Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection
Author(s): Kenny Easwaran
Reviewed work(s):
Source: *Philosophy of Science*, Vol. 80, No. 1 (January 2013), pp. 119-142
Published by: The University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/10.1086/668879
Accessed: 30/01/2013 13:41

Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection

Kenny Easwaran*†

Expected accuracy arguments have been used by several authors (Leitgeb and Pettigrew and Greaves and Wallace) to support the diachronic principle of conditionalization, in updates where there are only finitely many possible propositions to learn. I show that these arguments can be extended to infinite cases, giving an argument not just for conditionalization but also for principles known as 'conglomerability' and 'reflection'. This shows that the expected accuracy approach is stronger than has been realized. I also argue that we should be careful to distinguish diachronic update principles from related synchronic principles for conditional probability.

1. Introduction. The traditional 'Dutch book' and 'representation theorem' arguments for probabilism (the thesis that a rational agent's degrees of belief ought to satisfy the probability axioms) rely on a connection between credences and practical rationality that some have found problematic.[1] Starting at least with Joyce (1998), many formal epistemologists have sought to replace these arguments with more purely epistemic ones that focus on a notion of 'accuracy'. This is understood as a purely epistemic good, representing something like the 'distance' between one's credences and the actual truth values of the propositions one considers. In this article, I will extend some recent

Received April 2012; revised July 2012.

*To contact the author, please write to: University of Southern California, 3709 Trousdale Parkway, Los Angeles, CA 90089; e-mail: easwaran@usc.edu.

†I would like to thank the Coalition of Los Angeles Philosophers, audiences at the math department at the University of Southern California and the philosophy department at the University of Konstanz, Richard Pettigrew, and an anonymous referee for this journal, for their helpful comments.

1. I use the terms "degree of belief" and "credence" interchangeably throughout, depending on which seems more stylistically convenient.

Philosophy of Science, 80 (January 2013) pp. 119–142. 0031-8248/2013/8001-0009$10.00

arguments of this sort, to show that accuracy considerations can give stronger requirements than have been thought. This can be seen as either a strength or a weakness for the accuracy-based approach.

Hannes Leitgeb and Richard Pettigrew (2010a, 2010b) use accuracy conditions to support not only probabilism but also the diachronic norm of conditionalization (that one's credence in A after learning B is given by $P_{\text{new}}(A) = P_{\text{old}}(A \wedge B)/P_{\text{old}}(B)$, when that is defined), which is characteristic of Bayesianism. They show that their arguments work if the set of worlds that are taken to be possibilities by the agent is finite, and ask at the end of each paper how their arguments fare if the set of worlds is allowed to be infinite.

I will show that the argument does extend to infinite sets of worlds, if we make a few natural assumptions about how to understand various concepts that they only explicitly discuss for finite cases. Consideration of the methodology in infinite settings gives arguments for even more norms beyond the ones that Leitgeb and Pettigrew endorse. In particular, it seems to give arguments for principles like those known as 'conglomerability' and 'reflection'. (See sec. 3.2 for precise statements of these principles.) This may be seen as a further strength of the methodology, or (because of existing arguments against these principles) a reductio of the methodology, or a call to further investigate the methodology to find reasons to accept the arguments of Leitgeb and Pettigrew without accepting these further conclusions. It is not my goal in this article to decide between these three responses, but merely to propose them as questions for future research.

My arguments will be similar to Leitgeb and Pettigrew's argument for conditionalization (sec. 6.2 of their second paper) rather than to their argument for probabilism (sec. 6.1) or their more tentative arguments for uniform distribution (sec. 6.3) and a highly nonstandard alternative to Jeffrey conditionalization (sec. 7). This argument is very similar to the one given by Hilary Greaves and David Wallace (2006, 608) that, "under independently motivated conditions, conditionalization maximizes expected epistemic utility." As I set up my framework, I will show how my argument relates to each of these.

Leitgeb and Pettigrew (2010a, 209) say "there is nothing particularly philosophical about our decision to stick to the case of finitely many worlds in this article; we simply assume this to be so and postpone the discussion of the infinite case to another time." Similarly, Greaves and Wallace (2006, 611 n. 1) state, "We will assume throughout that $\mathcal{S}$ [the set of states] is finite. This is merely for simplicity of exposition." (They similarly assume that no nonempty proposition has credence 0.) As it turns out, even describing the argument requires some more substantive assumptions if $\mathcal{S}$ can be infinite rather than if it is finite. However, the appropriately modified arguments do work, giving an argument not just for conditionalization but for (versions of) conglomerability and reflection.

In section 2, I will outline the features used in the infinite version of the argument. The first four parts discuss the general framework, and the fifth discusses special issues surrounding the notion of expected value in an infinite setting. The section concludes with precise statements of my three main theorems: roughly, that a conditional probability function that minimizes expected inaccuracy satisfies the ratio formula (when that is defined) and takes on a range of values whose expectation is the unconditional probability, which must therefore lie within the range spanned by the conditional probabilities. (Proofs are given in the appendix.)

In section 3, I discuss the interpretation of these theorems. In particular, I consider their import for both synchronic coherence of an agent's attitudes and what I consider more speculative diachronic norms like those suggested by Leitgeb and Pettigrew, and Greaves and Wallace. I end with a brief conclusion and an appendix containing proofs of the theorems.

2. Assumptions of the Setting.

2.1. Propositions, Probability, and Measurability. For the purposes of this discussion, I treat the objects of doxastic attitudes as propositions and think of propositions as sets of possibilities.[2] Each proposition is identified with a subset of the set $\mathcal{S}$ of possibilities for the agent.

I will also assume that the synchronic norms of probabilism have already been established, perhaps by means of the argument that Leitgeb and Pettigrew use, or perhaps independently. Thus, I will assume that the agent's credences satisfy the axioms of probability theory, and, to emphasize this fact, I will use 'P' as the name for the function that represents the credences. To be explicit, this means:

1. Whenever $P(A)$ is defined, $P(A) \geq 0$.
2. $P(\mathcal{S}) = 1$.
3. If A and B are disjoint subsets of $\mathcal{S}$, then, whenever $P(A)$ and $P(B)$ are both defined, $P(A \cup B)$ is defined and equal to $P(A) + P(B)$.

The condition that these credences are defined turns out to be important for the generalization to the infinite case. It is well-known that in many infinite sets, some natural constraints on probabilities are incompatible with

2. I will sometimes refer to these possibilities as 'states' or 'worlds', but I am not committed to any particular metaphysical view of what they are, apart from the claim that an agent considers two propositions compatible if and only if the intersection of the corresponding sets is nonempty.

every subset having a probability. Instead, we must assign probabilities only to some collection $\mathcal{M}$ of the 'measurable' subsets of the space. In standard measure theory, it is assumed that $\mathcal{M}$ is a σ-field. This means:

1. $\mathcal{S}$ is a member of $\mathcal{M}$.
2. If A is in $\mathcal{M}$, then $\neg A$ (the complement of A in $\mathcal{S}$) is in $\mathcal{M}$.
3. If $A_1, A_2, \ldots$ are countably many members of $\mathcal{M}$, then their union $\cup_{i=1}^{\infty} A_i$ is in $\mathcal{M}$.

By suitable applications of these assumptions, one can show that countable intersections of members of $\mathcal{M}$ are in $\mathcal{M}$ too.

Thus, I will not assume that the agent has credences in every proposition, but I will let $\mathcal{F}$ be the set of propositions for which she does have credences, and I will assume that $\mathcal{F}$ is a σ-field. Of course, if the agent does have credences in every proposition, then $\mathcal{F}$ is trivially a σ-field. I suspect that the closure properties of $\mathcal{F}$ can be weakened somewhat, though this may complicate some of the arguments.

2.2. Updates and Experiments. Not all updates fall in the scope of my argument. I will only consider updates on which $\mathcal{F}$ remains fixed. Situations in which the agent comes to have credences in new propositions seem very different from the standard examples where an agent just learns that some proposition is true. For instance, two natural cases where $\mathcal{F}$ would be expected to change are the introduction of a new theory, and the discovery of confusion in one's old concepts, either one of which might be part of something like a scientific revolution. Bayesians already know (Earman 1992) that these cases are difficult ones to account for.

I will follow Greaves and Wallace in considering an update situation as being given in advance by some experiment: "[An experiment is] a situation in which the agent is to receive some new piece of information, from among a set of mutually exclusive and jointly exhaustive alternatives. Mathematically, an experiment is represented by a partition $\mathcal{E}$ of $\mathcal{S}$" (Greaves and Wallace 2006, 611). To say that $\mathcal{E}$ is a partition of $\mathcal{S}$ just means that $\mathcal{E}$ is a collection of propositions such that each member of $\mathcal{S}$ is in exactly one of the propositions in $\mathcal{E}$. Since their $\mathcal{S}$ is finite, so is any experiment $\mathcal{E}$ for them. But I will allow $\mathcal{E}$ to be infinite.

As a simple example, we can imagine that the agent is about to flip a coin and observe how it lands. We can think of $\mathcal{E}$ as the set $\{H, T\}$, where H is the set of all possibilities in which the coin lands "heads," and T is the set of all possibilities in which the coin lands "tails." For a more complicated example, consider the experiment of consulting an oracle to ask how many humans will ever have lived. In this case, it seems reasonable that any sufficiently large integer is an epistemically possible outcome, though we might imagine that there are other possible outcomes as well, such as one in which the oracle says, "I don't know." The experiment is represented by the set of distinguishable outcomes for the agent—in this case, rather than one member of $\mathcal{E}$ being the set of worlds in which there are exactly 10 trillion humans over the history of the universe, the relevant member of $\mathcal{E}$ is the set of worlds in which the oracle says "10 trillion" in answer to the question. Yet more complicated examples may arise in cases where the agent is about to throw an infinitely thin dart at a dart board—if the agent can distinguish each infinitely small point of the board, then the elements of $\mathcal{E}$ will correspond to each possible landing site of the dart, but if the agent cannot distinguish these points perfectly, then the elements of $\mathcal{E}$ will be much more complicated.

As suggested by my examples, the sort of experiment I consider is one in which the agent comes to have credence 1 in exactly one member of $\mathcal{E}$ and credence 0 in all the rest. Cases in which the update is generated in some other way (e.g., cases driven by conditional probability [van Fraassen 1981], cases involving potential memory loss [Arntzenius 2003], and cases calling for Jeffrey conditionalization—or the alternative to it that Leitgeb and Pettigrew [2010b, 252–60.] propose) are beyond the scope of the arguments considered here. However, the notion of conditional probability is traditionally thought to play an important role in these updates, and if the concept of conditional probability can be clarified by my arguments, then they will at least indirectly help us understand these more complicated updates.

I will also assume that $\mathcal{E}$ is a subset of $\mathcal{F}$—each distinguishable outcome of the experiment had some credence for the agent prior to the experiment being performed. In cases where some elements of $\mathcal{E}$ are not in $\mathcal{F}$, the agent has a possibility of coming to know some proposition that she does not already grasp—thus, they are outside the scope of my arguments. Updates involving indexicals like 'here' or 'today' have this feature, so my arguments do not directly apply to "Sleeping Beauty" or related cases (Meacham 2008; Titelbaum 2008; Bradley 2012).

2.3. Plans and Availability. I will consider an update as proceeding according to a set of "plans" for an experiment.[3] For a given experiment $\mathcal{E}$, and for each proposition A in $\mathcal{F}$, I will represent the agent's plan for updating her credence in A after performing $\mathcal{E}$ as a function $f_{A,\mathcal{E}}$ that gives a real number in each state. For any state s in $\mathcal{S}$, $f_{A,\mathcal{E}}(s)$ is the credence that the agent plans to have in A after learning the outcome of $\mathcal{E}$, if s is the

3. The terminology of "plans" seems to suggest a sort of voluntarism about belief that I do not mean to endorse. I will consider these "plans" as potential dispositions to update in a particular situation, with the apparatus of decision theory used to evaluate them. For more discussion of this issue, see Greaves and Wallace (2006, 612).

actual state. When $\mathcal{E}$ is understood from context, I will sometimes refer to just f_A.

This use of "plans" distinguishes my framework slightly from those of Leitgeb and Pettigrew, and Greaves and Wallace. Leitgeb and Pettigrew think of an update as given by a single proposition E that the agent becomes certain of. In the case where $\mathcal{E}$ is finite and all its members have positive prior credence, it will turn out that the difference is inessential. However, when some members of $\mathcal{E}$ have prior credence 0 (and in many cases, such as ones where the experiment consists of learning the precise real value of some unknown physical parameter, it may be the case that all members of $\mathcal{E}$ have prior credence 0), the partition will play an essential role in the argument. And once we have a partition rather than a single proposition, something like this notion of "plan" is essential.

Greaves and Wallace talk of "acts" whereas I talk of "plans." The main difference is that while their acts specify a complete probability distribution for the agent in each state, my plans only specify the agent's credence for a single proposition. This primarily helps simplify some of the discussion, but it is also connected to the distinction between global and local inaccuracy measures, to be discussed in section 2.4.[4]

Of course, given the notion of updates and experiments I am working with, not just any such function will correspond to a possible plan. Only a subset of these functions will be said to be "available." Since the experiment represents the information the agent comes to know during the course of the update, an agent's plans must make the same response in states where the outcome of the experiment is the same.

A plan f_A is said to be available (or "$\mathcal{E}$-available," if I need to make explicit which partition is relevant) if and only if it meets the following two conditions:

1. If $E_i \in \mathcal{E}$ and $s_1, s_2 \in E_i$, then $f_A(s_1) = f_A(s_2)$.
2. If x is any rational number, then $S_x \in \mathcal{F}$, where S_x is the set of states s for which $f_A(s) > x$.

A plan that did not satisfy the first constraint would require the agent's credence function to depend on information that she does not have, even once the experiment is performed. The terms s_1 and s_2 are both compatible with all the information that she has so far (which is what it means to be in $\mathcal{S}$), and

4. This also leads to a potential worry. An agent may have plans to update her credences in many propositions, and in some cases these plans may not be coherent—that is, there may be some state in which following all these plans simultaneously results in credences that do not satisfy the probability axioms. However, as I mention in sec. 2.6, this can always be avoided if the agent's initial credences are countably additive.

since what she learns on completing $\mathcal{E}$ is just which member of $\mathcal{E}$ is actual, she does not learn anything to distinguish them. Thus, she cannot plan to act differently in those two states.

The second constraint is a bit more subtle. If $\mathcal{E}$ is a finite or countable partition, then it is not a further requirement but in fact follows from the first one, together with the fact that $\mathcal{F}$ is a σ-algebra. So it is only a further assumption if $\mathcal{E}$ is uncountably infinite. In planning whether to have credence greater than x, the agent is responding to whether the actual world is in S_x or not, and it seems that this responsiveness should require that S_x be a proposition she can grasp. (Readers who are familiar with the language of measure theory will recognize this second constraint as the requirement that f_A be measurable with respect to the agent's credence function, which is essential for the notion of expected value to make sense.)

2.4. Inaccuracy, Global and Local. The arguments under consideration are all based on a function I that measures 'inaccuracy' of a belief state. The one I will use is what Leitgeb and Pettigrew call a 'local inaccuracy measure', which is a function such that $I(A, s, x)$ is a real number that measures how epistemically bad it is to have credence x in proposition A in state s. Greaves and Wallace instead describe their function as an 'epistemic utility function', but in the terminology of Leitgeb and Pettigrew, it behaves like a 'global inaccuracy measure'. This is a function U such that $U(s, p)$ is a real number that measures how good or bad the overall credence function p is in state s.

Leitgeb and Pettigrew argue that agents should have both local and global inaccuracy measures and that these should cohere in a particular way. As a result, they claim that the local inaccuracy function should be given by the Brier score, which is defined by $I(A, s, x) = (1 - x)^2$ if A is true in s and $I(A, s, x) = x^2$ if A is false in s. They also argue that the global inaccuracy function should be given by the sum of the local inaccuracies over all propositions. Their arguments are interesting, but they rely on some controversial assumptions about the "geometry of reason" (Leitgeb and Pettigrew 2010a, 210; see also their theorems 3, 4, and 5 on 222–29). In addition, these arguments cannot work when there are infinitely many propositions to be considered—their version of global inaccuracy will normally be infinite and thus cannot relate to local inaccuracy in any interesting way.[5]

5. As an anonymous referee pointed out, a semiglobal inaccuracy measure summing just over propositions in a single partition will be finite in this case. However, this involves choosing a second partition to evaluate, beyond the partition generating the update. The referee also points out that a more traditional definition of the Brier score, on which it is averaged over all propositions instead of summed, may be able to give some sort of finite limit, but this would still require a substantial modification of the argument of Leitgeb and Pettigrew.

Greaves and Wallace make no claims about local inaccuracy measures and very few claims about the global inaccuracy measure that they use. This gives them some greater generality but also requires that they consider an update in terms of "acts" that specify posterior credences for all propositions, rather than my "plans" that specify posterior credence for only a single proposition at a time.

I will instead work just with local inaccuracy measures. I will make five assumptions about the local inaccuracy function $I(A, s, x)$. (These assumptions are weaker than the ones that Leitgeb and Pettigrew make and have parallels in other discussions of inaccuracy, such as Joyce 1998.)

Extensionality. First, I assume that the inaccuracy of having credence x in proposition A does not depend on any feature of the state s other than whether A is true or false. Thus, the function can be thought of as two functions, $I(A, 1, x)$ and $I(A, 0, x)$, the former assigning inaccuracies in states where A is true and the latter in states where A is false. In cases where A is clear from context, I will suppress the first argument for the function and just talk about $I(1, x)$ and $I(0, x)$. (One might further assume that the function is in fact identical for each proposition A, but this is not needed for my arguments.)

Monotonicity. Second, I also assume that for each A, $I(A, 0, x)$ has its only local or global minimum when $x = 0$, while $I(A, 1, x)$ has its only local or global minimum when $x = 1$. That is, $I(A, 0, x)$ is monotonically increasing and $I(A, 1, x)$ is monotonically decreasing on $[0, 1]$. If we consider the possibility of credences outside of $[0, 1]$, then both functions should increase as x moves away from this interval.

Bounded Continuity. Third, I assume that for any positive ε and each A there is a set of finitely many values $0 = x_0 < x_1 < \ldots < x_n = 1$ such that, for every $i < n$,

$$I(A, 0, x_{i+1}) - I(A, 0, x_i) < \varepsilon$$

and

$$I(A, 1, x_i) - I(A, 1, x_{i+1}) < \varepsilon.$$

(Given Monotonicity, this is equivalent to the functions being bounded and continuous on $0 \leq x \leq 1$.)

Convexity. Fourth, for any $0 \leq x < y \leq 1$ and any $0 < \lambda < 1$, I assume that $I(0, \lambda x + (1 - \lambda)y) < \lambda I(0, x) + (1 - \lambda)I(0, y)$, and similarly $I(1, \lambda x + (1 - \lambda)y) < \lambda I(1, x) + (1 - \lambda)I(1, y)$. As illustrated in figure 1, this means

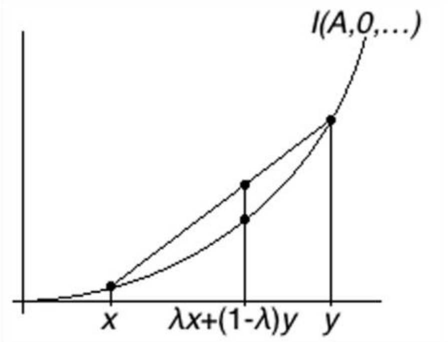

Figure 1.

that the inaccuracy of any credence intermediate between x and y is always less than predicted by linear interpolation between the inaccuracies of x and of y. Thus, not only is inaccuracy increasing as one's estimate gets farther from the truth value, but the rate of increase always increases as well.

Immodesty. Finally, I assume that if $0 \leq c \leq 1$ then the unique value of x that minimizes $cI(A, 1, x) + (1 - c)I(A, 0, x)$ is $x = c$. Once I introduce the notion of 'expected inaccuracy' (in sec. 2.5), we will see that, for someone who has $P(A) = c$, this is their estimate of the inaccuracy of someone who has credence x in A. If $x = c$ were not the unique minimum, then an agent with credence c would be 'modest', in the sense that she would think that some other credence was at least as accurate as her own. This condition is extensively discussed in the literature, sometimes under the name 'propriety' (Greaves and Wallace 2006; Joyce 2009; Myrvold 2012).

These five assumptions are all satisfied by the Brier score, so they follow from the assumptions made by Leitgeb and Pettigrew. They also follow from Joyce's assumptions on accuracy and are satisfied by the other scoring rules discussed on page 275 of Joyce (2009). Greaves and Wallace assume only the equivalent of Immodesty for their global inaccuracy measure. But since this is the most arbitrary-seeming assumption, my argument is no worse off than theirs.

2.5. Expected Inaccuracy. Before giving the proofs, I need to explain a bit more about 'expected inaccuracy' and show why this notion makes sense. If V is a random variable (i.e., a function that takes on a real value in each state s in $\mathcal{S}$), then the 'expected value of V', notated as $Exp(V)$, represents something like the agent's best estimate of V. It depends on the values for V in each possible state and the agent's credence function P. Although $Exp(V)$ may be distinct from every possible value of V (just as the average family may have 2.3 children even though no particular family has exactly that many), it traditionally plays a role in guiding one's actions in light of one's uncertainty about V.

We can stipulate a definition of the expected value of V with respect to P, as Leitgeb and Pettigrew do. However, they consider at least two different stipulations ($Exp(V) = \Sigma_{s \in S} P(s)V(s)$, or $Exp(V) = \Sigma_v vP(V = v)$), which give distinct results for some infinite cases, and both seem clearly wrong for others.[6] And there are further worries that arise in taking a sum of infinitely many terms, especially if the series is uncountable.

Thus, instead of stipulatively defining expectation as Leitgeb and Pettigrew do, I will adopt the axiomatic approach of Whittle (2000). For random variables $X, X_1, X_2, \ldots$, Whittle's axioms state:

1. If $X(s) \geq 0$ for all s, then $Exp(X) \geq 0$.
2. If c is a constant, then $Exp(cX) = cExp(X)$.
3. $Exp(X_1 + X_2) = Exp(X_1) + Exp(X_2)$.
4. For an event A, let 1_A be the random variable with $1_A(s) = 1$ if A is true in s and $1_A(s) = 0$ if A is false in s. Then $Exp(1_A) = P(A)$.[7]
5. If for every s the sequence $X_1(s), X_2(s), \ldots$ is monotonically increasing and converges to $X(s)$, then the sequence $Exp(X_1), Exp(X_2), \ldots$ is monotonically increasing and converges to $Exp(X)$.

Axiom 1 seems essential to the concept of an expected value. A random variable cannot have a negative expected value unless it is at least possible for it to take on a negative value. Axiom 4 gives the only direct connection between probability and expected value. Axioms 2, 3, and 5 provide linearity conditions on expected value. Axiom 3 is often a surprise to students when they learn that it holds regardless of whether X_1 and X_2 are independent or dependent. However, once one has convinced oneself that it holds, it is plausible that it is in fact essential to the notion of expected value and not merely a mathematical consequence of a formal definition.

If V is a measurable random variable that always takes either the value v or the value 0, then axioms 2 and 4 guarantee that $Exp(V) = vP(V = v)$. If V is a measurable random variable that takes on only finitely many pos-

6. Consider a case where S is infinite and every member of S has probability 0. For a case where the two stipulations are different, let V be a random variable that takes on the value 1 in every state. The first stipulation says $Exp(V) = 0$ (because it is a sum of terms that are all 0) while the second says $Exp(V) = 1$. For a case where both are wrong, let V' be a random variable that takes a different positive value in every state. Both stipulations say $Exp(V') = 0$, even though it seems clear that its expected value should be positive.

7. This is actually a combination of Whittle's (2000) axiom 4 on p. 15 and his definition of probability on p. 17.

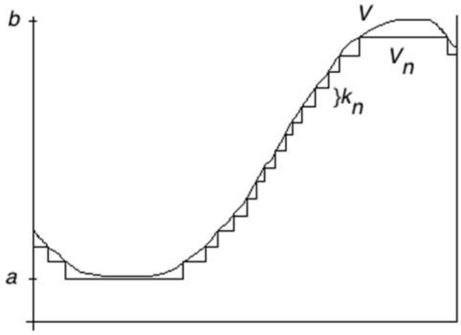

Figure 2.

sible values, then it is the sum of finitely many random variables of the previous sort, and thus, applying axiom 3, we see that $Exp(V) = \Sigma vP(V = v)$.

If V is a measurable random variable, and it is bounded (there are a and b with $a < V(s) < b$ for all s), then this previous result together with axiom 5 will be sufficient to uniquely determine $Exp(V)$. This is because we can define a sequence of random variables V_n that are monotonically increasing and converge to V, each one of which is measurable and takes on only finitely many possible values. In particular, we define $k_n = (b - a)/2^n$, and we define $V_n(s) = a + ik_n$ whenever $a + ik_n \leq V(s) < a + (i + 1)k_n$. (See fig. 2.)

The result of the previous paragraph suffices to define each $Exp(V_n)$, and axiom 5 guarantees that $Exp(V)$ is the limit of these values. (This is in fact equivalent to the standard measure-theoretic definition of the expected value of a bounded measurable random variable as a Lebesgue integral.)

If we let f_A be some available plan for updating one's credence in A for some experiment $\mathcal{E}$, and if $V(s)$ is the inaccuracy $I(A, s, f_A(s))$, then we can see that V is measurable (because $\mathcal{E}$ is measurable, so f_A is measurable, and I satisfies Extensionality and Monotonicity, so it is measurable as well) and bounded (because I satisfies Bounded Continuity). Thus, Whittle's axioms for expectation suffice to define the expected inaccuracy of any plan of the sort I consider.

Without Extensionality, inaccuracy might depend on some proposition for which the agent does not have a credence, so it would not have an expectation. Without Monotonicity and Bounded Continuity, inaccuracy could be unbounded in a way that results in its expected value being infinite or undefined. Thus, some version of these first three assumptions is surely needed for expected inaccuracy to be the sort of thing one could try to minimize.

The use of expected inaccuracy distinguishes the arguments of Leitgeb and Pettigrew, and that of Greaves and Wallace, from some other accuracy-based arguments. For instance, Lindley (1982) and Joyce (1998) both give arguments for probabilism (Lindley argues for the ratio analysis of conditional probability as well) based merely on the requirement that no other credence function dominate the one under consideration (one credence function "dominates" another if and only if the first has lower inaccuracy in every state). Since one credence function can have lower expected inaccuracy than another without being dominated, the Joyce and Lindley results are correspondingly stronger. All of these arguments, other than that of Greaves and Wallace, assume a strong connection between global and local accuracy.

My argument (like that of Leitgeb and Pettigrew) makes the stronger assumptions of each, that utility should be defined by a local inaccuracy measure and that agents should rule out all plans with higher expected inaccuracy rather than merely plans that are dominated. If there is some problem in principle with local inaccuracy measures or expected inaccuracy, then this would give a reason to reject my conclusions while accepting the conclusions of some of the others. However, my arguments apply even when $\mathcal{S}$ is infinite, and I give results leading to versions of conglomerability and reflection, in addition to the ratio formula for conditional probability. (As an anonymous referee has pointed out, the authors that give dominance arguments rather than expectation arguments are trying to establish probabilism as well as conditionalization. If probabilism is not already taken for granted, then the use of expectation is clearly problematic. But once probabilism is granted, as I have assumed here, the use of expectation rather than dominance may not be a real weakness.)

2.6. The Theorems. The following results are proved in the appendix.

> **Ratio Theorem.** If A is a measurable proposition, $\mathcal{E}$ is an experiment, E is an element of $\mathcal{E}$ with positive probability, and f_A has minimal expected inaccuracy among the $\mathcal{E}$-available plans for updating one's credence in A, then $f_A(s) = P(A \wedge E)/P(E)$ for any s in E.

This theorem is the same as theorem 3 in Leitgeb and Pettigrew (2010b, 250), except for my weaker assumptions on I and the measurability assumptions needed for the infinite case. This is also the same as corollary 1 that Greaves and Wallace state on page 625, except that they only assume Immodesty, while I use a local inaccuracy measure and the measurability assumptions for infinite sets.

> **Interval Theorem.** If A is a measurable proposition, $\mathcal{E}$ is an experiment, and f_A has minimal expected inaccuracy among the $\mathcal{E}$-available plans for updating one's credence in A, then there is no x with $P(A) < x$ and $x < f_A(s)$ for all s, and there is no x with $P(A) > x$ and $x > f_A(s)$ for all s. That is, $inf_{s\in\mathcal{S}} f_A(s) \leq P(A) \leq \sup_{s\in\mathcal{S}} f_A(s)$.

This theorem basically says that no optimal update plan guarantees an increase in credence, or guarantees a decrease in credence. For every experiment that could increase one's credence in A by some positive ε, there is always some other possible outcome of the experiment that would fail to do so.

In many cases, the content of this theorem can be better understood by looking at a more specific consequence:

> **Expectation Theorem.** If, for all propositions A in $\mathcal{F}$, the conditions of the Interval Theorem hold, and for each s, $f_{\ldots}(s)$ is a probability function (so that $f_A(s) \geq 0$ for all A, $f_S(s) = 1$, and $f_{A\cup B}(s) = f_A(s) + f_B(s)$ whenever A and B are disjoint), then for all A, $P(A) = Exp[f_A(s)]$.

That is, the expected result of an update plan for A should equal one's prior credence in A.

All of these theorems tell us what plans that minimize expected inaccuracy must be like, if they exist. Fortunately, such plans do exist. (The space of possible plans is the set of functions from $\mathcal{E}$ to $[0, 1]$, which is compact, by Tychonov's theorem. Expected inaccuracy is a continuous function on the space of plans and must therefore achieve its minimum value.)

The extra condition in the statement of the Expectation Theorem is a bit subtle. In particular, since I have a separate plan for each proposition, I have not assumed that the plans will combine to yield a coherent posterior probability function. And in fact, if $\mathcal{E}$ is a countable partition of propositions whose credences fail to sum to 1, then it is known that the conclusion of the Expectation Theorem must be violated, so there can be no set of plans that cohere probabilistically (Hill and Lane 1985, 369). However, if the unconditional credences satisfy countable additivity (as orthodox Bayesianism assumes), then the Radon-Nikodym theorem of real analysis guarantees the existence of coherent update plans that satisfy the conclusions of all three theorems. It remains to be seen if such plans must always minimize expected inaccuracy, but it seems plausible that they will.

When $\mathcal{E}$ is finite, the Interval Theorem and Expectation Theorem are trivial consequences of the Ratio Theorem. It is only in the infinite case that these provide additional content. Thus, the normative claims about updating and conditional credence that I discuss in the next section all apply to the finite case as a consequence of the arguments by Leitgeb and Pettigrew or Greaves and Wallace. My contribution is to show that they hold in the infinite case as well.

3. Normative Applications of the Theorems.

3.1. Accuracy Only? In much of the literature on accuracy, it is assumed that accuracy is the only normatively significant feature of a credence function. Leitgeb and Pettigrew (2010b, 244) defend this by arguing that "the ultimate desideratum for a belief function is that it be close to the truth" and that any other virtue of a belief function can be sacrificed if it conflicts with accuracy. In forthcoming work, Pettigrew suggests further that various other desiderata, like being appropriately responsive to one's evidence, may also be derived from accuracy, so that these other values cannot conflict with accuracy. Similarly, Joyce (2009, 275) suggests (at least tentatively) that epistemic value reflects considerations of accuracy alone, though he also suggests it may be merely a "useful fiction" (266). (Greaves and Wallace skip accuracy altogether and just talk about 'epistemic utility'.)

I will follow these authors in assuming that the inaccuracy measures include everything that is relevant. If there are other normatively significant values that could potentially overrule accuracy, then our results will all have to be interpreted quite differently (Easwaran and Fitelson 2012). But whatever support their results give to probabilism and conditionalization, my results give the same support to some additional norms for infinitary cases. The use of norms based only on accuracy is what I will call "the accuracy framework." As stated in the introduction, the point of my argument is just to illustrate the power of the accuracy framework in cases with infinitely many possibilities. It is a further question whether this power is a strength or a weakness.

3.2. Synchronic Norms on Conditional Probability. First I will see what my theorems tell us about the doxastic state of an agent at a time. I will start by assuming the following synchronic norm on plans (here, and in everything that follows, P will represent the agent's current credences, before actually performing the experiment):

> **Planning.** An agent considering an update on the result of an experiment $\mathcal{E}$ should plan to update in accord with some $\mathcal{E}$-available plan that minimizes expected inaccuracy with respect to her current credence function.

This then gives us norms on plans from our theorems. From the Ratio Theorem we get:

> **Plan Conditionalization.** For any experiment $\mathcal{E}$, and any E in $\mathcal{E}$ with $P(E) > 0$, it ought to be the case that the agent plans, in case experiment $\mathcal{E}$ gives outcome E, to update her credence in A to $P(A \wedge E)/P(E)$.

From the Interval Theorem we get:

> **Plan Reflection 1.** For any experiment $\mathcal{E}$, it ought to be the case that the agent's plans for future credence in A after updating on $\mathcal{E}$ span an interval that includes $P(A)$.

From the Expectation Theorem we get:

> **Plan Reflection 2.** For any experiment $\mathcal{E}$, it ought to be the case that the expected future credence in A given the agent's plan for updating on $\mathcal{E}$ is exactly $P(A)$.

Of course, an agent's plans for updating are not usually the sorts of things we talk about in discussing the synchronic status of an agent's credences. However, one might think that these plans have a normative connection to her current conditional credences. I propose:

> **CondProb.** When considering a prospective experiment $\mathcal{E}$, an agent's conditional credences should agree with her plan for updating on learning the outcome of the experiment.

Of course, this norm is trivial if one thinks that conditional credences *just are* one's plans for updating. For reasons discussed in Hájek (2003) and Eriksson and Hájek (2007), I suspect that update plans and conditional credence probably are not actually the same state. However, it seems plausible that there is a normative connection between the two.

Given both these norms, we can get the following additional conclusions. The Ratio Theorem gives:

> **Ratio Analysis.** If $P(E) > 0$, then $P(A|E)$ should be $P(A \wedge E)/P(E)$.

The Interval Theorem gives:

> **Conglomerability.** For any proposition A and any experiment $\mathcal{E}, P(A)$ should be within the range spanned by the $P(A|E_i)$, for E_i in $\mathcal{E}$.

The Expectation Theorem gives:

> **Disintegrability:** For any proposition A and any experiment $\mathcal{E}, P(A)$ should equal $Exp[P(A|E_i)]$.

Ratio Analysis is of course the standard account of conditional probability, and the other two are the principles known by those names in the literature (Dubins 1975). Philosophers and statisticians have given various arguments against Conglomerability and Disintegrability (Hill 1980; Hill and Lane 1985; Kadane et al. 1986; Arntzenius et al. 2004). If these challenges are taken to be decisive, then this raises problems for the accuracy framework. Defenders of the accuracy framework will have to use some of the defenses suggested by Easwaran (2008) or suggested by Dickey, Fraser, or Lindley in their responses included in Hill (1980).

3.3. Diachronic Norms on Updates. Leitgeb and Pettigrew, and Greaves and Wallace, claim to get more than just synchronic norms on credences. They both claim that over the course of an update, an agent should obey the following norm:

> **Conditionalization.** After learning a proposition B, an agent's new credence in A ought to be $P(A \wedge B)/P(B)$.

A natural diachronic norm that gives Conditionalization as a consequence of the Ratio Theorem would be:

> **Updating.** A rational agent always updates her credences in response to performing experiment $\mathcal{E}$ in accord with some plan that minimizes prior expected inaccuracy.

In light of my Interval Theorem and Expectation Theorem, we see that any update satisfying my conditions should also satisfy:

> **General Reflection 1.** For any experiment $\mathcal{E}$, a rational agent's current credence in A ought to lie in the range spanned by her possible future credences in A after updating on $\mathcal{E}$.

> **General Reflection 2.** For any experiment $\mathcal{E}$, a rational agent's current credence in A ought to be equal to her expected future credence in A after updating on $\mathcal{E}$.

Principles like this were discussed and attacked by van Fraassen (1984), and many other counterexamples have been proposed since then. But to the extent that these counterexamples involve updates that violate the assumptions I made in section 2.2, they are not challenges to my normative claims. Any counterexamples that involve only finite partitions are just as much challenges to Conditionalization as to the various reflection principles.

But regardless of the status of the counterexamples, I am somewhat suspicious of any such diachronic norm. As Michael Titelbaum pointed out in comments on Greaves's presentation of Greaves and Wallace (2006) at the 2006 Pacific meeting of the American Philosophical Association, if an agent just arbitrarily updates in a way that violates her prior plan, and we use the agent's new credence function to calculate the expected accuracy rather than the old function, then Immodesty guarantees that any update to a probabilistically coherent set of credences will end up minimizing expected inaccuracy. (Titelbaum recommended instead a synchronic interpretation of the significance of their theorem, like what I considered in sec. 3.2.) Thus, to get this conclusion, we need a real diachronic norm on credences.

Giving a diachronic norm on credence is a notoriously difficult problem. Despite the fact that a variety of arguments formally reach diachronic conclusions, Christensen (1991, 246) argues that "without some independent reason for thinking that an agent's present beliefs must cohere with her future beliefs," there can be no more support for a diachronic norm (like conditionalization) than for an interpersonal norm (e.g., the requirement that one's credences must match those of one's spouse).

Arguments based on decision theory and expected values are no better at providing reasons for intertemporal coherence. Various authors have tried to give arguments that the very fact that an agent has made a plan gives her a reason to stick with it when the appropriate time comes (McClennen 1990; Gauthier 1994; Hinchman 2003). However, there are some decision situations that are apparent counterexamples, where it seems an agent could do better by abandoning her plans (Kavka 1983).

Thus, van Fraassen (1995) moved from attacking General Reflection to defending principles more like Plan Reflection, saying that it is only one's synchronic commitments that are susceptible to these normative requirements, rather than one's actual diachronic updates.

Another perspective on reflection principles is provided by Skyrms (1990), who suggests that they should be taken as a way to classify certain updates as instances of "learning," even though they do not hold for all updates. (See the discussion on pages 97 and 99 of Skyrms [1990] of "Condition M," which is equivalent to General Reflection 2.) It is interesting that Myrvold (2012) has shown that, for an accuracy measure that satisfies Immodesty, any update that satisfies Condition M must decrease expected inaccuracy. My results can be seen as a sort of converse—an update that satisfies my conditions and minimizes expected inaccuracy must satisfy Condition M. The main difference is that they use Condition M to characterize the updates in question, while I show that it holds for updates that satisfy the conditions given in section 2.2.

3.4. Upshot. Thus, using the accuracy framework, we get some support for diachronic reflection principles and synchronic principles of Conglomerability and Disintegrability, in addition to the support for conditionalization given by Greaves and Wallace, and Leitgeb and Pettigrew. Although there are worries about the diachronic norms, the synchronic norms already provide interesting constraints on conditional probability, which may still be of use in understanding diachronic updates. I have shown that these results hold even without the idealizing assumption that all partitions are finite (in which case all the results follow from Ratio Analysis and Conditionalization).

Appendix
Theorems

Ratio Theorem. If A is a measurable proposition, $\mathcal{E}$ is an experiment, E is an element of $\mathcal{E}$ with positive probability, and f_A has minimal expected inaccuracy among the $\mathcal{E}$-available plans for updating one's credence in A, then $f_A(s) = P(A \wedge E)/P(E)$ for any s in E.

Proof. Let f' be a function that agrees with f_A everywhere outside E, while $f'(s) = P(A \wedge E)/P(E)$ for all s in E. Note that f' is also an available plan—it is still constant on all elements of $\mathcal{E}$ (since f_A was), and since it differs from f_A on only a single measurable set, we can see that the set of s with $f'(s) > x$ is still measurable, since it is either the set of s with $f_A(s) > x$, or the union of this set with E, or the intersection of this set with the complement of E. To prove the theorem, I will show that if $f_A \neq f'$, then the expected inaccuracy of f' is less than the expected inaccuracy of f. Thus, if f_A is an available plan that minimizes expected inaccuracy, then it must equal f', as required.

Now let $V(s) = I(A, s, f_A(s))$ and $V'(s) = I(A, s, f'(s))$. By the linearity of expectation, we see that $Exp(V) = Exp(V \cdot 1_E) + Exp(V \cdot 1_{\neg E})$, where 1_E is the variable that takes values 1 when E is true and 0 otherwise, and $1_{\neg E}$ is the opposite. Similarly, $Exp(V') = Exp(V' \cdot 1_E) + Exp(V' \cdot 1_{\neg E})$. Because V and V' are identical on $\neg E$, we see that $V \cdot 1_{\neg E} = V' \cdot 1_{\neg E}$. Thus, $Exp(V') \leq Exp(V)$ (which is what we need to show) if and only if $Exp(V' \cdot 1_E) \leq Exp(V \cdot 1_E)$.

Now let x be the value that f_A takes on E, and let $x' = P(A \wedge E)/P(E)$ be the value that f' takes on E. Now, on E, $V(s) = I(A, s, x)$ and $V'(s) = I(A, s, x')$. By linearity again, $Exp(V \cdot 1_E) = Exp(V \cdot 1_E \cdot 1_A) + Exp(V \cdot 1_E \cdot 1_{\neg A})$, and $Exp(V' \cdot 1_E) = Exp(V' \cdot 1_E \cdot 1_A) + Exp(V' \cdot 1_E \cdot 1_{\neg A})$. But for any s in $E \wedge A$, Extensionality gives us $V(s) = I(A, 1, x)$ and $V'(s) = I(A, 1, x')$, and, similarly, for any s in $E \wedge \neg A$, $V(s) = I(A, 0, x)$ and $V'(s) = I(A, 0, x')$. Thus, since V and V' are constant on these sets, axioms 2 and 4 of expectation give us that

$$Exp(V \cdot 1_E \cdot 1_A) = P(E \wedge A)I(A, 1, x),$$

$$Exp(V \cdot 1_E \cdot 1_{\neg A}) = P(E \wedge \neg A)I(A, 0, x),$$

$$Exp(V' \cdot 1_E \cdot 1_A) = P(E \wedge A)I(A, 1, x'),$$

$$Exp(V' \cdot 1_E \cdot 1_{\neg A}) = P(E \wedge \neg A)I(A, 0, x').$$

Since $P(E \wedge A) = x'P(E)$ and $P(E \wedge \neg A) = (1 - x')P(E)$ (the former by definition, and the latter because P is a probability function), we have:

$$Exp(V \cdot 1_E) = x'P(E)I(A, 1, x) + (1 - x')P(E)I(A, 0, x),$$

$$Exp(V' \cdot 1_E) = x'P(E)I(A, 1, x') + (1 - x')P(E)I(A, 0, x').$$

Dividing through by $P(E)$ (which is positive, by assumption), we see that Immodesty then says $Exp(V' \cdot 1_E) \leq Exp(V \cdot 1_E)$, so $Exp(V') \leq Exp(V)$, with equality holding if and only if $x = x'$. Thus, if V has minimal expectation, so that f_A has minimal expected inaccuracy, then $x = x'$ and $f_A = f'$. QED

> **Interval Theorem.** If A is a measurable proposition, $\mathcal{E}$ is an experiment, and f_A has minimal expected inaccuracy among the $\mathcal{E}$-available plans for updating one's credence in A, then there is no x with $P(A) < x$ and $x < f_A(s)$ for all s, and there is no x with $P(A) > x$ and $x > f_A(s)$ for all s. That is, $inf_{s \in \mathcal{S}} f_A(s) \leq P(A) \leq \sup_{s \in \mathcal{S}} f_A(s)$.

Proof. This proof comes in two parts. First I will show that this is true when $\mathcal{E}$ is a finite partition, and then I will use this fact to show that it is true in general.

When $\mathcal{E}$ is finite, we can write it as $\{E_1, \ldots, E_n\}$. Since f_A is constant on each E_i, I will slightly abuse the notation and write $f_A(E_i)$ rather than $f_A(s)$ for s in E_i. By the previous result, if $P(E_i) > 0$, then $f_A(E_i) = P(A \wedge E_i) / P(E_i)$, so $P(A \wedge E_i) = f_A(E_i)P(E_i)$. Because P is a probability function, if $P(E_i) = 0$, then $P(A \wedge E_i) = 0 = f_A(E_i)P(E_i)$. Thus, in any case, we have $P(A \wedge E_i) = f_A(E_i)P(E_i)$. Since P is a probability function and $\mathcal{E}$ is a finite partition, we see that $P(A) = P(A \wedge E_1) + \ldots + P(A \wedge E_n) = f_A(E_1)P(E_1) + \ldots + f_A(E_n)P(E_n)$. Since the $P(E_i)$ add up to 1, we see that $P(A)$ is a weighted average of the $f_A(E_i)$, and thus at least one of them must be no larger than $P(A)$, and at least one of them must be no smaller than $P(A)$, which entails what we want to prove. So this theorem (as well as the Expectation Theorem) is proved when $\mathcal{E}$ is finite.

One observation that will be useful—the proof establishes in fact that there are E_i and $E_{i'}$ with $f_A(E_i) \leq P(A) \leq f_A(E_{i'})$, where $P(E_i)$ and $P(E_{i'})$ are both positive.

I will define $I(c, x) = cI(A, 1, x) + (1 - c)I(A, 0, x)$ to be the expected inaccuracy of having credence x in A, when evaluated by an agent whose credence in A is c. (Conveniently, although I have earlier used $I(1, x)$ as an abbreviation for $I(A, 1, x)$ and $I(0, x)$ as an abbreviation for $I(A, 0, x)$, this definition agrees with those abbreviations when $c = 1$ or $c = 0$.) Immodesty is the claim that $I(c, c) \leq I(c, x)$, with equality holding if and only if $x = c$.

I will also define $\Delta(c, x) = I(c, x) - I(c, c)$, which is the amount that the expected inaccuracy of credence x exceeds the expected inaccuracy of c itself, again when evaluated by an agent whose credence in A is c.

I will start the proof for the infinite case by using Convexity (together with implicit appeal to Monotonicity and Immodesty) to prove two lemmas:

Lemma 1. If $c < x < x'$, then $\Delta(c, x) < \Delta(c, x')$.

Proof. Let $\lambda = (x - c)/(x' - c)$ measure the proportion to which x is between c and x', so that $x = \lambda x' + (1 - \lambda)c$. By Convexity, we know that $I(1, x) - I(1, c) < \lambda[I(1, x') - I(1, c)]$ and that $I(0, x) - I(0, c) < \lambda[I(0, x') - I(0, c)]$. Multiplying the first inequality by c and the second by $(1 - c)$, and adding them together, tells us that

$$c[I(1, x) - I(1, c)] + (1 - c)[I(0, x) - I(0, c)] \\ < \lambda\{c[I(1, x') - I(1, c)] + (1 - c)[I(0, x') - I(0, c)]\}.$$

By substituting definitions, this means that

$$\Delta(c, x) < \lambda\Delta(c, x').$$

Since all of these are positive numbers, and $0 < \lambda < 1$, we see that $\Delta(c, x) < \Delta(c, x')$. QED

Lemma 2. If $c' < c < x$, then $\Delta(c, x) < \Delta(c', x)$.

Proof. Let $\lambda = (x - c)/(x - c')$ measure the proportion to which c is between x and c', so that $c = \lambda c' + (1 - \lambda)x$. By Convexity, we know that $I(1, x) - I(1, c) < \lambda[I(1, x) - I(1, c')]$, with both sides negative. Thus, if we multiply the first by a larger positive number than the second, then we preserve the inequality, so we see that $c[I(1, x) - I(1, c)] < \lambda c'[I(1, x) - I(1, c')]$. By Convexity, we know that $I(0, x) - I(0, c) < \lambda[I(0, x) - I(0, c')]$, with both sides positive. Thus, if we multiply the first by a smaller positive number than the second, then we preserve the inequality, so we see that $(1 - c)[I(0, x) - I(0, c)] < \lambda(1 - c')[I(0, x) - I(0, c')]$. Adding these two resulting inequalities together tells us that

$$c[I(1, x) - I(1, c)] + (1 - c)[I(0, x) - I(0, c)] \\ < \lambda\{c'[I(1, x) - I(1, c')] + (1 - c')[I(0, x) - I(0, c')]\}.$$

By substituting definitions, this means that

$$\Delta(c, x) < \lambda\Delta(c', x).$$

Since all of these are positive numbers, and $0 < \lambda < 1$, we see that $\Delta(c, x) < (c', x)$. QED

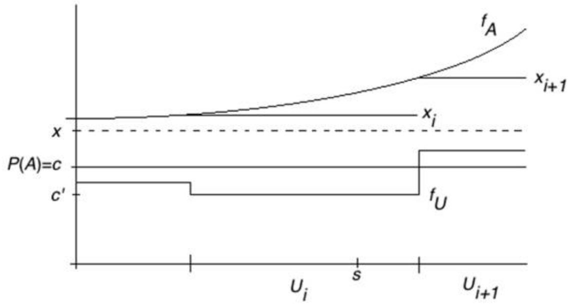

Figure A1.

I will now return to the proof when $\mathcal{E}$ is infinite. This part of the proof will proceed by contradiction. Figure A1 will help keep track of the relevant values that are introduced.

So assume that there is $x > P(A)$ such that $f_A(s) > x$ for all s. (The dual case with $x < P(A)$ such that $f_A(s) < x$ for all s works similarly.) Let $c = P(A)$. Let $\varepsilon = \Delta(c, x)$. By Immodesty, ε is positive.

By Bounded Continuity, there are finitely many values $0 = x_0 < x_1 < \ldots < x_n = 1$ such that $I(A, 0, x_{i+1}) - I(A, 0, x_i) < \varepsilon$ and $I(A, 1, x_i) - I(A, 1, x_{i+1}) < \varepsilon$. Let U_i be the set of s such that $x_i \leq f_A(s) < x_{i+1}$. (By refinement of the partition if necessary, we can ensure that any nonempty U_i will have $x < x_i$, since $x < f_A(s)$ for all s.) Because f_A is constant on each element of $\mathcal{E}$, each U_i must be a union of members of $\mathcal{E}$. If $\mathcal{U}$ is the partition into the U_i, then this means that $\mathcal{U}$ is a coarsening of the partition $\mathcal{E}$.

Let f_U be a $\mathcal{U}$-available plan with minimal expected inaccuracy. Since $\mathcal{U}$ is a finite partition, the finite part of this proof establishes that there is an s with $f_U(s) \leq c = P(A)$. Let U_i be the element of $\mathcal{U}$ containing s, and let $c' = f_U(s)$, so that $c' \leq c < x < x_i \leq f_A(s)$ for all s in U_i. By the observation made after proving the finite version, we can ensure that $P(U_i) > 0$. Thus, by the Ratio Analysis result, we see that $c' = P(A \wedge U_i)/P(U_i)$. Let f'_U be the function that agrees with f_A outside of U_i and agrees with f_U inside U_i.

Now consider the set F_i of all $\mathcal{E}$-available functions that agree with f_A outside U_i. Note that f_A and f'_U are both in F_i. (The term f_A is not $\mathcal{U}$-available, but it is $\mathcal{E}$-available.) I will also consider the function f_i, which agrees with f_A outside of U_i and takes on the value x_i for all s in U_i. By linearity of expectations,

$$Exp[I(A, s, f(s))] = Exp[I(A, s, f(s)) \cdot 1_{U_i}] + Exp[I(A, s, f(s)) \cdot 1_{\neg U_i}].$$

Because all of these functions are identical for s in $\neg U_i$, their total expected inaccuracy will depend only on $Exp[I(A, s, f(s)) \cdot 1_{U_i}]$.

Because $f'_U(s) = c'$ and $f_i(s) = x_i$ for all s in U_i, we can see that

$$Exp[I(A, s, f'_U(s)) \cdot 1_{U_i}] = P(U_i \wedge A)I(A, 1, c') + P(U_i \wedge \neg A)I(A, 0, c')$$

and

$$Exp[I(A, s, f_i(s)) \cdot 1_{U_i}] = P(U_i \wedge A)I(A, 1, x_i) + P(U_i \wedge \neg A)I(A, 0, x_i).$$

But we observed earlier that $P(U_i \wedge A)/P(U_i) = c'$, and thus since P is a probability function, $P(U_i \wedge \neg A)/P(U_i) = (1 - c')$, so we have

$$\frac{Exp[I(A, s, f'_U(s)) \cdot 1_{U_i}]}{P(U_i)} = c'I(A, 1, c') + (1 - c')I(A, 0, c') = I(c', c')$$

and

$$\frac{Exp[I(A, s, f_i(s)) \cdot 1_{U_i}]}{P(U_i)} = c'I(A, 1, x_i) + (1 - c')I(A, 0, x_i) = I(c', x_i).$$

Since $c' \leq c < x < x_i$, our two lemmas together tell us that $I(c', x_i) - I(c', c') = \Delta(c', x_i) > \Delta(c, x) = \varepsilon$. Thus, the expected inaccuracy of f_i is at least $\varepsilon P(U_i)$ greater than the expected inaccuracy of f'_U.

But for any s in U_i, the inaccuracy of f_A on s is within ε of the inaccuracy of f_i on s (because f_A is always between x_i and x_{i+1}, which are chosen to make sure that regardless of whether A is true or false, the inaccuracy of any value in this range is within ε of the inaccuracy of any other value in this range—this is the one use of Monotonicity). Thus, the expected inaccuracy of f_A is within $\varepsilon P(U_i)$ of the expected inaccuracy of f_i, which means that the expected inaccuracy of f_A must be greater than the expected inaccuracy of f_U.

Thus, since f_A and f'_U agree outside of U_i, we see that the expected inaccuracy of f_A is greater than the expected inaccuracy of f'_U, which contradicts the assumption that f_A had minimal expected inaccuracy of any $\mathcal{E}$-available plan. Thus, the original assumption that there is x with $P(A) < x$ and $f_A(s) > x$ for all s must be false, which proves one direction of the inequality. The other direction is proven by the corresponding argument with all inequalities reversed. QED

Remark: if f'_A differs from f_A only on a set of probability 0, then the expected inaccuracy of both plans is the same, so that one has minimal expected inaccuracy if and only if the other does. Therefore, as a corollary, we have that for any $x > P(A)$, the set of s with $f_A(s) \leq x$ must have nonzero probability—otherwise we could find an f'_A that also has minimal expected inaccuracy but violates the conclusion of the theorem.

Expectation Theorem. If for all propositions A in $\mathcal{F}$ the conditions of the Interval Theorem hold, and for each s, $f_{\ldots}(s)$ is a probability function (so that $f_A(s) \geq 0$ for all A, $f_S(s) = 1$, and $f_{A \cup B}(s) = f_A(s) + f_B(s)$ whenever A and B are disjoint), then for all A, $P(A) = Exp[f_A(s)]$.

Proof. This is a straightforward consequence of the Interval Theorem together with theorem 1 of Dubins (1975), which states that an expectation is conglomerable if and only if it is disintegrable. The Interval Theorem shows that the expectation here is conglomerable, and the conclusion of this theorem is the claim that the expectation is disintegrable. QED

REFERENCES

Arntzenius, F. 2003. “Some Problems for Conditionalization and Reflection.” *Journal of Philosophy* 100 (7): 356–70.

Arntzenius, F., A. Elga, and J. Hawthorne. 2004. “Bayesianism, Infinite Decisions, and Binding.” *Mind* 113 (450): 251–83.

Bradley, D. 2012. “Four Problems about Self-Locating Belief.” *Philosophical Review* 121 (2): 149–77.

Christensen, D. 1991. “Clever Bookies and Coherent Beliefs.” *Philosophical Review* 100 (2): 229–47.

Dubins, L. 1975. “Finitely Additive Conditional Probabilities, Conglomerability and Disintegrations.” *Annals of Probability* 3 (1): 89–99.

Earman, J. 1992. *Bayes or Bust?* Cambridge, MA: MIT Press.

Easwaran, K. 2008. “The Foundations of Conditional Probability.” PhD thesis, University of California, Berkeley.

Easwaran, K., and B. Fitelson. 2012. “An ‘Evidentialist’ Worry for Joyce’s Argument for Probabilism.” *Dialectica* 66 (3): 425–33.

Eriksson, L., and A. Hájek. 2007. “What Are Degrees of Belief?” *Studia Logica* 86:185–215.

Gauthier, D. 1994. “Assure and Threaten.” *Ethics* 104 (4): 690–721.

Greaves, H., and D. Wallace. 2006. “Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.” *Mind* 115 (459): 607–32.

Hájek, A. 2003. “What Conditional Probability Could Not Be.” *Synthese* 137:273–323.

Hill, B. M. 1980. “On Some Statistical Paradoxes and Non-disintegrability.” *Trabajos de estadística y de investigación operativa* 31 (1): 39–66.

Hill, B. M., and D. Lane. 1985. “Conglomerability and Countable Additivity.” *Sankhyā: Indian Journal of Statistics* A 47 (3): 366–79.

Hinchman, E. 2003. “Trust and Diachronic Agency.” *Noûs* 37 (1): 25–51.

Joyce, J. 1998. “A Nonpragmatic Vindication of Probabilism.” *Philosophy of Science* 65 (4): 575–603.

———. 2009. “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief.” In *Degrees of Belief*, ed. F. Huber and C. Schmidt-Petri, 263–97. Berlin: Springer.

Kadane, J. B., M. J. Schervish, and T. Seidenfeld. 1986. “Statistical Implications of Finitely Additive Probability.” In *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, ed. P. K. Goel and A. Zellner, 59–76. Amsterdam: Elsevier.

Kavka, G. 1983. “The Toxin Puzzle.” *Analysis* 43:33–36.

Leitgeb, H., and R. Pettigrew. 2010a. “An Objective Justification of Bayesianism I: Measuring Inaccuracy.” *Philosophy of Science* 77 (2): 201–35.

———. 2010b. “An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy.” *Philosophy of Science* 77 (2): 236–72.

Lindley, D. 1982. "Scoring Rules and the Inevitability of Probability." *International Statistical Review* 50 (1): 1–11.
McClennen, E. 1990. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge: Cambridge University Press.
Meacham, C. 2008. "Sleeping Beauty and the Dynamics of De Se Beliefs." *Philosophical Studies* 138 (2): 245–69.
Myrvold, W. 2012. "Epistemic Values and the Value of Learning." *Synthese* 187 (2): 547–68.
Skyrms, B. 1990. *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.
Titelbaum, M. 2008. "The Relevance of Self-Locating Beliefs." *Philosophical Review* 117 (4): 555–606.
van Fraassen, B. 1981. "A Problem for Relative Information Minimizers in Probability Kinematics." *British Journal for the Philosophy of Science* 32 (4): 375–79.
———. 1984. "Belief and the Will." *Journal of Philosophy* 81 (5): 235–56.
———. 1995. "Belief and the Problem of Ulysses and the Sirens." *Philosophical Studies* 77:7–37.
Whittle, P. 2000. *Probability via Expectation*. 4th ed. Berlin: Springer.