The Coherence Argument Against Conditionalization

Matthias Hild
Christ's College, Cambridge

18 January, 1998

Abstract. I re-examine Coherence Arguments (Dutch Book Arguments, No Arbitrage Arguments) for diachronic constraints on Bayesian reasoning. I suggest to replace the usual game–theoretic coherence condition with a new decision–theoretic condition ('Diachronic Sure Thing Principle'). The new condition meets a large part of the standard objections against the Coherence Argument and frees it, in particular, from a commitment to additive utilities. It also facilitates the proof of the Converse Dutch Book Theorem.

I first apply the improved Coherence Argument to van Fraassen's (1984) Reflection principle. I then point out the failure of a Coherence Argument that is intended to support Conditionalization as a naive, universal, update rule. I also point out that Reflection is incompatible with the universal use of Conditionalization thus interpreted. The Coherence Argument therefore defeats the naive view on Bayesian learning that it was originally designed to justify.

Key words: Bayesianism, Dutch Book Arguments, Conditionalization, Jeffrey's Rule, Reflection.

Introduction

This paper rehearses what appears to be the strongest argument in support of certain diachronic constraints on Bayesian reasoning. The argument runs as follows: 'If you violate the diachronic constraint under discussion, then the series of your successive decisions will suffer from diachronic incoherence.' Such Diachronic Coherence Arguments (also referred to as 'Dutch Book Arguments' or 'No Arbitrage Theorems') have first been proposed in an attempt to establish Conditionalization — or, at least, a naive interpretation thereof — as a universally applicable update rule. More recently, van Fraassen (1984) has used the Coherence Argument in defense of his Reflection principle. Coherence Arguments proceed in two steps. The first, non–technical, part of the argument proposes a coherence condition for rational decision making. Their second, technical, part assumes that decisions are based on the maximization of expected utility. It consists of a so-called 'Dutch Book Theorem' and a 'Converse Dutch Book Theorem'. These theorems attempt to relate the coherence condition from the first part of the argument to certain probabilistic principles, and *vice versa*. These are the probabilistic principles to be supported by the Coherence Argument.

Both parts of these arguments have come under attack. In Section 2, I will defend the Coherence Argument for Reflection from some unfair criticism. I will firstly argue that most problems with the technical part of the arguments disappear if the coherence condition is not formulated game–theoretically but decision–theoretically. Instead of the usual game theoretic condition, I suggest a diachronic version of Savage's (1954) Sure Thing Principle for coherent decision making. The usual game theoretic argument merely helps to illustrate the importance of this decision–theoretic condition.[1] The shift towards a decision–theoretic notion of diachronic coherence has the additional advantage of simplifying the proof of the Converse Dutch Book Theorem.

In Section 2, I secondly address criticism that is directed against the non-technical part of the argument, namely the notion of coherence and its importance. Some authors (e.g. Maher (1993)) argue that diachronic coherence looses its importance when it conflicts with the desire to limit the effects of future irrational judgements.[2] I will argue that this criticism overlooks that diachronic coherence is part of a framework which consciously confines itself to ideally rational agents. Once we abandon this framework of ideal rationality, we can and must re-formulate both the Coherence Argument and the Reflection principle (cf. Jeffrey's (1988) Refined Reflection and considerations in Spohn (1978)).

In Section 3, I examine the Coherence Argument for Conditionalization more closely. This application of the Coherence Argument is normally accompanied by a perhaps naive, but *prima facie* appealing, account of Bayesian learning from evidence. On this account, a Bayesian reasoner uses a single update rule (Conditionalization) in all possible learning situations regardless of the epistemic process that has generated her evidential input.[3] In connection with the Coherence Argument, this view has, for instance, been criticized by Shafer (1985). The Coherence Argument slips information about the process that generates the reasoner's evidential input into its technical apparatus. Briefly, the Coherence Argument is restricted to those special learning situations where the agent can be certain that she will either reliably learn that A or that $-A$. This I call the 'Partitioning Assumption'.

In Section 4, I point out that Reflection validates the use of Conditionalization precisely when this Partitioning Assumption holds. Whenever the reasoner distinguishes between the evidence that A has occurred and the circumstances in which she obtains this evidence, the Partitioning Assumption may well fail and Bayesian reasoning grounded on Reflection will then conflict with naive Conditionalization. The Coherence Argument in support of Reflection therefore defeats the naive account of Bayesian learning and the interpretation of Conditionalization as a universally applicable update rule. According to Bayesian folklore, a universally applicable update rule is the antidote that counter-acts the subjectivism arising from a free choice of prior probabilities. This view is untenable. Bayesian subjectivism not only holds for prior probabilities but also extends to the choice of update rules for these prior probabilities. This conclusion is by no means new, but I hope to show that it is a consequence of the very Coherence Argument that was designed to justify the naive interpretation of Conditionalization.

Let me finally mention the probably most fundamental criticism of the Coherence Argument which outrightly questions the significance of of any notion of diachronic coherence (e.g., Christensen (1991)). This criticism raises issues of personal identity across time that I do not address here. I am, however, inclined to think that my decision theoretic coherence condition is as good a prescriptive condition as we can ever hope to find. In my opinion, it is not necessary to deny the importance of diachronic coherence in order to avoid a universal commitment to Conditionalization. If methodological freedom from Conditionalization should be the secret motive behind the rejection of diachronic coherence, my analysis shows that one can have one's cake and eat it (or rather, another cake).

1. The Probabilistic Model

Probabilities

In the present framework, opinions (or, probabilistic information) are represented by unique probability measures. Let $\mathcal{A}$ be a σ–algebra over a non–empty set of possible worlds Ω. I shall refer to $\mathcal{A}$ as the reasoner's 'language' since $\mathcal{A}$ represents the reasoner's ability to discriminate between different possible worlds. Let $\mathcal{P}_{\mathcal{A}}$ be the set of all σ–additive probability measures on $\mathcal{A}$. Let I be a diachronic index that is totally ordered.[4] A *diachronic trajectory* is a sequence $\langle P_i \rangle_{i \in I}$ ($P_i \in \mathcal{P}_{\mathcal{A}}$ for all $i \in I$) that describes the diachronic changes in probabilities over different stages of reasoning. The totality of such sequences describes the possible trajectories along which the reasoner's opinion can evolve. For any (bounded) random variable $X : \Omega \rightarrow \Re$, let $E_i(X) := E_{P_i(.)}(X) := \int_\Omega X(.)\, dP_i$ be the expectation of X at stage i. In order to simplify the discussion, I will avoid using integration theory and restrict the discussion to probability measures that assign a non–zero probability only to countably many events from $\mathcal{A}$.

Conditional probabilities $P_i(.|.)$ may either be reduced to (regular) unconditional probabilities (cf. (1) below) or be axiomatized independently.[5] In either case, they are required to merge with unconditional probabilities as follows:

$$\text{(CP)} \qquad P_i(A \cap B) = P_i(A|B) \cdot P_i(B)$$

Hence,

$$P_i(A|B) = \frac{P_i(A \cap B)}{P_i(B)}, \qquad \text{if} \quad P_i(B) > 0. \tag{1}$$

Conceptually and formally, conditional probabilities must not be confused with updated probabilities. It is a non–trivial claim to identify new, updated, probabilities with old conditional probabilities.

Auto–epistemology

Let '$\mathrm{SP}_i(Q)$' express the event that the reasoner's *s*ubjective *p*robability measure at i is Q. (We postulate that $\mathrm{SP}_i(Q) \in \mathcal{A}$, for all $Q \in \mathcal{P}_\mathcal{A}$ and all $i \in I$.) With the introduction of this auto–epistemic vocabulary into $\mathcal{A}$, a probabilistic reasoner can now form opinions and expectations about her own opinions. Consider two diachronic trajectories $\langle P_i \rangle_{i \in I}$ and $\langle Q_i \rangle_{i \in I}$ that intersect at i (i.e., $P_i = Q_i$). For $i \leq j$, Q_j then represents one possible way in which the reasoner's probability measure P_i may evolve at j. The diachronic principle of Reflection suggests that the reasoner's expectations about the future evolution of her probabilities should be in balance with her present opinion (cf. van Fraassen (1984)). For all diachronic trajectories $\langle P_i \rangle_{i \in I}$ and $\langle Q_i \rangle_{i \in I}$ with $P_i = Q_i$:[6]

$$\text{(Reflection)} \qquad P_i(A|\mathrm{SP}_j(Q_j)) = Q_j(A), \qquad i \leq j.$$

If conditional probabilities are by definition reduced to the ratio of unconditional probabilities, Reflection can only apply when $P_i(\mathrm{SP}_j(Q_j)) > 0$. When I need to emphasize that Reflection is restricted to cases where $P_i(\mathrm{SP}_j(Q_j)) > 0$, I will also speak of 'Non–Zero Reflection'.

Reflection implies that the reasoner is auto–epistemically transparent to herself in that she knows exactly what opinion she currently holds:[7]

$$\text{(AE–Transparency)} \qquad P_i(\mathrm{SP}_i(Q_i)) = \begin{cases} 1, & \text{if } P_i(.) = Q_i(.) \\ 0, & \text{otherwise} \end{cases}$$

Auto–Epistemic Transparency essentially requires the reasoner to have a Cartesian awareness of her own opinion. Hild/Jeffrey/Risse (1998) mention a variant of Reflection that does not imply transparency.[8] In the present setting, the transparency condition plays, however, an important role in the proof of the Dutch Book Theorem for Reflection.

Conditionalization

Bayesian update rules offer an account of how prior probability measures (should) change as a result of new input or evidence.[9] We start with the set $\mathcal{EV}$ of possible evidential input with which the reasoner may be presented. Any item $e \in \mathcal{EV}$ is intended to represent a piece of total evidence that may be obtained by the agent. Let us first think of possible evidential input as a primitive notion. In particular, let us not yet assume that evidential input necessarily has the form of propositions.

On what I shall call the *'naive view'*, all that counts in the application of a given update rule is the actually received piece of evidence. Update rules do not consider the epistemic processes that connect the reasoner with her environment and generate the reasoner's evidence. They do not consider what evidence the reasoner would have received under circumstances different from the actual. Nor does an update rule in this sense depend on any pre-conditions on the type of process that generates the evidential input (apart perhaps from an implicit reliability assumption that excludes scepticism). Formally, an *update rule* is therefore a function $\star : \mathcal{P}_{\mathcal{A}} \times \mathcal{EV} \mapsto \mathcal{P}_{\mathcal{A}}$ that specifies the agent's updated probability measure $[P^{\star}e](.)$ in reaction to the actually received piece of evidence $e \in \mathcal{EV}$ and on the basis of the prior probability measure $P(.)$. This is the formal meaning of updated probabilities which it is important not to confuse with conditional probabilities (cf. below).

We now add additional structure to the notion of evidence. A particular type of evidential input constrains the new probabilities in terms of the probabilities to be assigned to certain propositions. Such evidence can be represented by a constraint family $\langle B_k, p_k \rangle_{k \in K}$ with $B_k \in \mathcal{A}$ and $p_k \in [0,1]$. Such a family $\langle B_k, p_k \rangle_{k \in K}$ is a *Jeffrey constraint* if and only if (i) K is countable, (ii) the propositions B_k of this family are mutually exclusive and jointly exhaustive (i.e., if $\{B_k | k \in K\}$ partitions Ω), and (iii) the values p_k are compatible with the probability calculus (i.e., if $\sum_{k \in K} p_k = 1$).

For Jeffrey constraints, *Generalized Conditionalization* (Jeffrey's Rule) denotes the following equation:

(Gen.Cond) $$[P_i^\star \langle B_k, p_k\rangle_{k\in K}](A) = \sum_{k\in K} p_k \cdot P_i(A|B_k).$$

In the special case of propositional input B with maximum probability (i.e., $B_1 = B$, $p_1 = 1$, $B_2 = -B$, $p_2 = 0$), this equation simplifies to *Conditionalization*:

(Cond) $$[P_i^\star B](A) = P_i(A|B)$$

(Generalized) Conditionalization is equivalent to the combination of the following two conditions. The first condition requires that $\star$ successfully incorporates evidential constraints into updated probabilities:

(Success) $$[P_i^\star \langle B_k, p_k\rangle_{k\in K}](B_k) = p_k$$

The second condition requires that $\star$ keeps the conditional probabilities $P_i(.|B_k)$ constant during the update with $\langle B_k, p_k\rangle_{k\in K}$:

(Rigidity) $$[P_i^\star \langle B_{k'}, p_{k'}\rangle_{k'\in K}](.|B_k) = P_i(.|B_k), \qquad \text{for all } k \in K.$$

The naive view interprets (Generalized) Conditionalization as the recommendation that updated probabilities should always satisfy Success and Rigidity (*'Naive Conditionalization'*). It is this interpretation that I will later criticize. That Success should always hold is more or less taken for granted in the literature, whereas a large number of arguments have been offered in support of Rigidity as a prescriptive ideal of rationality. Among these arguments are the Coherence Arguments (Dutch Book Arguments, No Arbitrage Arguments) to be examined in Section 3 (cf. Teller (1973), Armendt (1980)). Other arguments include axiomatic approaches (Teller (1973), Gärdenfors (1988)), van Fraassen's (1986) symmetry argument, entropy maximization (cf. Williams (1980)), scoring rules (de Finetti (1974)), and Carnap's (1950) frequentist and logical probability arguments.

Jeffrey (1983), (1992) explicitly rejects the idea that (Generalized) Conditionalization is a universally applicable update rule. In his understanding, (Generalized) Conditionalization is less of an update rule than a probabilistic tautology: If the reasoner's updated probabilities fulfill Success and Rigidity, then they are determined by (Generalized) Conditionalization. This statement is a simple theorem of the probability calculus and hence trivially non–controversial. Notice that this interpretation renders (Generalized) Conditionalization compatible with any diachronic evolution of the reasoner's epistemic states. Failure of (Generalized) Conditionalization thus interpreted simply means (and is equivalent to the fact) that the two conditions for its applicability, Success and Rigidity, do not hold.

This consequence of Jeffrey's interpretation underlines some perhaps appealing feature of the naive view. For, if Conditionalization is not universally applicable regardless of the epistemic processes that generate the evidence, it becomes crucial when the reasoner judges a learning situation to qualify for an application of Conditionalization. Unfortunately, standard Bayesianism (without auto–epistemic vocabulary) makes no explicit provisions for representing such judgements, and Jeffrey adds such judgements externally to the standard formalism. In auto–epistemic reasoning, such judgements become an internal part of the Bayesian model itself. The logic of auto–epistemic reasoning that is embodied in Reflection thus provides a logic for how to apply update rules.

Decision Theory

I now introduce the fundamental notions of Bayesian decision theory, both in the language of choice functions and of preference relations. The formulation of the decision–theoretic coherence condition that I will propose in the next section in terms of choice functions paves the way for the proof of the Dutch Book Theorem and its converse. Let $\mathcal{F}$ be the set of all possible options. Let $F \subseteq \mathcal{F}$ be a choice situation, i.e., a set of available options among which the agent has to make a choice. If the agent exhibits regularities in her decision making, these regularities can for each $i \in I$ be encoded in a choice function $C_i : 2^{\mathcal{F}} \to 2^{\mathcal{F}}$ which selects a subset of the options available to the agent ($C_i(F) \subseteq F$, for $F \subseteq \mathcal{F}$).[10] In any given choice situation F, the set $C_i(F)$ represents those available options that the agent's decision criteria cannot eliminate. If $C_i(F)$ contains more than one element, it is left open which option in $C_i(F)$ the agent will actually perform. The sequence $\langle C_i \rangle_{i \in I}$ describes the diachronic changes in the agent's choice function.

Every choice function induces a preference relation $\succeq \subseteq \mathcal{F} \times \mathcal{F}$ over the set of possible options. The following two conditions define the preference relation $\succeq_i$ that is *induced by* a choice function C_i.

(Preference) $\quad f \succeq_i g$ if and only if $f \in C_i(\{f, g\})$.

$f \succeq_i g$ expresses weak preference for f over g. Indifference $f \sim_i g$ is defined as $f \succeq_i g$ and $g \succeq_i f$. Strong preference $f \succeq_i g$ is defined as $f \succeq_i g$ and $g \not\succeq_i f$. A normal choice function in the sense of Sen (1971) is completely determined by the binary choices made between all pairs of options. Any normal choice function can therefore be replaced by the preference relation which it induced.

So far, we have treated options as a primitive notion. Savage's framework adds more detail to the description of options. Savage conceives of options as acts whose outcome depends on external circumstances. The range of possible external circumstances is represented by a non–empty set Ω of possible worlds. Let Γ be a non–empty set of possible consequences. An act is then described as a function $f : \Omega \to \Gamma$ that maps possible worlds into possible consequences. From now on, I will understand acts in this technical sense. Let $u_i : \Gamma \to \Re$ be the agent's utility function at i (for all $i \in I$). Let $\mathcal{F}$ be the set of all functions $f : \Omega \to \Gamma$ such that $u_i \circ f$ is $\mathcal{A}$–measurable for all $i \in I$. For $f \in \mathcal{F}$, let $U_i(f) := E_i(u_i \circ f)$ be the *expected utility of* f. Applied to choices between acts, Bayesian decision theory recommends maximizing the expected utility of the available acts:[11]

(EU–Max) $$C_i(F) = \{f \in F | U_i(f) = \max_{g \in F} U_i(g)\}, \quad \text{for all } F \subseteq \mathcal{F}.$$

Expected–utility maximization induces a normal choice function. Walley (1991) shows how the Coherence Argument can be adapted to the framework of upper and lower probabilities/expectations. The Coherence Argument therefore proves robust against variations of the basic idea that decision under uncertainty should be reached by weighing and comparing alternative options.

Many acts can be described by their payoffs in terms of quantities like money, number of surviving off–spring, energy intake etc. which change depending on external circumstances. The agent's utility function assigns a certain value to different levels of these primary quantities. In many applications, the agent's utility is non–linear and, especially, non–additive in these quantities[12] because they take into account the global features of the conditions in which these primary payoffs are realized. The utility of combining two acts f and g can therefore differ from the sum of the utilities of f and g in isolation. This is not to say, of course, that they can never be additive. Recreational monetary gambles with small stakes are perhaps an example of near additivity.

I will present the Coherence Argument in a form that does not presuppose additive utilities simply because it does not involve combining two acts into one. This operation is only required in the game–theoretic illustration of the decision–theoretic argument. In order to blame the game–theoretic predicament on deviant probabilities, the utilities of the considered acts would have to be additive (and remain constant diachronically.)[13]

2. The Coherence Argument for Reflection

Diachronic Coherence

In defending Reflection, I will draw on a diachronic version of Savage's (1954) Sure Thing Principle and show its equivalence with Reflection. Imagine an option f which you are sure you will under any circumstances strictly prefer to g tomorrow. It would then not be prudent of you today to choose g rather than f, since you are certain that tomorrow you will want to undo g and bring about f. In variation of Savage's example,[14] you should buy a property today if you believe that you will still want to buy it tomorrow regardless of what happens in the meantime (whether a Republican or a Democrat is elected president). In short,

(Diachronic Coherence) Prefer weakly (strongly) today what you are certain you will prefer weakly (strongly) in the future.

This principle presupposes that the options f and g are described in a time invariant way. They must not produce different effects if chosen at different times. When necessary their description, conditions, and consequences must be time–indexed in order to avoid ambiguities. Part of the same requirement is that their utilities must remain constant over time. In our example, if you expect the price of the property to fall, you should certainly wait.

It is important to realize that this coherence condition is formulated from the point of view of the reasoner herself. It therefore involves her auto–epistemic opinions about her future preferences. In the literature, coherence conditions are normally formulated from the point of view of an external observer (with omission of the phrase 'the reasoner is certain that'). The auto–epistemic treatment of the coherence condition could, however, be harmonized with the observer's point of view. We could easily shift back into talk from the point of view of an external observer and interpret the reasoner's auto–epistemic views about her own epistemic states as the observer's views about the reasoner's epistemic states. We could also think of the reasoner's auto–epistemic views as an echo of the observer's views from whom she could have borrowed her views about her epistemic situation.

Notice also that Diachronic Coherence does not impose any restrictions on single future choice criteria. It rather restricts what future choice criteria the reasoner can coherently judge admissible, given her present choice criterion. Diachronic Coherence does therefore not require an agent obstinately to abide by her old decisions no matter what. She is fully entitled to update her old decisions, for example, when she receives new evidence about the probable effect of her options. A violation of Diachronic Coherence requires much more than revising prior decisions by updating. Diachronic Coherence is only violated if the reasoner believes with full certainty that her future choice criteria will under any circumstances overthrow her present choices. It is the anticipated systematic posterior disregard for prior preferences that constitutes a diachronic incoherence.

Insofar as Diachronic Coherence regulates the anticipated diachronic evolution of preferences, it is a diachronic principle. Insofar as this criterion of coherence takes the point of view of the reasoner's current prior beliefs about her future preferences, Diachronic Coherence could also be regarded as a synchronic principle. In an equivalent formulation that stresses its relativity to the prior point of view, Diachronic Coherence outlaws any prior epistemic state in which future preferences are expected systematically to disregard prior preferences. (The same holds *mutatis mutandis* for Reflection.)

Decision–theoretic incoherence may be dramatized in a game–theoretic setup with two players since a violation of decision–theoretic Diachronic Coherence can under suitable conditions give rise to game–theoretic incoherence. Player 1 is *game–theoretically incoherent* if and only if a second player with the same information as Player 1 can follow a strategy that necessarily exploits Player 1. In other words, Player 2 can make money off Player 1 simply because of the way Player 1 plays the game and regardless of any contingent external conditions and without additional information. Such a strategy for Player 2 is called a 'Dutch book'. In short, Player 1 is diachronically incoherent if and only if a Dutch book can be made against him. It is not difficult to see how a decision–theoretic incoherence can lead Player 1 to play the game in such a way that he also becomes game–theoretically incoherent. Imagine that Player 2 can impose a choice between f and g on Player 1. Suppose that tomorrow Player 1 will necessarily choose f, but today he chooses g. Then Player 2 has successfully lured Player 1 into a sure loss because tomorrow he will be able to charge Player 1 for helping him to undo f. Notice that there is no risk for Player 2 involved. He knows that he will be able to charge Player 1 tomorrow because Player 1 will always prefer g. In terms of bets, Player 2 could sell Player 1 a bet today and buy it back tomorrow at a lower price no matter what happens in the meantime. Buying and selling would cancel each other out and Player 2 could keep the price difference — presupposing, of course, that it is always possible to find bets with additive utilities. Recreational monetary bets are the most plausible candidate as long as their stakes are sufficiently small.

In the literature, the Coherence Argument and the notion of diachronic coherence have only been discussed in a game–theoretic setting. This has earned the Coherence Argument some unfair criticism that is easily avoided by a decision–theoretic formulation of Diachronic Coherence. Christensen (1991), for example, argues that the game–theoretic Coherence Argument relies on the existence of an 'Evil Super–Bookie' (Player 2). The idea is that only the threat posed by Player 2 can rationally prevent Player 1 from having incoherent preferences. If Player 2 is not on guard, Player 1 may as well be incoherent. Howson (1995), however, worries that the Player 2 might not be told how Player 1 is going to play the game. In the betting game, Player 2 might not know Player 1's betting odds. The fact that Player 2 could possibly make a Dutch book against Player 2 even provides a reason for Player 1 to conceal his betting odds (i.e., how he will play the game). Furthermore, it is often remarked that the Coherence Argument depends on Player 1's inclination to play the game, e.g. to bet. No game, no Dutch book. Finally, it is argued that the Coherence Argument only works for preferences over options whose payoff conditions can in a finite amount of time be decided to be true or false. If their truth could not be determined at all or if it took infinitely long to do so, incoherence involving these options would not affect the payoff structure of any real game and would hence be innocuous.

These critical remarks only apply if the Coherence Argument is taken at its game–theoretic face value. The essential problem behind game–theoretic incoherence is that it arises from the intrinsic features of the agent's choice function. These intrinsic problems vex the agent even if they are not brought into public by game–theoretic interactions with other players. It is therefore better to understand Diachronic Coherence as a condition on personal preferences rather than on how to play a game. In the decision–theoretic setting, all the above objections no longer work because they all construct situations where the incoherence game would not be played.

The Equivalence Theorem

The shift to a decision–theoretic formulation of Diachronic Coherence has the advantage of bringing out the more fundamental incoherence that is only illustrated in the game–theoretic formulation. When I prove the equivalence of Diachronic Coherence and Reflection, it will become apparent that an additional advantage of the decision–theoretic formulation lies in the technical ease with which this proof can be carried out. Similar proofs in the game–theoretic setup are much more complicated. In preparation of this proof, I now present a formal version of Diachronic Coherence. For this purpose, we need to define the set $\mathcal{C}(i,j)$ of choice criteria at time j that are possible from the point of view of stage i. I assume that the agent has probabilities for the future evolution of his choice criterion. Let $P_i(\mathrm{CF}_j(C'_j))$ be the probability at stage i for the event that the agent's choice function evolves into C'_j at stage j. We define the set of choice functions at j that are subjectively possible from the point of view of stage i as follows:

(Def.$\mathcal{C}$) $$\mathcal{C}(i,j) := \{C'_j | C_i = C'_i, P_i(\mathrm{CF}_j(C'_j)) > 0\}$$

In terms of choice functions, the coherence condition then becomes:[15]

(Diachronic Coherence) $\forall i, j \ (i \leq j)$:

$$\bigcap_{C'_j \in \mathcal{C}(i,j)} C'_j(\{f,g\}) \ \subseteq \ C_i(\{f,g\}) \ \subseteq \ \bigcup_{C'_j \in \mathcal{C}(i,j)} C'_j(\{f,g\})$$

The restriction of Diachronic Coherence to binary choices $\{f,g\}$ is essential. Imagine today's preference ranking: $f \succ g \sim h$. The binary coherence condition is satisfied if we have the following two subjectively possible rankings in the future: $h \succ f \succ g$ and $g \succ f \succ h$. For example, g and h could be two fair bets (hence, $g, h \sim 0$), one on A and one against it (hence, $g = -h$), and f could be an advantageous bet on B (hence, $f \succ 0$). If we will learn tomorrow whether or not A is the case, we will tomorrow either prefer g over h ($g \succ h$) or h over g ($h \succ g$). But although f is today better than either of the fair bets g and h ($f \succ g, h$), it may tomorrow be worse than the winning bet and better than the loosing bet ($g \succ f \succ h$, if A turns out to be true and g is the winning bet). Diachronic Coherence can therefore not be extended to choice sets of a size > 2 or else the above rankings would incorrectly be ruled out as incoherent.

The following two subsections will be devoted to the proof of the central theorem which relates Diachronic Coherence to Reflection:

Theorem 2.1 (Equivalence Theorem). Given AE–Transparency, Diachronic Coherence is equivalent to Reflection for expected–utility maximizers with constant utilities.

Diachronic incoherences can stem from two sources, either from ill–balanced diachronic probabilities or from a diachronic change in utilities. I already pointed out that the principle of Diachronic coherence presupposes the diachronic constancy of utilities. This is the reason why Theorem 2.1 has to make the same assumption. I want to emphasize, however, that the justification of Reflection does not depend on the assumption that the utilities of *all* options remain constant. The theorem will yield a justification of Reflection whenever we can find *some* options whose utilities remain constant. If we want to have coherent preferences over such options, the theorem tells us that it is necessary and sufficient to obey Reflection. Now suppose we obey Reflection. Then we might nonetheless be incoherent over options with diachronically changing utilities. What is more, Reflection prevents us from adjusting our probabilities in a way that would avoid these incoherences over options with non–constant utilities because these adjustments would induce incoherences over options with constant utilities. In order to complete our justification of Reflection we must thus accept the additional thesis that incoherences induced by non–constant utilities are irrelevant. Another way of expressing this additional thesis is to say that changes in values (utilities) amount to adopting an altogether different choice criterion. Under this perspective, incoherences induced by non–constant utilities are irrelevant to the coherence of the present choice criterion because all futures of the present choice criterion must have the same constant utilities *by virtue of being* a future of the present choice criterion. Notice that this is not an argument for the irrelevance of incoherences induced by non–constant utilities, but just a reformulation of the same thesis.

For the proof of the Equivalence Theorem, we need one final assumption that allows us to construct acts with arbitrary utility payoffs:

(Fullness) For any $r \in \Re$, any u_i $(i \in I)$, and any $\omega \in \Omega$, there is a $f \in \mathcal{F}$ such that $u_i(f(\omega)) = r$.

This is an assumption very much in the style of the decision–theoretic literature.

The Dutch Book Theorem

The Dutch Book Theorem for Reflection was first presented by van Fraassen (1984) in the game–theoretic context. Goldstein (1983) had reached a similar result in terms of scoring rules and for only finitely additive probability measures. The Coherence Argument for Reflection exploits the fact that information about the reasoner's opinion can be expressed in the auto–epistemic vocabulary of $\mathcal{A}$. The agent can consider acts that depend on auto–epistemic facts, such as bets on the reasoner's opinion.[16] Let $\beta_0 = 0$ be the empty act which has zero payoffs under all circumstances. If Reflection does not hold between two stages i and j, we can construct an act β_1 with zero expectation at i and strictly negative expectation at j. Thus, the agent will be indifferent between β_0 and β_1 at i $(C_i(\{0, \beta_1\}) = \{0, \beta_1\})$, but will strictly prefer β_0 over β_1 at

j ($C'_j(\{0, \beta_1\}) = \{0\}$, for all subjectively possible $C'_j \in \mathcal{C}(i, j)$). Failure of Reflection therefore gives rise to a violation of Diachronic Coherence. By contraposition, Diachronic Coherence implies Reflection.

Theorem 2.2 (Dutch Book Theorem). For expected–utility maximizers with constant utilities, AE–Transparency and Diachronic Coherence imply Reflection.

Proof: Throughout the argument it is assumed that the (synchronic) probability axioms hold for personal probabilities (i.e., that personal probability assignments are, mathematically speaking, probability functions). Suppose that the probability measures at i and j ($i \leq j$) do not obey Reflection for some proposition A. Without loss of generality, we assume that $P_i(A|\mathrm{SP}_j(Q_j)) = Q_j(A) + \delta$ with $\delta > 0$ and $P_i(\mathrm{SP}_j(Q_j)) > 0$.[17] We will now construct an act β_1 with payoffs in units of utility that will lead to a diachronic incoherence. For notational and intuitive convenience, we can break the definition of β_1 down into two subclauses with $u(\beta_1) := u(\beta_{1.1}) + u(\beta_{1.2})$:[18]

$$u(\beta_{1.1}) = \left\{ \begin{array}{rl} 1 - P_i(A|\mathrm{SP}_j(Q_j)), & \text{if } A \\ \\ -P_i(A|\mathrm{SP}_j(Q_j)), & \text{otherwise} \end{array} \right| \mathrm{SP}_j(Q_j)$$

$$u(\beta_{1.2}) = \left\{ \begin{array}{rl} \delta(1 - P_i(\mathrm{SP}_j(Q_j))), & \text{if } \mathrm{SP}_j(Q_j) \\ \\ -\delta P_i(\mathrm{SP}_j(Q_j)), & \text{otherwise} \end{array} \right.$$

By Fullness, there is a $\beta_1 \in \mathcal{F}$ that satisfies these conditions. Since the payoffs of $\beta_{1.1}$ and $\beta_{1.2}$ are directly given in units of utility, this procedure does not presuppose that utilities are additive in, for example, monetary payoffs. The act β_1 is simply defined to have payoffs in utility that are equal to the sum of utility payoffs in $\beta_{1.1}$ and $\beta_{1.2}$. Alternatively, we could define β_1 in one big swoop. If utilities are non–additive in monetary payoffs, β_1 can therefore not be interpreted as the combination of two acts $\beta_{1.1}$ and $\beta_{1.2}$ with payoffs in monetary units.

Now assume that utilities remain constant over time:

(Constant Utilities) $\quad u_i = u_j = u$

$\beta_{1.1}$ is an act conditional on $\mathrm{SP}_j(Q_j)$. If this condition fails, its payoff is zero units of utility. Clearly, $E_i(u(\beta_1)) = 0$. When calculating the future expectations of β_1, we have to distinguish two cases:

Case 1: The personal probability function at j is identical to $Q_j(.)$. If this function is auto– epistemically transparent, we have $E_{Q_j}(u(\beta_1)) = (1 - P_i(A|\mathrm{SP}_j(Q_j))) \cdot Q_j(A) - P_i(A|\mathrm{SP}_j(Q_j))) \cdot (1 -$

$Q_j(A))+\delta\cdot(1-P_i(\mathrm{SP}_j(Q_j))$. Hence, $E_{Q_j}(u(\beta_1)) = -\delta\cdot P_i(\mathrm{SP}_j(Q_j)) < 0$. (It is here that we need the assumption that $P_i(\mathrm{SP}_j(Q_j)) > 0$ which restricts the proof to 'Non–Zero' Reflection.)

Case 2: The personal probability function at j is not identical to $Q_j(.)$. AE–Transparency implies that $E_{Q_j}(u(\beta_1)) = E_{Q_j}(u(\beta_{1.2})) = -\delta P_i(\mathrm{SP}_j(Q_j)) < 0$. An expected–utility maximizer would thus accept β_1 at i (prefer it to the empty act), although he would be certain that he will reject it at j (prefer empty act to β_1). Failure of Reflection therefore violates Diachronic Coherence. □

Converse Dutch Book Theorem

In order to prove the Converse Dutch Book Theorem for Reflection, I will first show that Diachronic Coherence is equivalent to the Generalized Iteration principle. It will then be easy to see that Reflection implies Generalized Iteration and hence Diachronic Coherence. In the decision–theoretic context, the proof of the Converse Dutch Book Theorem is particularly elegant because it does not involve the game–theoretic notion of a strategy. Converse Dutch Book Theorems in the game–theoretic setup have been studied by Skyrms (1987c). Their proof becomes very complicated because it has to consider all possible strategies against Player 1.

Generalized Iteration requires that present expectations should lie within the subjectively possible range of future expectations (cf. van Fraassen (1995), Hild (1998)). We define the subjectively possible range of j–opinions as $\mathcal{P}(i,j) := \{Q_j | P_i = Q_i, P_i(\mathrm{SP}_j(Q_j)) > 0\}$.[19]

(Gen.Iteration) $\forall i, j \ (i \leq j), \quad \forall X \ (X : \Omega \to \Re)$:

$$\min_{Q_j \in \mathcal{P}(i,j)} E_{Q_j}(X) \leq E_{P_i}(X) \leq \max_{Q_j \in \mathcal{P}(i,j)} E_{Q_j}(X)$$

Lemma 2.3. For expected–utility maximizers, Generalized Iteration and Diachronic Coherence are equivalent.

Proof: (i) Assume that Generalized Iteration holds. Hence, if $E_{P_i}(X) \geq 0$, there is a $Q_j \in \mathcal{P}(i,j)$ such that $E_{Q_j}(X) \geq 0$.
Right Half of Diachronic Coherence: Suppose that $E_{P_i}(u(f)) \geq E_{P_i}(u(g))$ and, hence, $E_{P_i}(u(f) - u(g)) \geq 0$. (Expectations are by definition additive in utilities, even if utilities are not additive in monetary payoffs.) Thus, there is a $Q_j \in \mathcal{P}(i,j)$ such that $E_{Q_j}(u(f) - u(g)) \geq 0$ and therefore $E_{Q_j}(u(f)) \geq E_{Q_j}(u(g))$. Hence, there is a subjectively possible choice function $C'_j \in \mathcal{C}(i,j)$ such that $f \in C'_j(\{f,g\})$.

Left Half of Diachronic Coherence: Suppose that for all E_{Q_j}: $E_{Q_j}(u(f)) \geq E_{Q_j}(u(g))$. Then, if $E_{P_i}(u(g)) > E_{P_i}(u(f))$, there would have to be a E_{Q_j} such that $E_{Q_j}(u(g)) > E_{Q_j}(u(f))$. By *reductio ad absurdum*, it follows that $E_{P_i}(u(f)) \geq E_{P_i}(u(g))$. Hence, $f \in C_i(\{f, g\})$ if $f \in C'_j(\{f, g\})$ for all subjectively possible choice functions C'_j.[20]

(ii) Assume that Diachronic Coherence holds. Let $f, g \in \mathcal{F}$ be acts such that $u(f) = X$ and $u(g) = E_{P_i}(X)$. By Fullness, such f and g exist. We have $C_{P_i}(\{f, g\}) = \{f, g\}$. Suppose that for all $Q_j \in \mathcal{P}(i, j)$: $E_{P_i}(X) > E_{Q_j}(X)$. Then for all choice functions $C_{E_{Q_j}} \in \mathcal{C}(i, j)$ based on any such E_{Q_j}, $f \notin C_{E_{Q_j}}(\{f, g\})$. Hence, $f \notin \bigcup C_{E_{Q_j}}(\{f, g\})$. Contradiction! If, on the other hand, $E_{Q_j}(X) > E_{P_i}(X)$ for all E_{Q_j}, then $g \notin C_{E_{Q_j}}(\{f, g\})$ for all choice functions $C_{E_{Q_j}}$. Hence, $g \notin \bigcup C_{E_{Q_j}}(\{f, g\})$. Contradiction! □

Lemma 2.4. Reflection implies Generalized Iteration.

Proof: Due to the Total Probability Theorem for σ–additive probabilities, Reflection implies

$$P_i(A) = \sum_{Q_j \in \mathcal{P}(i,j)} Q_j(A) \cdot P_i(\mathrm{SP}_j(Q_j))$$

This means that $P_i(.)$ is a convex combination of the elements in $\mathcal{P}(i, j)$ which implies Generalized Reflection. □

Theorem 2.5 (Converse Dutch Book Theorem). For expected–utility maximizers with constant utilities, Reflection implies Diachronic Coherence.

Relative subjective rationality

Diachronic Coherence incorporates diachronic aspects into the notion of decision–theoretic rationality by imposing constraints on the rational diachronic evolution of choice criteria. It postulates that the evolution of a choice criteria cannot be rational if it necessarily results in a diachronic conflict between synchronic choices. So far, we have discussed a framework whose only standard of rationality is (synchronic and diachronic) coherence. Within Bayesianism, there is potentially space for a notion of rationality that extends the standard core of coherence. A flagrant example of irrationality even by Bayesian lights are beliefs that result from the consumption of mind and perception altering drugs. It seems to be an implicit assumption in much of the Bayesian literature that such cases should be excluded and that the discussion should be limited to agents who are ideally rational by some vague external standard. I will now briefly consider Diachronic Coherence and Reflection in a framework that allows the reasoner to form opinions about which future opinions seem to be rational or irrational from her present point of view.

If you think that tomorrow you will irrationally come to believe that pigs can fly, you should certainly not adopt this belief today just because you want to satisfy Diachronic Coherence and Reflection. Rather, you should prudently limit the effects of your future irrationality, stick to your present belief that pigs cannot fly, and violate Diachronic Coherence. (Of course, it would be even better if you could prevent yourself from adopting this irrational belief tomorrow.) In a liberalized epistemological setting which allows the possibility of irrationality, Reflection and Diachronic Coherence can at times not be realized. Their violation then indicates the presence of a (yet to be localized) fault in rationality which makes it impossible to establish Diachronic Coherence.[21]

Diagnosing a violation of Diachronic Coherence is only a first step that leaves open at exactly what diachronic state this fault has occurred. For all we know, the diachronic incoherence may result from the prior epistemic state as much as from the posterior epistemic state. In a second step, the attribution of this incoherence to a specific stage in the diachronic sequence of epistemic states is a matter of relative judgment. From the prior point of view, a posterior epistemic state may seem irrational and unendorsable, but the prior epistemic state may create the same impression from the point of view of this posterior epistemic state. This leads to a notion of relative subjective rationality, or of relative endorsement.

If the reasoner herself is given the means of expressing failures of relative rationality, a refined version of Reflection can again be defended by a Coherence Argument in the liberalized setting. Let '$\text{Rational}_i(Q_j, j)$' express that, by the standards of reasoner's opinion at i, it would be rational to adopt Q_j as the opinion at j.[22] In other words, Q_j at j is endorsed by the opinion at i. For the time being, simply consider the special case of a self–devaluing reasoner who does not believe that his current opinion is rational ($P_i(-\text{Rational}_i(P_i, i)) = 0$). When I am under the influence of alcohol, for instance, I might doubt whether I am still rational. I would then abstain from certain actions and decisions. I would, for example, not drive a car. It is only when I believe that I am still rational that the trouble starts. Relative subjective rationality therefore has an immediate impact on decision making.

In particular, the agent must not believe that all her opinions are rational if they lead her into diachronic incoherence. If the agent runs into diachronic incoherence, she must step back and tell herself 'Wait a minute, there must be something wrong here'. There must be at least one stage where she believes that there is a fault in her synchronic decisions somewhere. Equivalently, her decisions must not be diachronically incoherent within an interval I if her decisions during that interval are all subjectively rational from the point of view of each other.

I will now briefly state the principles of Refined Reflection and Refined Coherence in an (half–hearted) attempt to make the idea of relative subjective rationality more precise. Define $\mathcal{RC}(i,j) := \{C'_j | C_i = C'_j, P_i(\mathrm{CF}_j(C'_j) \cap \mathrm{Rational}_i(C'_j, j)) > 0\}$ as the set of posterior choice functions that are potentially rational from the prior point of view. The only difference between the previous, crude, coherence condition and Refined Coherence is that we replace the set $\mathcal{C}(i,j)$ of subjectively possible posterior choice functions by the set $\mathcal{RC}(i,j)$ of posterior choice functions that are potentially rational from the prior point of view. This Refined Coherence condition supports the principle of Refined Reflection proposed by Jeffrey (1988):

(Ref.Reflection) $\forall i, j (i \leq j), \quad P_i = Q_i$
$P_i(.|\mathrm{SP}_j(Q_j) \cap \mathrm{Rational}_i(Q_j, j)) = Q_j(.)$

A modified Dutch book can be constructed simply by betting on $\mathrm{SP}_j(Q_j) \cap \mathrm{Rational}_i(Q_j, j)$ instead of $\mathrm{SP}_j(Q_j)$ (cf. below).

3. Conditionalization

The supposed theorems

Historically, (naive) Conditionalization was the first principle to be defended by a Diachronic Coherence Argument.[23] Hacking (1967) pointed out that the synchronic argument in support of the axiom (CP) for conditional probabilities does not justify the diachronic update rule of Conditionalization. Nonetheless, David Lewis proposed a Dutch Book Theorem for Conditionalization (reported in Teller (1973)) in the game–theoretic setting which uses the presence of an update rule to exploit the potential imbalance between present and future probabilities.

Supposed Dutch Book Theorem: Assume that a change in personal probabilities from $P_i(.)$ to $P_{i+1}(.)$ occurs due to an update rule $\star$ under evidential input B with maximum probability ($P_{i+1}(.) = [P_i^\star B](.)$) where $P_i(B) > 0$ and that this update rule differs from Conditionalization on B. Without loss of generality, we can assume that

$P_i(A|B) = P_{i+1}(A) + \delta$ with $\delta > 0$. Consider the following acts:

$$u(\beta_{1.1}) = \left\{ \begin{array}{ll} 1 - P_i(A|B), & \text{if } A \\ \\ -P_i(A|B), & \text{if } -A \end{array} \right| B$$

$$u(\beta_{1.2}) = \begin{cases} \delta(1 - P_i(B)), & \text{if } B \\ \\ -\delta P_i(B), & \text{otherwise} \end{cases}$$

$$u(\beta_2) = \begin{cases} 1 - P_{i+1}(-A), & \text{if } -A \\ \\ -P_{i+1}(-A), & \text{otherwise} \end{cases}$$

Again, we define $u(\beta_1) := u(\beta_{1.1}) + u(\beta_{1.2})$. The agent expects zero payoff from β_1 at i and thus accepts it at i. If she learns that B, Player 2 will offer her β_2 at $i + 1$ which she will accept. If she learns that $-B$, Player 2 offers no further bets. Player 2's strategy is supposed to lead to a game that necessarily terminates with a the negative payoff $-\delta P_i(B)$ for Player 1.[24]

Supposed Converse Dutch Book Theorem: If there were a dynamic Dutch book against a Conditionalizer, then there would also be a synchronic Dutch book against her because the updated probabilities at j equal conditional probabilities at i. For a Conditionalizer, diachronic coherence hence reduces to synchronic coherence (Skyrms (1987c), 15).

The fallacy

These supposed theorems make much stronger assumptions than they explicitly reveal. This has been argued before, e.g., by Shafer (1985). The proofs of the supposed theorems implicitly utilize information about the particular type of learning situation in which Conditionalization is to be applied and about the epistemic processes that generate the reasoner's evidence. Since these non–trivial assumptions do not hold in general, the corresponding Coherence Argument is not valid — at least if it is taken to support a universally applicable update rule.

To begin with, the Dutch Book Theorem for Conditionalization assumes that the reasoner[25] makes certain assumptions about the timing and the reliability of her evidence.[26] These assumptions are quite plausible for the type of bets that are usually considered in the literature. Paradigmatically, let B be the proposition that Trafalgar will win the horse-race on Saturday, and A the proposition that Trafalgar will also win the horse-race on Sunday.

More importantly, we find the much more problematic *Partitioning Assumption* at the core of the argument:

(PA) The reasoner is certain that she receives at $i+1$ either the total evidence B or the total evidence $-B$. ($P_i(\mathrm{Ev}^{tot}_{i+1}(B) = -\mathrm{Ev}^{tot}_{i+1}(-B)) = 1$.)[27]

Suppose, in the given example, that in addition to the proposition $B =$ 'Trafalgar wins on Saturday', Player 1 also learnt the proposition C at $i+1$ that Nelson, Trafalgar's strongest competitor, was injured in the race on Saturday. Presumably, this new information leads to a posterior probability $[P^{\star}_i(B \cap C)](A)$ for A that is higher than the reasoner's probability conditional on B (i.e., $P_i(A|B)$) and — in case this probability is different — also higher than her probability updated on B (i.e., $P_{i+1}(A) = [P^{\star}_i B](A)$). Rationality at $i+1$ would then require Player 1 to decline the bet β_2 against Trafalgar. Hence, the reasoner will not accept Player 2's Dutch book. She will even coherently expect to make a profit from β_1 because she now considers it more likely that Trafalgar will win on Sunday than she had anticipated when she accepted β_1 at i.

We can easily restate this observation in terms of decision–theoretic coherence. Obviously, $E_i(u(\beta_1)) = 0$. If B is received as total evidence, β_1 has a negative expected utility at $i+1$. The same is true if the reasoner receives $-B$ as total evidence.[28] But unless we exclude the (subjective) possibility of any other evidence different from B and $-B$, Diachronic Coherence will not be violated. If, for instance, the total evidence is $B \cap C$, then the expected utility of β_1 at $i+1$ may well be non–negative. The agent would then remain diachronically coherent. Similar considerations show that the Converse Dutch Book Theorem relies on the presence of the Partitioning Assumption (PA).[29]

This example illustrates that future evidence is not always packed into neat parcels that partition Ω. In other words, 'Either total evidence B or total evidence $-B$' is not a complete alternative for future total evidence ($\mathrm{Ev}^{tot}_j(B) \cup \mathrm{Ev}^{tot}_j(-B) \neq \Omega$) — as the Supposed Dutch Book Theorem incorrectly assumes. The agent may well experience a learning situation in which she will learn either (i) the proposition B or (ii) the proposition $-B$ or (iii) the proposition $B \cap C$. All we can therefore presuppose is the logically exhaustive alternative 'Either B is total evidence or not' ($\mathrm{Ev}^{tot}_j(B) \cup -\mathrm{Ev}^{tot}_j(B) = \Omega$).

If the Coherence Argument for (Generalized) Conditionalization fails, why should the corresponding argument for Reflection be immune from the same criticism? For the very simple reason that the Coherence Argument for Reflection assumes only that either Q is the future probability measure or not ($\mathrm{SP}_j(Q)$ or $-\mathrm{SP}_j(Q)$) — *tertium non datur.*

This criticism of the Coherence Argument for Conditionalization differs from Jeffrey's (1988) objection that this argument presupposes Rigidity. If that were true, the argument would be blatantly circular because Rigidity is precisely the non–logical element that constitutes the classical update rule (in the presence of Success). The objection against (PA) does not yet show that Rigidity (and therefore Conditionalization) fails. For all I have argued so far, it might very well hold in the horse race example. My objection against (PA) rather shows that the argument for Conditionalization is *not valid*. Objecting to the implicit assumption (PA) of the Coherence Argument does not suffice to show that its conclusion is false. This is what I will argue in the following final section.

4. Auto–Epistemic Conditionalization

The grain of truth in the proposed Dutch book for Conditionalization emerges from a possible reaction to the criticism of the partitioning assumption (PA): Simply bet on the proposition $\mathrm{Ev}_{i+1}^{tot}(B)$ that the reasoner receives B as total evidence instead of B and make use of the logically exhaustive disjunction 'Either B is total evidence or not' ($\mathrm{Ev}_{i+1}^{tot}(B) \cup -\mathrm{Ev}_{i+1}^{tot}(B) = \Omega$). We then obtain a Dutch Book Theorem for the more sophisticated principle that I call *'Auto–Epistemic Conditionalization'*. It requires that the effect of updating on evidence $e \in \mathcal{EV}$ is encoded in the present probabilities conditional on the proposition that $e \in \mathcal{EV}$ will be total evidence:

(AE–Cond) $$[P_i^\star e](.) = P_i(.|\mathrm{Ev}_j^{tot}(e))$$

In those special cases where future opinions are entirely evidence driven (i.e., where for all possible futures Q_{i+1} of P_i take the form $P_i^\star e$, for some piece of evidence e), AE–Conditionalization is just another way of writing Reflection.[30]

Despite their superficial resemblance, naive Conditionalization and AE–Conditionalization are two fundamentally different principles. Auto–epistemic reasoners can form beliefs about the processes that generate their evidence B and about the circumstances under which they receive which piece of total evidence B. The conditional probabilities $P_i(B|\mathrm{Ev}_{i+1}^{tot}(B))$, $P_i(\mathrm{Ev}_{i+1}^{tot}(B)|B)$ represent the reasoner's view about this relationship. AE–Conditionalization enjoins the reasoner to draw on this prior assessment of her learning situation when choosing an update rule. Although this prescription is very much in line with the general Bayesian outlook, it is all too often neglected in the naive account of Bayesian learning.

The naive view interprets Conditionalization as a universal update rule that is insensitive to the epistemic process generating the actually obtained evidence. From the posterior point of view, there is no mention of what other propositions the reasoner would have learnt under different circumstances. The conditional probabilities $P_i(B|\mathrm{Ev}^{tot}_{i+1}(B))$, $P_i(\mathrm{Ev}^{tot}_{i+1}(B)|B)$ are in general hard to come by, at least if the process that determines the evidence lies beyond the reasoner's control. I therefore find the naivety of Conditionalization — and indeed the naivety of any universal update rule — quite appealing. From the more sophisticated auto–epistemic viewpoint, this naivety pretends that Conditionalization holds for any prior assessment $P_i(B|\mathrm{Ev}^{tot}_{i+1}(B))$, $P_i(\mathrm{Ev}^{tot}_{i+1}(B)|B)$ of the reasoner's evidence (apart, perhaps, from the Reliability constraint $P_i(B|\mathrm{Ev}^{tot}_{i+1}(B)) = 1$).

It is therefore not surprising that we find clashes between the ambitions of universal update methodologies like Conditionalization and the verdict of the auto–epistemic criteria for their admissibility. What particular prior assessment of a learning situation validates the application of Conditionalization? The following simple theorem gives two necessary and jointly sufficient conditions (cf. Skyrms (1980)).[31]

Theorem 4.6. Let $e_{GC} = \langle B_k, p_k \rangle_{k \in K}$ be a Jeffrey constraint. Given AE–Conditionalization, Success is equivalent to belief in the reliability of evidential input:

(Reliability) $$P_i(B_{k'}|\mathrm{Ev}^{tot}_j(e_{GC})) = p_{k'}$$

Given AE–Conditionalization, Rigidity is equivalent to the probabilistic independence of A and $\mathrm{Ev}^{tot}_j(e_{GC})$ given B:

(Evidential Indep.) $$P_i(A|B_{k'} \cap \mathrm{Ev}^{tot}_j(e_{GC})) = P_i(A|B_{k'})$$

Given AE–Conditionalization, Conditionalization thus holds if and only if Reliability and Evidential Independence hold.

Reliability and Evidential Independence are conditions on the reasoner's prior assessment of her learning situation. Failures of Reliability and Evidential Independence describe updates which go against Success and Rigidity and therefore against (Generalized) Conditionalization.[32] In the case of simple Conditionalization, these two conditions are met, for instance, if the reasoner judges her evidence reliable and makes the Partitioning Assumption.[33]

Neither Reliability nor Evidential Independence hold universally true. It is easy enough to imagine how Reliability may fail. I will now briefly present the example of Freund's (1965) Puzzle in which Evidential Independence fails although Reliability holds.[34] In order to describe a learning situation unambiguously, we must specify (i) the representation $\mathcal{EV}$ of evidence and (ii) the epistemic process that generates the reasoner's evidence. In my opinion, the confusion surrounding Freund's Puzzle and similar scenarios stems from the fact that Conditionalization as a naive update rule merely talks about the actually received evidence but neglects to speak of the reasoner's subjective assessment of her overall learning situation. I assume that evidence is presented to the reasoner in form of utterances made to her.

Example 4.7. In Freund's Puzzle, two cards are drawn from a four–card deck with the aces and deuces of spades and hearts. First the reasoner is told that at least one ace has been drawn. At stage 1 she assigns an equal probability of $\frac{1}{5}$ to the remaining possible outcomes. Then it is also revealed to her that the ace of spades has been drawn. She trusts fully in what she is told (Reliability). She also believes fully that if only one ace has been drawn, then she will be told the colour of that ace. Supposing that two aces have been drawn, she believes with probability p that she will be told that the ace of spades has been drawn ($P_1(\mathrm{Ev}_2^{tot}(\spadesuit A)|\spadesuit A \cap \heartsuit A) = p$).
Then her updated probability at stage 2 for both aces having been drawn is $\frac{p}{p+2}$. Only if $p = 1$, will the event $\spadesuit A \cap \heartsuit A$ be probabilistically independent from $\mathrm{Ev}_2^{tot}(\spadesuit A)$ given $\spadesuit A$. In other words, Evidential Independence fails unless $p = 1$. The required update therefore takes the form of Conditionalization on $\spadesuit A$ only in the extreme case of $p = 1$. A more detailed analysis can be found in Hild (1998).

In an attempt to save naive Conditionalization from such counter–examples, we could adjust our notion of total evidence and require that B should only count as total evidence* if B is probabilistically equivalent to $\mathrm{Ev}_{i+1}^{tot}(B)$ (i.e., only if $P_i(B = \mathrm{Ev}_{i+1}^{tot}(B)) = 1$).[35] This adjusted notion of evidence* would allow us simply to re-interpret Conditionalization as AE–Conditionalization. Put differently, the reasoner should update* on the evidence* 'My evidence is B'/'I learn that B'. In learning from B, the reasoner would hence be expected to draw on her theory $P_i(B|\mathrm{Ev}_{i+1}^{tot}(B))$, $P_i(\mathrm{Ev}_{i+1}^{tot}(B)|B)$ about her overall learning situation.

Apart from verbal disagreements, the central issue in distinguishing naive Conditionalization and AE–Conditionalization is the following: From the sophisticated auto–epistemic viewpoint, the reasoner needs a prior assessment for how she epistemically relates to her environment before she can even start to update her prior probabilities with the input that she receives. The needed prior assessment $P_i(B|\mathrm{Ev}_{i+1}^{tot}(B))$, $P_i(\mathrm{Ev}_{i+1}^{tot}(B)|B)$ is in general difficult to come by. The statistical literature usually bypasses this difficulty and assumes that the reasoner knows that the experimental setup satisfies Reliability and the Partition Assumption.[36] This difficulty becomes salient, however, if Nature instead of the reasoner controls what evidence the reasoner receives. More concretely, this difficulty becomes salient if we wish to build an artificial reasoner to explore uncharted territory. Updating* on the evidence* 'My evidence is B' simply restates the problem that the reasoner absolutely needs prior probabilities for the type of her learning situation before she can even start to learn.

Auto–epistemic reasoning highlights the notorious Bayesian problem of finding the 'right' prior probabilities — a consequence that is too often ignored in Bayesian accounts of learning from evidence. It turns out that a Bayesian 'theory' of learning depends entirely on a theory of confirmation that constrains prior probabilities. Bayesian learning does not only depend on prior probabilities as far as prior probabilities represent the epistemic state that is to be updated by some preferred rule. What is at stake is the very admissibility of a single preferred update rule $\star$ for all prior points of view. AE–Conditionalization and Reflection are auto–epistemic criteria that regulate the admissibility of such update rules. Instead of one universal update rule $\star$ for all prior points of view, coherent auto–epistemic reasoners choose a different update rule for every single learning situation that they encounter. If Bayesianism wants to take the notion of decision–theoretic coherence seriously, it cannot offer a substantial theory of learning or updating (unless this theory takes the special form of constraints on prior probabilities). Coherent Bayesianism can only offer a *logic* of updating and of auto–epistemic reasoning with Reflection as its central principle. This conclusion is in line with the subjectivist tradition but it adds the twist that learning from evidence cannot function as an antidote to subjectivism. If the choice for prior probabilities is free, then even how a Bayesian learns from evidence is subjective.

5. Conclusion

Since neither Reliability nor Evidential Independence hold universally, Reflection (AE–Conditionalization) is incompatible with the universal use of naive Conditionalization. Ironically, the Coherence Argument is therefore an argument *against* Conditionalization as a universal update rule.

Let me summarize:

(P1) The Coherence Argument for Reflection (AE–Conditionalization) is valid.

(P2) The Coherence Argument for Conditionalization as a universal update rule is invalid because it implicitly presupposes the Partitioning Assumption (PA).

(P3) Reflection (AE–Conditionalization) is incompatible with Conditionalization *qua* universal update rule.

(C) *Ergo:* The Coherence Argument defeats Conditionalization as a universal update rule.

Notice that this criticism of Conditionalization applies to a much wider range of cases than the two objections discussed by van Fraassen (1984) who constructs (by the help of auto–epistemic vocabulary) self–referential sentences for which Conditionalization fails. The present criticism does not refer to such self–referential exceptions but applies to arbitrary propositions from $\mathcal{A}$ and is therefore considerably stronger.

As the example shows, it is easy to construct probability assignments that obey Conditionalization but violate AE–Conditionalization/Reflection and hence Diachronic Coherence. The key to these constructions is the fact that the reasoner assigns probabilities to auto–epistemic events, such as receiving posterior evidence. Of course, the reasoner might in many cases have only incomplete or even no opinions about her future probabilities. Her auto–epistemic capacities might then be too weak to determine a unique way of learning in a given learning situation or even to rule out Conditionalization. But if she were to accept Conditionalization universally, she may become unable ever to extend her auto–epistemic opinions coherently (cf. Skyrms (1987a)).

Acknowledgements: I would like to thank Richard Jeffrey and Isaac Levi for long hours of intense discussion. I am very grateful to Mathias Risse, Glenn Shafer, and Wolfgang Spohn for their generous help.

Notes

[1] This illustration shows how incoherent decision making can be systematically exploited by a second player who is aware of this incoherence.

[2] Consider the following example: If you are sure that tomorrow you will believe that pigs can fly, you should not adopt this belief today just because you want to preserve your diachronic coherence.

[3] The interpretation of Conditionalization as a universal update method seem to have been endorsed, e.g., by Teller (1973), but not, e.g., by Jeffrey (1983).

[4] The index I roughly resembles a time index. It should, however, be born in mind that the location in time of an epistemic state is not of primary interest. A certain time might elapse between two immediately subsequent epistemic states.

[5] For independent axiomatizations, cf. Spohn (1986) and van Fraassen (1976).

[6] If we simply wrote Reflection as '$P_i(A|\mathrm{SP}_j(Q)) = Q(A)$', (where Q is an arbitrary probability measure) it would follow that $Q(\mathrm{SP}_j(Q)) = P_i(\mathrm{SP}_j(Q)|\mathrm{SP}_j(Q)) = 1$ where Q can be a personal probability measure at any arbitrary stage; especially $P_i(\mathrm{SP}_j(P_i)) = 1$ for $i \neq j$. Similarly, if Reflection were not restricted to trajectories that cross each other ($P_i = Q_i$).

We may think of the two trajectories $\langle P_i \rangle_{i\in I}$ and $\langle Q_i \rangle_{i\in I}$ as the succession of the reasoner's probabilities in two different possible worlds ω_1, ω_2 and write $P_i = P_{\omega_1,i}$ and $Q_i = P_{\omega_2,i}$ (for all $i \in I$), where $P_{(.,.)} : \Omega \times I \mapsto \mathcal{P}_{\mathcal{A}}$ is a function that attaches a probability measure to every possible world at every possible point in time. Then $\mathrm{SP}_j(Q_j) = \mathrm{SP}_j(P_{\omega_2,j}) = \{\omega' | P_{\omega',j} = P_{\omega_2,j}\}$.

[7] *Proof*: For $i = j$, Reflection gives us $P_i(\mathrm{SP}_i(P_i)) = P_i(\mathrm{SP}_i(P_i)|\,\mathrm{SP}_i(P_i)) = 1$. □

[8] Hild/Jeffrey/Risse do not assume that AE–Transparency holds throughout Ω. Let $\Omega_{AE} \subseteq \Omega$ be the set of worlds $\omega \in \Omega$ in which AE–Transparency holds. For propositions $\mathrm{SP}_i(Q_i)$, define *auto–epistemic knowledge* as the operator $\mathrm{AK}(\mathrm{SP}_i(Q_i)) := \mathrm{SP}_i(Q_i) \cap \Omega_{AE}$. Thus, the proposition $\mathrm{AK}(\mathrm{SP}_i(Q_i))$ expresses that the reasoner is auto–epistemically transparent about $\mathrm{SP}_i(Q_i)$. We can then formulate the following modified principle:

(AK–Reflection) $\qquad P_i(A|\mathrm{AK}(\mathrm{SP}_j(Q_j))) = Q_j(A), \qquad i \leq j, P_i = Q_i$

[9] I use the terms 'input', 'evidence', and 'evidential input' synonymously.

[10] For further constraints on $C_i(.)$, cf. Sen (1971).

[11] Jeffrey (1983) includes the description of acts into the states of the world and maximizes the proportion $E_i(u_i(.)|f)$ of expected utility that is assigned to the states in which f is performed.

[12] Linearity is the combination of the following two conditions:

(Homogeneity) $\qquad u(s \cdot x) = s \cdot u(x)$

(Additivity) $\qquad u(x + y) = u(x) + u(y)$

[13] Schick (1986), following Davidson et al. (1955), tries to argue that the game–theoretic predicament could and maybe should *always* be blamed on non–additive utilities instead of deviant probabilities. But as it would be unrealistic always to require additive utilities, it would not be any more realistic never to allow additivity.

[14] Savage's example does not concern future choices or updates but the synchronic coherence of choice functions.

[15] We show that the formulation of Diachronic Coherence in terms of choice functions is equivalent to the earlier formulation in terms of preferences. Notice that the left and right half of the formulation in terms of choice functions are equivalent to (L) and (R), respectively:

(L) $\qquad$ If for all $C'_j \in \mathcal{C}(i,j)$: $f \in C'_j(\{f,g\})$, then $f \in C_i(\{f,g\})$.

(R) $\qquad$ If $f \in C_i(\{f,g\})$, then for some $C'_j \in \mathcal{C}(i,j)$: $f \in C'_j(\{f,g\})$.

(R), in turn, is equivalent to

(R') $\qquad$ If for all $C'_j \in \mathcal{C}(i,j)$: $f \notin C'_j(\{f,g\})$, then $f \notin C_i(\{f,g\})$.

(L) states: If it is subjectively certain that f is weakly preferred over g at j, then f is weakly preferred over g at i. (R') states: If it is subjectively certain that f is strongly preferred over g at j, then f is strongly preferred over g at i. This is the earlier version of Diachronic Coherence.

[16] It might, in practice, be difficult to find acts other than bets whose payoffs quantitatively depend on the agent's epistemic states.

[17] Remember that Reflection is based on the usual definition of conditional probabilities and therefore presupposes $P_i(\mathrm{SP}_j(Q_j)) > 0$. This restriction is essential for the proof since we could otherwise not conclude that $E_{Q_j}(u(\beta_1)) < 0$ (see below).

[18] $\beta_{1.1}$ is a conditional act of the form $\beta_{A|\mathrm{SP}_j(Q_j)}$. A conditional act $\beta_{A|B}$ is cancelled (has zero payoff) if condition B does not hold; otherwise it has the same payoff as the corresponding unconditional act β_A.

[19] Again, Generalized Iteration puts restrictions on the possible diachronic trajectories of probability measures. Reference to $\mathcal{P}(i,j)$ prevents inconsistencies.

[20] For expected–utility maximizers, the right half of Diachronic Coherence $(C_i(\{f,g\}) \subseteq \bigcup C'_j(\{f,g\}))$ implies the left half: Assume that $E_{P_i}(u(f)) > E_{P_i}(u(g))$. Hence $\exists\delta > 0 : E_{P_i}(u(f)) = E_{P_i}(u(g)) + \delta = E_{P_i}(u(g)+\delta)$. The right half implies that there is a $Q_j \in \mathcal{P}(i,j)$ such that $E_{Q_j}(u(f)) \geq E_{Q_j}(u(g)+\delta) = E_{Q_j}(u(g)) + \delta > E_{Q_j}(u(g))$.

[21] I take this to be the valid point behind the 'counter–examples' in Skyrms (1987b) Talbott (1991), Christensen (1991), and Maher (1992), (1993) .

[22] Spohn (1978) discusses very similar self–applications of the rationality concept of decision theory. Jeffrey (1988) calls this relationship between P_i and Q_j a 'reasonable transition'.

[23] Coherence Arguments were pioneered for the synchronic principles of the probability calculus by de Finetti (1936) and Ramsey (1926). Kemeny (1955) and Lehman (1955) contributed the corresponding Converse Dutch Book Theorems. Teller's (1973) seminal paper then presented a diachronic Coherence Argument by David Lewis in support of Conditionalization. Armendt (1980) supplied a similar argument for Generalized Conditionalization (Jeffrey's Rule) and Skyrms (1987c) proposes the corresponding Converse Dutch Book Theorem.

[24] Instead of AE–Transparency in the case of Reflection, these calculations are based on the Success condition $[P_i^\star B](B) = 1$ and $[P_i^\star - B](-B) = 1$.

[25] In the literature, the coherence condition is normally formulated from the point of view of an external observer with omission of the phrase 'the reasoner believes that'. The objections of this section against the supposed Dutch Book and Converse Dutch Book Theorem for Conditionalization are independent of my use of an auto–epistemic coherence condition. Simply shift back into talk from the point of view of an external observer.

[26] Let '$\mathrm{Ev}_i(e)$' express that e is part of, but not necessarily identical to, the total evidence obtained at i. Then these assumptions are:

(F1) B is reliably settled at i.
($\mathrm{Ev}_i(B) \subseteq B$, $\mathrm{Ev}_i(-B) \subseteq -B$.)

(F2) B is not settled at i or before.
($-\mathrm{Ev}_k(B) \cap -\mathrm{Ev}_k(-B) = \Omega$, if $k \leq i$.)

(F3) A is not settled at $i+1$ or before.
($-\mathrm{Ev}_k(A) \cap -\mathrm{Ev}_k(-A) = \Omega$, if $k \leq i+1$.)

(F4) The condition B settles consistently at $i+1$.
($\mathrm{Ev}_{i+1}(B) \cap \mathrm{Ev}_{i+1}(-B) = \emptyset$.)

[27] In what follows, I shall make use of another auto–epistemic addition to $\mathcal{A}$. For any piece of total evidence $e \in \mathcal{EV}$ — whether a proposition or not —, let '$\mathrm{Ev}_i^{tot}(e)$' express that e is the total evidence newly obtained at i.

[28] For the sake of the argument, let us presuppose Success and, hence, that $[P_i^\star B](B) = 1$ and $[P_i^\star - B](B) = 0$. Then,

$$E_{[P_i^\star B]}(u(\beta_1)) \;=\; P_{i+1}(A)\,(1 - P_i(A|B)) + (1 - P_{i+1}(A))\,(-P_i(A|B))$$

$$
\begin{aligned}
&+\delta\,(1-P_i(B))\,P_{i+1}(B) - \delta\,P_i(B)\,(1-P_{i+1}(B)) \\
&= -\delta\,P_i(B) \quad < 0 \\
E_{[P_i^\star - B]}(u(\beta_1)) &= -\delta P_i(B) \quad < 0.
\end{aligned}
$$

[29] If future pieces of evidence are believed to be reliable and to form a partition $\mathcal{B}$, then the probabilities for $B \in \mathcal{B}$ and for receiving B as evidence coincide ($P_i(\mathrm{Ev}_j^{tot}(B)) = P_i(B)$). By the Total Probability Theorem, Conditionalization therefore yields

$$
\begin{aligned}
P_i(A) &= \sum_{B\in\mathcal{B}} P_i(B)\cdot P_i(A|B) \\
&= \sum_{B\in\mathcal{B}} P_i(\mathrm{Ev}_j^{tot}(B))\cdot [P_i^\star B](A)
\end{aligned}
$$

Assuming that changes in probabilities are evidence driven (cf. footnote 5), this yields Generalized Iteration — as was the case with Reflection. If future evidence is not believed to form a partition or is not considered reliable, we will have $P_i(\mathrm{Ev}_j^{tot}(B)) \neq P_i(B)$.

[30] In those cases, Reflection can be written as $P_i(A|\mathrm{SP}_j(P_i^\star e)) = [P_i^\star e](A)$. Since we also have $P_i(\mathrm{SP}_{i+1}(P_i^\star e) = \mathrm{Ev}_{i+1}^{tot}(e)) = 1$ (i.e., $P_i(\mathrm{SP}_{i+1}(P_i^\star e) - \mathrm{Ev}_{i+1}^{tot}(e)) = P_i(\mathrm{Ev}_{i+1}^{tot}(e) - \mathrm{SP}_{i+1}(P_i^\star e)) = 0$), this is equivalent to AE–Conditionalization. For more details on how auto–epistemic reasoning relates to updating, cf. Hild (1998).

[31] Armendt (1980) bases his supposed Dutch Book Theorem for Generalized Conditionalization on a version of Evidential Independence. Skyrms (1987a) provides conditions under which AE–Conditionalization validates Entropy Maximization.

[32] Remember that Success plus Rigidity is equivalent to (Generalized) Conditionalization.

[33] I write $P(A \subseteq B) = 1$ if and only if $P(A - B) = 0$. By Reliability, we have $P(\mathrm{Ev}_i^{tot}(B) \subseteq B) = 1$ and $P(\mathrm{Ev}_i^{tot}(-B) \subseteq -B) = 1$. By (PA), $P(\mathrm{Ev}_{i+1}^{tot}(B) = -\mathrm{Ev}_{i+1}^{tot}(-B)) = 1$. This yields $P(B = \mathrm{Ev}_{i+1}^{tot}(B)) = 1$, and hence Evidential Independence.

[34] Similar examples are the Three Prisoners Puzzle and the Monty Hall Problem.

[35] Notice that (i) this notion of total evidence* automatically satisfies Evidential Independence and that (ii) it is the only notion of evidence that is guaranteed to satisfy this assumption.
Carnap's requirement of total evidence, if construed broadly enough, can be taken to support this notion of evidence.

[36] Cf. Good's (1950) story of the 'statistician's stooge'. The stooge is a reliable measuring device that satisfies the Partitioning Assumption.

References

Armendt, B.: 1980, 'Is there a Dutch Book Argument for Probability Kinematics?', *Philosophy of Science* **47**, 583-588.

Carnap, R.: 1950, *Logical Foundations of Probability*, Chicago: University of Chicago.

Christensen, D.: 1991, 'Clever Bookies and Coherent Beliefs', *Philosophical Review* **100**, 229-247.

Diaconis, P., Zabell, S.L.: 1982, 'Updating Subjective Probability', *Journal of the American Statistical Association* **77**, 822-830.

Davidson, D., McKinsey, J.C.C., Suppes, P.: 1955, 'Outlines of a Formal Theory of Value', *Philosophy of Science* **22**, 140-160.

de Finetti, B.: 1936, 'Foresight: Its Logical Laws, Its Subjective Sources', in: Kyburg, H.E., Smolker, H.E. (eds.): 1980, *Studies in Subjective Probability*, Huntington (N.Y.): Krieger, 2nd edition.

de Finetti, B.: 1974, *Theory of Probability*, vol. I, New York/London/Sydney: Wiley.

Freund, J.E.: 1965, 'Puzzle or Paradox?', *American Statistician* **19**, 29-44. For the discussion that followed cf. **20**, 34-37 and **21**, 38-51.

Gärdenfors, P.: 1988, *Knowledge in Flux*, Cambridge (Mass.): MIT.

Goldstein, M.: 1983, 'The Prevision of Prevision', *Journal of the American Statistical Association* **78**, 817-819.

Good, I. J.: 1950, *Probability and the Weighing of Evidence*, London: Charles Griffin.

Hacking, I.: 1967, 'Slightly More Realistic Personal Probabilities', *Philosophy of Science* **34**, 311-325.

Hild, M.: 1998, 'Auto–Epistemology and Updating', forthcoming in *Philosophical Studies*.

Hild, M., Jeffrey, R., Risse, M.: 1998, 'Aumann's 'No Agreement Theorem' Generalized', forthcoming in: Bicchieri, C., Jeffrey, R., Skyrms, B. (eds.): 1998, *The Logic of Strategy*, Oxford: OUP.

Howson, C.: 1995, 'Theories of Probability', *British Journal for the Philosophy of Science* **46**, 1-32.

Jeffrey, R.C.: 1983, *The Logic of Decision*, 2nd edn, Chicago: University of Chicago Press.

Jeffrey, R.C.: 1988, 'Conditioning, Kinematics, and Exchangeability', in: Skyrms, B., Harper, W.L. (eds.), *Causation, Chance, and Credence*, vol. I, 221-255, Dordrecht: Kluwer; also in: Jeffrey (1992), 117-153.

Jeffrey, R.C.: 1992, *Probability and the Art of Judgment*, Cambridge: CUP.

Kemeny, J.G.: 1955, 'Fair Bets and Inductive Probabilities', *Journal of Symbolic Logic* **20**, 263-273.

Lehman, R.S.: 1955, 'On Confirmation and Rational Betting', *Journal of Symbolic Logic* **20**, 251-262.

Maher, P.: 1992, 'Diachronic Rationality', *Philosophy of Science* **59**, 120-141.

Maher, P.: 1993, *Betting on Theories*, Cambridge: CUP.

Ramsey, P.: 1926, 'Truth and Probability', in: Braithwaite, R.B. (ed.), *The Foundations of Mathematics and Other Essays*, New York: Harcourt Brace, 1978, 58-100.

Savage, L.J.: 1954, *The Foundations of Statistics*, 1972, 2nd edn, New York: Dover.

Schick, F.: 1986, 'Dutch Bookies and Money Pumps', *The Journal of Philosophy* **83**, 112-119.

Sen, A.: 1971, 'Choice Functions and Revealed Preference', *Review of Economic Studies* **38**, 307-317.

Shafer, G.: 1985, 'Conditional Probability', *International Statistical Review* **53**, 261-277.

Skyrms, B.: 1980, *Causal Necessity*, New Haven/London: Yale University.

Skyrms, B.: 1987a, 'Updating, Supposing, and Maxent', *Theory and Decision* **22**, 225-246.

Skyrms, B.: 1987b, 'The Value of Knowledge', in: Savage, C.W. (ed.), *Justification, Discovery and the Evolution of Scientific Theories*, Minneapolis: University of Minnesota Press.

Skyrms, B.: 1987c, 'Dynamic Coherence and Probability Kinematics', *Philosophy of Science* **54**, 1-20.

Spohn, W.: 1978, *Grundlagen der Entscheidungstheorie*, Kronberg (Taunus): Scriptor.

Spohn, W.: 1986, 'The Representation of Popper Measures', *Topoi* **5**, 69-74.

Talbott, W.J.: 1991, 'Two Principles of Bayesian Epistemology', *Philosophical Studies* **62**, 135-150.

Teller, P.: 1973, 'Conditionalization and Observation', *Synthese* **26**, 218-258.

van Fraassen, B.: 1976, 'Representation of Conditional Probabilities', *Journal of Philosophical Logic* **5**, 417-430.
van Fraassen, B.: 1980, 'A Temporal Framework of Conditionals and Chance', *Philosophical Review* **89**, 91-108.
van Fraassen, B.: 1984, 'Belief and the Will', *The Journal of Philosophy* **81**, 235-256.
van Fraassen, B.: 1995, 'Belief and the Problem of Ulysses and the Sirens', *Philosophical Studies* **77**, 7-37.
van Fraassen, B.: 1986, 'A Demonstration of the Jeffrey Conditionalization Rule', *Erkenntnis* **24**, 17-24.
Walley, P.: 1991, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall.
Williams, P.M.: 1980, 'Bayesian Conditionalisation and the Principle of Minimum Information', *British Journal for the Philosophy of Science* **31**, 131-144.

Address for correspondence: Christ's College, Cambridge CB2 3BU (GB), email: mh247@cam.ac.uk