book made against it. This is impossible by the results of the previous section.

The larger game is constructed along the lines suggested by section 4. The bookie has a larger probability space, which is the product of a space of N elements $\langle p_1, \dots, p_n \rangle$ (for the N possible Pr_2 's over the color-flavor space) with the original color-flavor space. The bookie has a prior over this space that gives each atom positive probability. At time t_2 a p_i is

announced and the bookie must move to a new probability, pr_2 , by a rule which makes pr_2 a function of the announced p_i . Pr_2 must be a non-zero for each color–flavored pair. At time t_3 a color is announced and the bookie must move to a new probability, Pr_3 , by a rule that makes Pr_3 a function of Pr_2 and COLOR. This larger game is the composition of two Lewis games for the move from t_1 to t_2 and the move from t_2 to t_3 , for which we know that the strategy of conditionalization is both a necessary and sufficient condition for immunity from Dutch book.

The possible probability assignments over a discrete space of m objects can be thought of as represented by the points in an $m-1$ dimensional polyhedron in m dimensional space. The n p_i 's that are considered by the bookie's strategy in the original game to be possible pr_2 's are to be thought of as n points in the interior of such a polyhedron (the interior because they must all give each color–flavor pair non-zero probability.) The INTERIOR condition on the bookie's strategy in the original game is that his pr_1 over the color–flavor space be in the interior of the convex hull of the p_i 's. If the INTERIOR condition is met, it follows that the bookie's pr_1 can be represented as a non-trivial mixture of his p_i 's, that is, one which gives each p_i non-zero mixing coefficient.

The bookie's pr_1 in the smaller game is extended to PR_1 over the probability space of the bigger game as follows:

- (1) Let $PR_1(C\&F/p_i) = p_i(C\&F)$ for each C, F, i .⁴
- (2) Distribute PR_1 over the p_i 's such that each p_i gets non-zero probability and $PR_1(C\&F) = pr_1(C\&F)$ for each C, F ; that is, represent pr_1 as a non-trivial mixture of the p_i 's.

Notice that we then have:

- (3) $PR_1(F/C\&p_i) = PR_1(p_i)PR_1(F\&C/p_i)/PR_1(p_i)PR_1(C/p_i)$
 $= PR_1(F\&C/p_i)/PR_1(C/p_i)$
 $= p_i(F\&C)/p_i(C)$ by (1)
 $= p_i(F/C)$
 $= pr_1(F/C)$ because the p_i 's figure in a strategy of belief change by probability kinematics on the partition of colors in the original game.
 $= PR_1(F/C)$ by (2).

((3) and its relation to probability kinematics are discussed in Armendt (1980), Good (1981), and Skyrms (1980a, 1980b).)

Now let a bookie's strategy in the big game be belief change by con-

⁴In a slight abuse of notation, ' $p_i(C\&F)$ ' is used for ' $pr_2(C\&F)$ ' according to the i^{th} possible pr_2 '. The construction models the random variable pr_2 as probability conditional on the partition of the p_i 's.

ditionalization. His behavior *vis-à-vis* the color–flavor space will be indistinguishable from that of a bookie pursuing a strategy of belief change by probability kinematics in the smaller game. The initial probabilities at t_1 are the same by (2). At t_2 when the p_i resulting from the observation interaction becomes known the probabilities are the same by (1). At t_3 when the true color becomes known as well, the probabilities are the same by (3). Consequently, a bettor’s strategy that will *score* against a strategy of belief change by probability kinematics that meets the INTERIOR condition in the observation game will *score* against the strategy of conditionalization in the bigger game. But a *score* against the conditionalizing strategy in the bigger game could be turned into an unconditional dynamic Dutch book if the course of play leading to that *score* has positive prior probability, by making a suitable sidebet against that course of play in the standard way. Each course of play does have positive prior probability. $PR(p_i)$ is positive by (2). $PR_1(C/p_i)$ is positive by (1) together with the requirement of the observation game that each possible probability at t_2 gives each atom of the color–flavor space positive probability. We know that a bettor’s strategy that constitutes an unconditional dynamic Dutch book against the conditionalization strategy in the larger game is impossible from the previous section. So we have shown:

THEOREM VI. *In the case in which the bookie has a finite number of possible pr_2 ’s in the observation game, if his strategy proceeds by probability kinematics on the partition of colors and meets the INTERIOR condition, his strategy is bulletproof.*

What can we say about the general case where the bookie’s strategy may recognize infinitely many pr_2 ’s as possible? Call a strategy *catholic* if it contains pr_2 ’s that distribute probabilities among the colors in every possible way consistent with giving each color non-zero probability. If we idealize our observer so that he has no trouble dealing with arbitrary real numbers, a catholic strategy will be for him a sign of open-mindedness about what observation will bring. For each possible pr_1 (which gives each atom of the color–flavor space positive probability) there is a unique *catholic probability kinematics* strategy for the bookie. (The quadruples have the given pr_1 as first coordinate. The second coordinate is a pr_2 determined by an arbitrary pr_2 for color, extended to flavor by probability kinematics. The third coordinate is an arbitrary color. The fourth coordinate is a pr_3 gotten from the second coordinate by conditionalization on the third.)

THEOREM VII. *A catholic strategy of belief change by probability kinematics in the observation game is bulletproof.*

Proof. Suppose that it is not bulletproof. Then some bettor’s strategy can score against it. The bettor’s strategy scores against that of the

bookie if (1) there is a sequence of play in which the bettor ends up at t_3 with bets whose net payoff is positive for him (negative for the bookie) for every flavor, and (2) for every sequence of play, the bettor's winnings are non-negative for every flavor. So if a bettor's strategy will score against the bookie's infinite strategy, it will score against any finite substrategy of that strategy that includes (one of) the quadruples of the type described under (1). Let $q = \langle pr_1, pr_2, C, pr_3 \rangle$ be such a quadruple relative to the hypothesized score against the infinite strategy. There is a finite substrategy of the original strategy that meets the INTERIOR condition. Since pr_1 gives each color-flavor atom non-zero probability, it is in the interior of the convex hull of the probabilities that can be gotten from pr_1 by conditionalization on a color. Consider the probability measures that can be gotten by from pr_1 by "almost conditionalizing" on colors, that is, by probability kinematics on the partition of colors that gives the color "almost conditionalized on" probability $1-e$. Call these the V 's. Given pr_1 we can choose e small enough so that pr_1 is in the interior of the convex hull of the V 's. As our possible pr_2 's of the finite substrategy, we take the V 's together with the pr_2 of the quadruple q . The finite substrategy consists of all the quadruples in the original strategy having one of these pr_2 's. This is a finite number because there are only a finite number of possible colors, C , and Pr_3 is a function of C . So if the bettor's strategy will score against the original infinite strategy, it will score against this finite substrategy, but this is impossible by the previous theorem.

6. Diachronic Coherence and Probability Kinematics. Conditionalization is such a natural way to for updating degrees of belief that Hacking needed to remind us that it required a justification. When it is regarded as a *rule* or *strategy* applicable to a certain sort of situation that is commonly approximated by our experience, a justification is forthcoming. In such situations, adoption of any alternative rule, strategy, or habit leads an agent to dynamic incoherence regarding conditional bets. In effect, he adopts two distinct fair betting quotients for conditional bets depending on how they are described, where the equivalence of the descriptions is a simple consequence of Boolean logic together with his own rule. In such situations, adoption of the rule of conditionalization guarantees coherence.

But not every learning situation is of the kind to which conditionalization applies. The situation may not be of the kind that can be correctly described or even usefully approximated as the attainment of certainty by the agent in some proposition in his degree of belief space. The rule of belief change by probability kinematics on a partition was designed to apply to a much broader class of learning situations than the rule of conditionalization. In those situations for which it was designed, it preserves

the virtues of conditionalization.⁵ It is coherent and any rule which conflicts with it is not.

This is not to say that we can build an adequate epistemology solely on the rule of belief change by probability kinematics. That rule has its own limits of applicability; its justification here occurs within the context of the observation game. Other models of learning situations are possible, and each poses the question: "What rule or rules for belief change are dynamically coherent for this sort of situation?"

REFERENCES

- Adams, E. (1962), "On Rational Betting Systems", *Archive für Mathematische Logik und Grundlagenforschung* 6: 7–29, 112–28.
- Armendt, B. (1980), "Is There a Dutch Book Argument for Probability Kinematics?" *Philosophy of Science* 47: 583–88.
- Diaconis, P., and Zabell, S. (1982), "Updating Subjective Probability", *Journal of the American Statistical Association* 77: 822–30.
- de Finetti, B. (1937), "La Prévission: ses lois logiques, ses sources subjectives", *Annales de l'Institut Henri Poincaré* 7: 1–68. (Translation published as "Foresight: its logical laws, its subjective sources", in *Studies in Subjective Probability*, Henry E. Kyburg, Jr. and H. Smokler (eds.). Huntington NY: Krieger, 1980.)
- Freedman, D. A., and Purves, R. A. (1969), "Bayes Method for Bookies", *The Annals of Mathematical Statistics* 40: 1177–86.
- Goldstein, M. (1983), "The Prevision of a Prevision", *Journal of the American Statistical Association* 78: 817–19.
- Good, I. J. (1981), "The Weight of Evidence Provided by Uncertain Testimony or from an Uncertain Event", *Journal of Statistical Computation and Simulation* 13: 56–60.
- Hacking, I. (1967), "Slightly More Realistic Personal Probability," *Philosophy of Science* 34: 311–25.
- Jeffrey, R. (1965), *The Logic of Decision*. New York: McGraw-Hill. Second edition. Chicago: University of Chicago Press, 1983.
- Ramsey, F. P. (1931), "Truth and Probability", in *The Foundations of Mathematics and Other Essays*, R. B. Braithwaite (ed.). New York: Harcourt Brace. (Reprinted in *Studies in Subjective Probability*, Henry E. Kyburg, Jr. and H. Smokler (eds.), Huntington NY: Krieger, 1980.)
- Skyrms, B. (1980a), *Causal Necessity*. New Haven: Yale University Press, Appendix 2.
- . (1980b), "Higher-Order Degrees of Belief", in *Prospects for Pragmatism*, D. H. Mellor (ed.). Cambridge: Cambridge University Press.
- 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*, W. Harper and C. Hooker (eds.). Dordrecht: D. Reidel, pp. 205–53.
- van Fraassen, B. (1984), "Belief and the Will", *Journal of Philosophy* 81: 235–56.

⁵This includes some other virtues, in addition to dynamic coherence. See Teller (1973, 1976).