Generalized Probabilism: dutch books and accuracy domination

J. Robert G. Williams

Last revised: 3rd March 2011

Abstract

This paper explores De Finetti's generalized versions of Dutch Book and Accuracy Domination theorems. Following proposals due to Jeff Paris, we construe these as underpinning a *generalized* probabilism appropriate to belief states against either a classical or a non-classical background. Both results are straightforward corollaries of the separating hyperplane theorem; their geometrical relationship is examined. It is shown that each point of Accuracy Domination for b induces a Dutch Book on b; but Dutch Books may need to be 'scaled' in order to find a point of Accuracy-Domination. Finally, diachronic Dutch Book defences of conditionalization are examined in the general setting. The formulation and limitations of the generalized conditionalization this delivers are examined.

Contents

Probabilism is the thesis that ideal belief states should be structured probabilistically. Familiar arguments for probabilism presuppose a background classical logic and semantics. *Generalized probabilism* is not subject to these limitations; it claims that even in non-classical settings ideal belief states should be structured via a non-classical analogue of probability.

Section 1 briefly reviews relevant literature: Bruno De Finetti's very general Dutch Book and Domination arguments, and Jeff Paris's interpretation of the former as an argument for generalized probabilism. In sections 2-4 De Finetti's arguments are presented; sections 5-6 draw out the geometrical relationships between the Dutch Books and Accuracy Domination. Sections 7-9 examine an extension of these results. We give a geometrical characterization of the constraints on updating degrees of belief provided by Lewis-style 'diachronic Dutch Books'. In the classical case, the non-dutchbookable update strategy is conditionalization; we examine the characterization of non-dutchbookable updates in the generalized setting.

1 Generalized probabilism

Dutch Book theorems have long been used to argue for probabilism: the view that ideal belief states should be structured probabilistically. More recently, Jim Joyce (1998, 2009) has argued for probabilism on the basis of 'Accuracy Domination'—considerations that are supposed to turn simply on rational agents' aspiration to have beliefs that are close to the truth, rather than their love of money. The Dutch Book argument for probabilism contends that having non-probabilistic degrees of belief commits one to regarding as fair a sure-loss book of bets. The Accuracy Domination argument contends that if one has non-probabilistic degrees of belief, there will be a rival probabilistic belief-state that is *guaranteed* to be 'closer to the truth' than one's own.

Versions of the theorems that underpin each argument are presented by De Finetti (1974, pp.87-90). De Finetti's theorems are interestingly general: they are formulated in terms of expectations ('previsions') over an arbitrary set of random variables over the reals. They show that something analogous to 'Dutch Booking' or 'domination' will occur unless the expected values of the variables meet certain formal constraints, allowing us to derive standard constraints on expectation values. Once we have this general result, as a special case one can consider random variables that correspond to classical propositions—each one taking value 1 if its associated proposition is true, and 0 if it is false. We identify degrees of belief in a proposition with the expectation of the associated random variable or 'prevision of truth value'. In the general case, De Finetti's theorems show that a set of 'coherent' (non-dutchbookable/non-dominated) expectations must be a convex combination of the set of values taken by the random variables. In the special case, this means that 'coherent' (non-dutchbookable/non-dominated) degrees of belief must be convex combinations of classical truth value distributions. It turns out that the convex combinations of classical truth values are exactly the classical probabilities; so the special case provides the resources for the arguments for (classical) probabilism mentioned above.[1]

Paris (2001) argues that a more generalized probabilism is in the offing. Paris uses De Finetti's general Dutch Book result to study coherent degrees of belief in *non-classical* settings. Rather than expectations of classical truth values (random variables taking values 1 and 0, classically distributed) he studies *nonclassical* truth value distributions and the associated degrees of belief (again construed as previsions of truth value). More specifically, on minimal assumptions about the nature of the truth-value distributions, we will be able to show that non-dutch-bookable belief states must satisfy the following generalizations of the familiar axioms of probability (Choquet, 1953):

[1] Unlike Joyce, De Finetti does not interpret the domination results in terms of a norm of gradational accuracy on credence, but instead uses it to construct a second sure-loss argument for probabilism. See De Finetti op cit.

$$
\begin{array}{llll}
(\mathcal{L}1) & \models A & \Rightarrow & b(A)=1 \\
 & A\models & \Rightarrow & b(A)=0 \\
(\mathcal{L}2) & A\models B & \Rightarrow & b(A)\leq b(B) \\
(\mathcal{L}3) & \multicolumn{3}{l}{b(A\wedge B)+b(A\vee B)=b(A)+b(B)}
\end{array}
$$

The assumptions required to give this result are that (i) the propositions take truth values from $\{1,0\}$; (ii) the consequence relation $\models$ is given by '1-preservation'; i.e. $A\models B$ iff on every truth value distribution, if the truth value of A is 1, the truth value of B is 1; (iii) the distribution of truth values satisfy the following:

$$
\begin{array}{llll}
(\mathcal{T}2) & V(A)=1\wedge V(B)=1 & \iff & V(A\wedge B)=1 \\
(\mathcal{T}3) & V(A)=0\wedge V(B)=0 & \iff & V(A\vee B)=0.
\end{array}
$$

This lovely result shows that (for a large range of cases) coherent belief states against a non-classical backdrop are subject to *the same* local constraints as in the classical case, so long as these are suitably formulated to make their dependence on the background logic explicit.

However, the results so far still fail to cover some interesting non-classical settings. For example, there are comparatively natural examples where $(\mathcal{T}3)$ fails; and in 'fuzzy' semantics the range of truth values may fall within $[0,1]$ rather than just within $\{0,1\}$. This motivates a programme finding axiomatizations of coherent belief states/convex combinations of truth values are available in those cases falling outside Paris's characterization. Here I will mention just one of the examples given by Paris (crediting Jaffray (1989)) that will be of special interest to us later.

The space of 'value assignments' on which we focus can be thought of as 'partial' assignments of truth and falsity. More specifically, for each assignment V there is some proposition B such that $V(A)=1$ iff $B\models A$; otherwise $V(A)=0$. Note that when neither A nor its negation are classically entailed by B, both get value 0. The interpretation of these values Paris suggests is epistemic—the value 1 being read as 'known true' and 0 as 'known false': a situation in which A and $\neg A$ both taking value 0 is one where there is a gap in your evidence. An alternative, on-epistemic interpretation of such truth-value distributions is provided by the widely-discussed *supervaluational* semantics (cf. van Fraassen, 1966; Fine, 1975; Keefe, 2000; Field, 2000), where such cases are interpreted as *truth value gaps*.[2]

Shafer (1976) and Jaffray (1989) tell us exactly what the non-dutchbookable belief states look like against this backdrop: they are the *Dempster-Shafer belief functions*, axiomatized thus:

$$
\begin{array}{llll}
(\mathcal{DS}1) & \models A & \Rightarrow & b(A)=1, b(\neg A)=0 \\
(\mathcal{DS}2) & \models (A\leftrightarrow B) & \Rightarrow & b(A)=b(B) \\
(\mathcal{DS}3) & \multicolumn{3}{l}{b(\bigvee_{i=1}^{m}A_i)\geq \sum_S(-1)^{|S|-1}b(\bigwedge_{i\in S}A_i)}
\end{array}
$$

(where S ranges over non-empty subset of $\{1,\ldots,m\}$).

So the project of generalized probabilism may be extended to provide a potential underpinning for a well-known rival to classical probabilism itself.

In discussing generalized probabilism, we followed Paris in focusing on Dutch Book arguments. But one can equally study coherent non-classical degrees of belief using De Finetti's general domination result; and non-domination delivers the same set of generalized probabilities as non-dutch-bookability, if one works with the same input set of non-classical truth-value distributions.[3] So not only the familiar pragmatic arguments for probabilism, but also the more recent alethic arguments that Joyce has proposed can be put at the service

[2]This is further explored in (Author).

[3]This is the approach taken in (Author), where a generalized Accuracy-Domination theorem is interpreted in the Joycean manner as a 'de-pragmatised' argument for generalized probabilism. (Some generality is added in that presentation. There are various potential ways of measuring the key notion of 'distance from the truth values', but De Finetti's result are formulated with one particular measure—the so-called Brier score. Joyce instead argues for a set of axioms providing constraints on what can count as an 'accuracy measure', and proves a domination result for any measure satisfying the constraint; (Author) shows Joyce's strategy generalizes to the non-classical settings.)

of generalized probabilism.

That Dutch Books and Accuracy Domination results both hold over such a wide range of settings is extremely striking. The project in what follows is to examine the nature of the connection between them, expanding upon De Finetti's brief proofs of the two results. We will look into the details of the construction to extract information about how specific Dutch Books and 'dominating belief states' relate to one another. Finally, we push forward the project of generalized probabilism by providing an additional result in the spirit of De Finetti: a generalized diachronic Dutch Book theorem.

2 The key lemmas

We model an arbitrary belief state as a function from a set of N propositions to real numbers—the number assigned to the proposition θ representing the degree of belief that it is true. Belief states, so construed, will be identified with vectors in an N-dimensional space, equipped with the Euclidean inner product and (hence) a notion of the 'nearness' or 'distance' between two arbitrary belief states. For now, we can think of this as a purely abstract notion—but in the arguments below, it will receive a definite interpretation. Within this space of belief states, we have a convex set C containing a subset W. We'll call C the 'coherent belief states' or 'probabilities', and W the 'truth value distributions' or 'worlds'. And indeed, on a standard interpretation of W and C, W will be vectors that represent an exhaustive, bivalent, truth value distribution (of 1's and 0's) over the propositions; and C will be the corresponding set of classical probabilities (which, recall, can be identified with the convex combinations of elements of W). But the results below will not require this exact interpretation, which allows us to apply to the theorem with W non-classical truth value distributions and C the convex combinations thereof. Other readings are also possible, though we will assume throughout that W is finite.

The first key result is this. Given a belief state b not in C, there is a state $c \in C$ which is at least as close to b as any other state in C (this follows straightforwardly from the closure of C). 'Nearness' here is understood via the standard Euclidean inner-product on the vectors—so what we're saying is that $\forall x \in C$, $||b-x|| \geq ||b-c||$. [4]

The second key lemma can be expressed in several ways. But the underlying idea is this. Take the vector s that goes from c to b ($s = b-c$). And consider a vector t that goes from c to some arbitrary belief state a in C ($t = a-c$). The lemma says that the angle between s and t is not acute. In vector terminology, $s \cdot t \leq 0$. We can equally express the result as: $s \cdot (a-c) \leq 0$ or $s \cdot a \leq s \cdot c$. The situation is easy to picture geometrically: see figure 1.

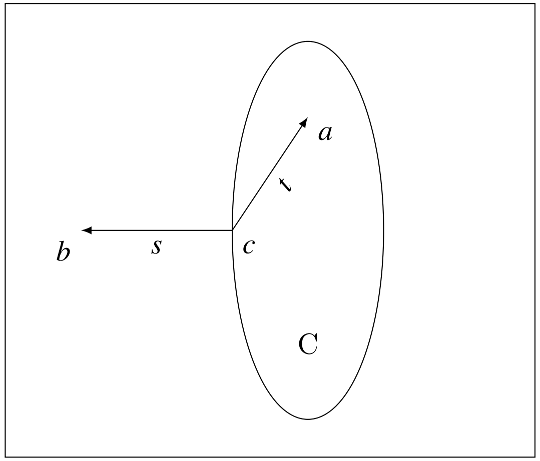

Figure 1: The key lemma: the angle between s and t is non-acute.

[4] Let X be the intersection of C with the set of vectors x such that $||b-x|| < r$, for some r large enough to make X non-empty. Note that X is closed and bounded. The function $f(x) = ||x-b||$ is continuous over the closed bounded set X, so it achieves a minimum by the extreme value theorem. This minimum is the required c. See Border (manuscript 2009, p.5) for a proof of this result that works for general Hilbert spaces, rather than just the Euclidean N-dimensional space we're working with here.

We can prove this algebraically.[5] To do so, consider an arbitrary point $x = c + \lambda(a - c)$ on the line segment between a and c. By convexity, $x \in C$. By choice of c, we have $||b - x|| \geq ||b - c||$. Hence $(b - x)^2 \geq (b - c)^2$. We write:

$$
\begin{array}{rrll}
 & 0 & \geq & ||b-c||^2 - ||b-x||^2 \\
 & & = & ||b-c||^2 - ||b-c-\lambda(a-c)||^2 & \text{construction of } x \\
 & & = & ||b-c||^2 - [(b-c-\lambda(a-c))\cdot(b-c-\lambda(a-c))] \\
 & & = & ||b-c||^2 - [||b-c||^2 - 2\lambda(b-c)\cdot(a-c) + \lambda^2||a-c||^2] \\
 & & = & 2\lambda(b-c)\cdot(a-c) - \lambda^2(a-c)^2 \\
\therefore & 0 & \geq & (b-c)\cdot(a-c) - \frac{1}{2}\lambda(a-c)^2 & \text{dividing by } 2\lambda \\
\therefore & 0 & \geq & (b-c)\cdot(a-c) & \text{letting } \lambda \to 0 \\
\therefore & 0 & \geq & s \cdot t
\end{array}
$$

For illustration, let us consider the situation in two dimensions, to see how the geometry (at least in that special case) confirms the algebraic derivation above (see figure 2). What regions can C occupy, seen from b? Well, we constructed c as a maximally closest point in C from b. So if we draw a circle around b, of a radius equal to the distance between c and b, none of C falls within that circle. Now draw the tangent to the circle at c (the dashed line in figure 2). Can any portion of C fall into region 2, the part to the left side of the tangent? It cannot. For if there's any point in C in region 2—call it x—the line from c to x will intersect the circle. The convexity of C assures us that every point on the line cx, including those within the circle, are elements of C. But this now contradicts the earlier constraint, that no part of C is within the circle. By reductio, every point in C must be in region 1. Since C is confined to region 1, the angle formed between s and the line between c and any other point in C, will be at worst right-angled, if not obtuse. Hence the result.

[5]The proof below follows Border (manuscript 2009, p.6). It forms the core of the well-known 'separating hyperplane theorem' (due to Minkowski).

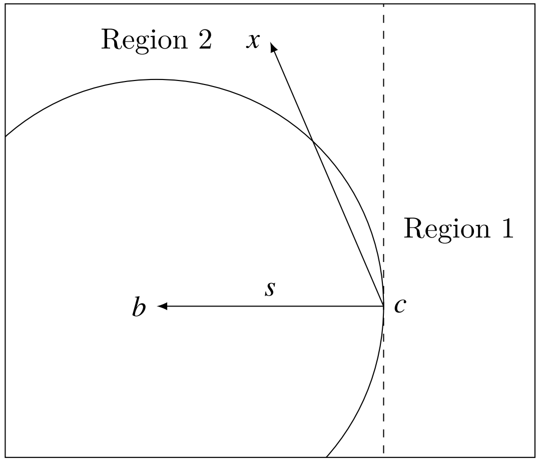

Figure 2: A geometrical version of the argument, in two dimensions. C is confined to region 1.

3 Generalized Dutch Books

With the key lemmas in hand, let's prove the Dutch Book result. The theorem (shorn of philosophical interpretation) is simply the following. Let $s = b - c$ (i.e. the difference between the belief state b and the nearest point to b in C). Then, for all $w \in W$, $(s \cdot w) < (s \cdot b)$. To prove this, recall that from the lemmas above, we already have that for every a in C, $(s \cdot a) \leq (s \cdot c)$. Since $W \subseteq C$, we have $(s \cdot w) \leq (s \cdot c)$. It will suffice to show, therefore, that $(s \cdot c) < (s \cdot b)$. But this is straightforward: consider the vector $s = b - c$ once more. It has positive length (else $b = c$, which contradicts the assumption that b is outside C). So we know that $||s|| > 0$. But this means that $0 < (s \cdot s) = (s \cdot (b - c))$; hence $0 < (s \cdot b) - (s \cdot c)$. Rearranging, this gives us the result.

Why would this be called the 'Dutch Book' theorem? To see this, we need to consider how we might represent the returns of a book of bets, and the price of that bet, within our space. Our assumption will be that if we bet on a particular proposition θ_i with a prize of z_i dollars/utils, then we will gain z_i if the proposition is true, and 0 otherwise (take the term 'prize' with a pinch of salt: prizes in our sense can be negative). Or more generally, we will obtain at world $w \in W$ an amount equal to the prize multiplied by the truth value of w in W (this characterization, note, covers non-classical as well as classical truth values). A complete book of bets, one for each proposition, is thus given by a vector $z = (z_i)$, which will live within our vector space. The overall prize of a book of bets z at world w is given by multiplying each z_i by the truth value of θ_i at w and summing. But the process just described is simply that of taking the Euclidean inner product between z and w. Hence, if z describes a book of bets, $z \cdot w$ gives its return at w.

As well as the returns of a book, we need to consider how much we pay to get the chance of that prize—the price of a book of bets. As is standard in Dutch Book arguments, we assume that the fair price for an individual bet with unit prize for an individual with belief state b is specified by the degree of belief that b assigns to the proposition bet upon; and that the fair price for a book of bets is given by summing the fair prices of the individual bets. For variable prizes, there is a parallel assumption: the fair price for an individual bet on θ_i with prize z_i is given by the degree of belief in θ_i multiplied by z_i; and the fair price of a book of bets with prize z is given by taking the individual fair prices and summing. (Of course, this is an ideal point for those sceptical of the philosophical significance of Dutch Book arguments to resist; but since we are interested not in whether the argument works, but with its geometrical structure, I will not pursue this further). Again, we note that this has a geometrical expression: the fair price relative to belief state b of a book of bets z is $b \cdot z$.

A book of bets with be a 'sure loss' for the buyer if the price paid exceeds the returns in every world. And a Dutch Book is one which is a sure loss even when bought at a fair price. Let's sum up: any vector z can represent a book of bets, where $z \cdot w$ represents the prize you receive at w. The fair price for this, at belief state b, is $z \cdot b$. A Dutch Book (for belief state p) is one that represents a sure-loss when bought at a fair-price: that is, $z \cdot b > z \cdot w$, for all $w \in W$. But now we see the significance of the earlier theorem: for it showed that whenever the belief state b is outside C, we can find a vector s (which we will interpret as a book of bets) that meets the above conditions: we can find, that is, a Dutch Book against b. (We can also show quite simply that there's no Dutch Book against credences in C, for the special case where C is the set of (generalized) probabilities, but that's another story—see Paris (2001)).

4 Generalized Accuracy domination

A belief state is Accuracy-Dominated (by belief state c, say) if, no matter which world is actual, c is 'closer to the truth' or 'more accurate' than the original belief state.[6] The obvious first question is: how do we measure 'closeness to the truth'? In the current setting, there's an obvious candidate: we measure the distance between the belief state and the relevant truth values within our Euclidean space. In the literature on accuracy measures, this Euclidean distance measure is known as the 'Brier Score'. It is a substantive philosophical assumption that it describes the relevant notion of 'distance from the truth' that rational agents, as truth-lovers, aim to minimize.[7]

Accuracy domination can be represented geometrically via a slight adaption of the picture we used for the key lemma earlier. Rather than an arbitrary point in C, a, we now focus on an arbitrary world

[6]See Joyce (1998) and Joyce (2009) for philosophical arguments for probabilism based on Accuracy-Domination, in a classical setting.

[7]See Joyce (2009) for a description of the Brier score as a measure of accuracy. Joyce is particularly concerned, however, to prove results that are as neutral as possible on the nature of the accuracy-measure. He provides axioms for the accuracy measure and shows that any measure of accuracy satisfying those axioms will give rise to an Accuracy-Domination theorem. Author (suppressed) follows Joyce in this respect.

w. And we add to our diagram the 'third side' of the triangle cbw—the vector $u = b - w$. At each such w, we need the distance between b and w to exceed that between c and w, i.e. $||u|| > ||t||$. That is to say, we need to show that $u \cdot u > t \cdot t$.

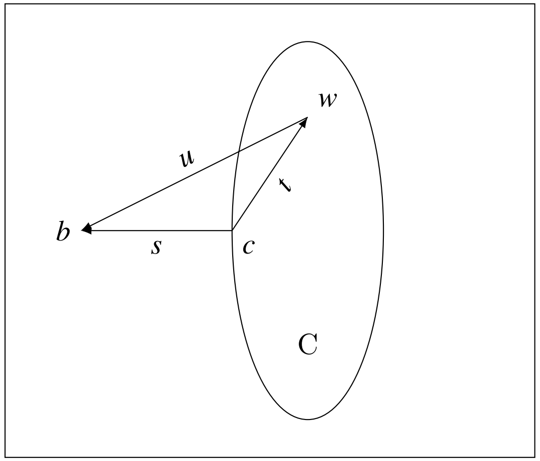

Figure 3: Accuracy domination: the length of u exceeds that of t.

This is elementary geometry, given the earlier lemma that the angle between s and t is non-acute. For completeness, let us run through the proof. Note first that $u = t - s$, hence:

$$\begin{aligned} u \cdot u &= (s-t) \cdot (s-t) \\ &= t \cdot t + s \cdot s - 2t \cdot s \\ &\geq t \cdot t + s \cdot s \\ &> t \cdot t \end{aligned}$$

The penultimate move here is secured by the key lemma—that $t \cdot s \leq 0$. The final move is the positivity of s, which, recall, is gotten by observing that s is the vector between b and the nearest point of C, which must be of positive length else b would be a member of C. This demonstrates that the square of the length of u exceeds the square of the length of t, whence our result.

(Note that the above argument shows that if the angle of the vertex of a triangle is non-acute, the opposite side is longer than the other two. We'll be appealing to this simple result several times below).

5 Dutch Book and Accuracy Domination

What is the relation between Dutch Books and points of Accuracy-Domination? We will see that every time you have a point c that Accuracy-Dominates b, you can use that point to construct a Dutch Book for b. Indeed $b - c$ will be a Dutch Book for b. The converse doesn't hold. There are Dutch Books that can be written in the form $b - c'$ where c' doesn't Accuracy-Dominate b. Ultimately however, we will see that every Dutch Book for b can be scaled so that it terminates at a point that Accuracy-Dominates b. The algebraic proof is given in the final paragraph of this section; but in order to see why this holds, it helps to first consider the case from an informal geometrical point of view.

The necessary and sufficient condition for s to Dutch Book b (relative to a set of worlds W) is that $s \cdot b > s \cdot w$ for all $w \in W$. Likewise, the necessary and sufficient condition for c to Accuracy-Dominate b (relative to W) is that $b - w > c - w$ for all $w \in W$. One easy consequence of the above characterization will be useful later: if $s \cdot b > s \cdot w$ then for any scalar $k > 0$, $k(s \cdot b) > k(s \cdot w)$; and hence $(ks) \cdot b > (ks) \cdot w$. Thus, when s is a Dutch Book, any positive scalar multiple of it is also a Dutch Book (intuitively, you just expand or shrink your guaranteed losses; but you can never reduce them to zero).

We first want to show that for every point c that accuracy dominates b (with respect to W), $b - c$ is a Dutch Book. We demonstrate this by proving the contrapositive: we start from the assumption that $s = b - c$ is not a Dutch Book, and show that c does not accuracy dominate b. So suppose that s is not a Dutch Book. This means that for some w, we must have $s \cdot w \geq s \cdot b$, or equivalently $0 \geq s \cdot (b - w)$.

We can simplify the statement of the problem by setting $u = b - w$ and set $t = c - w$ just as before, so that the situation is as depicted in figure 4.

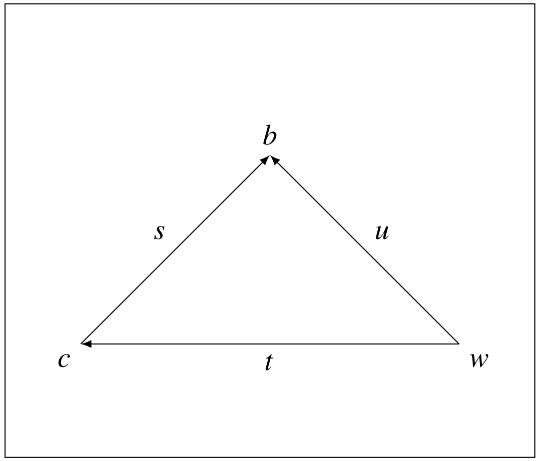

Figure 4: A situation where w witnesses that s does not Dutch Book b, i.e. $s \cdot (b - w) \leq 0$.

The assumption that s is not a Dutch Book translates as $0 \geq s \cdot u$. And what we want to prove is that c is just as far away from w as b is; i.e. $(c - w) \cdot (c - w) \geq (b - w) \cdot (b - w)$. More concisely, we need: $t \cdot t \geq u \cdot u$. Algebraically, we first note that

$$t = c - w = (b - w) - (b - c) = u - s.$$

We then argue that $t \cdot t = (u - s) \cdot (u - s) = u \cdot u + s \cdot s - 2u \cdot s$. Our assumption tells us that the the last term is positive, and hence $t \cdot t \geq u \cdot u + s \cdot s \geq u \cdot u$, just as required.

This argument has intuitive geometrical content. The starting point, $u \cdot s \leq 0$ tells us that the angle at b is non-acute. This must mean that the opposite side, t, is the longest in the triangle, and in particular, longer than u (the algebraic argument in effect uses the triangle equality to show this). Since the length of u is the distance from b to w, and the length of t is the distance from c to w, this is exactly what we need to demonstrate that c does not Accuracy-Dominate b.[8]

The converse does not hold. There are Dutch Books that can be written in the form $b - c'$, where c' does not Accuracy-Dominate b. The diagram in figure 5 below illustrates one such case.

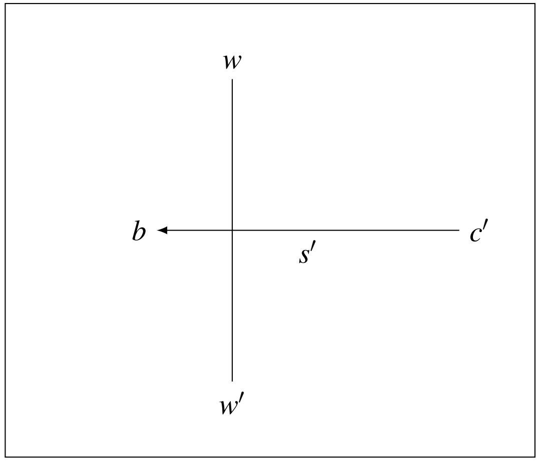

Figure 5: s' is a Dutch Book but c' is further from w, w' than is b.

[8]It is well known that no probability function admits of a Dutch Book. A corollary of the above is that no probability function can be Accuracy-Dominated either. For suppose that there are no Dutch Books for p. Then p cannot be Accuracy-Dominated. For suppose it was, by c, say. Then $p - c$ will be a Dutch Book for p, contrary to our assumptions. Thus, the immunity of probabilities to Accuracy-Domination (given the Brier score for measuring accuracy) follows immediately from the immunity of probabilities to Dutch Books.

6 A partial converse

In order to see what kind of converse is possible to the above, it helps to translate the algebraically specified necessary and sufficient conditions for Dutch Books and Accuracy Domination given above into something more geometrical. Those wanting to see the algebraic proof immediately may skip to the final paragraphs of this section.

For dutch-books s, the relevant condition is that $s \cdot (b - w) > 0$. Geometrically, that is to say that the angle at b between these two vectors is acute. Since this must hold for each world, all the worlds must all lie within some 'cone' whose tip is b and which expands in the opposite direction from which s points. See figure 6.

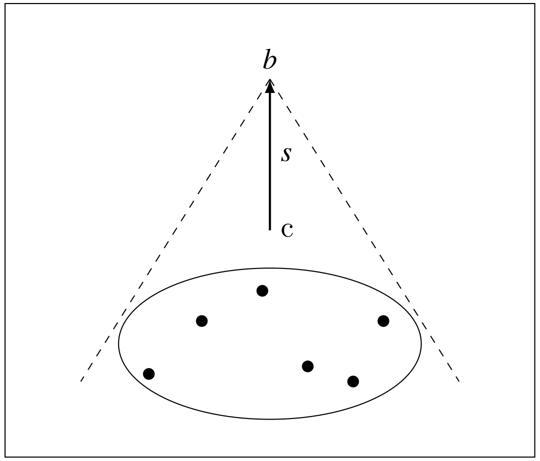

Figure 6: Dutch Book: All the worlds must lie within some cone projected from b by s

Notice that quite generally a sufficient condition for the vertex bcw (for arbitrary world w) at b to be acute, is that its angle at c is non-acute. Now in the De Finetti/Paris Dutch Book argument, we constructed a c where the latter condition held. But the condition while sufficient is not necessary. Geometrically, if s is long enough relative to $w - c$, the angles at c and b can both be acute, for some or all of the w. See figure 7 for a diagram where this occurs in a simple two-world setting.

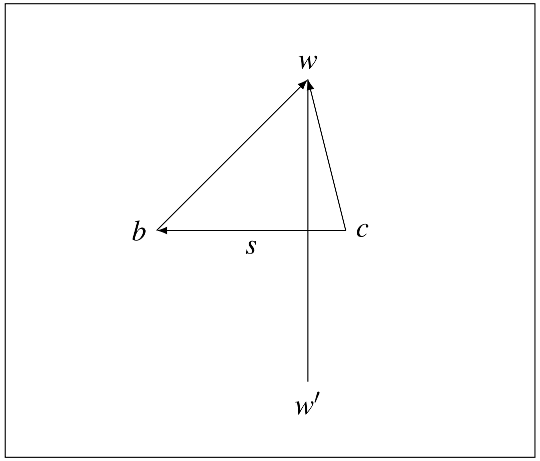

Figure 7: A two-world case in which s is a Dutch Book and c Accuracy-Dominates b even though the angle of bcw at c is acute. The points on ww' are the coherent belief states.

We can likewise find a geometrical characterization of the sets of points that Accuracy-Dominate b. Recall that the condition for this was just that $c - w < b - w$ for all $w \in W$. This has a very obvious geometrical interpretation: all the worlds must be nearer c than b. If we take the hyperplane that contains all points equidistant from b and c, all worlds must be on the same side—the side that also contains c. See figure 8. (It follows that if c Accuracy-Dominates b, then any point lying on the line between c and b will Accuracy-Dominate b).[9]

[9]Suppose we have a given set of worlds W and a belief state b. What are the Dutch Book vectors, and what are the points of Accuracy Domination, relative to these worlds and that belief state? To find the points of Accuracy Domination,

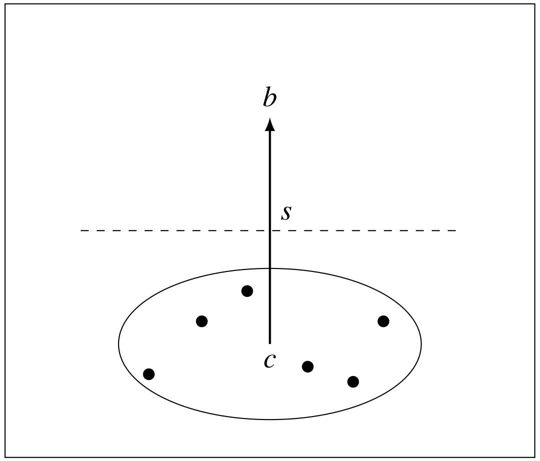

Figure 8: Accuracy-domination: all the worlds must lie below the line of equidistance between b and c

Once more, a sufficient condition for this to hold is that the angle at c be non-acute. But it isn't necessary—if the side cb (=s) is long enough (relative to $w-c$), then the relationship between sides can hold even if the angle at c is acute. That is why we do not have a full converse to the result proved in the previous section. A vector s can be a Dutch Book (i.e. the angle at b can be acute) while s's length is such that the condition for Domination is not met. This is the situation illustrated in figure 5 in the previous section.

However, the geometrical characterization also suggests a way to construct an Accuracy-Dominating point from any Dutch Book. The informal, geometrical argument is as follows. Recall that whenever we have a Dutch Book s for b with respect to W, ks is a Dutch Book b relative to W, for any positive scalar k. We will be able to find a k such that (a) ks is a Dutch Book; (b) such that $c' = b + ks$ Accuracy-Dominates b with respect to W; one simply chooses k small enough that its length is smaller than the components of each $w-b$ in the s-direction. This ensures that the angle at c' is obtuse, at which point our earlier proof of Accuracy-Domination kicks in.

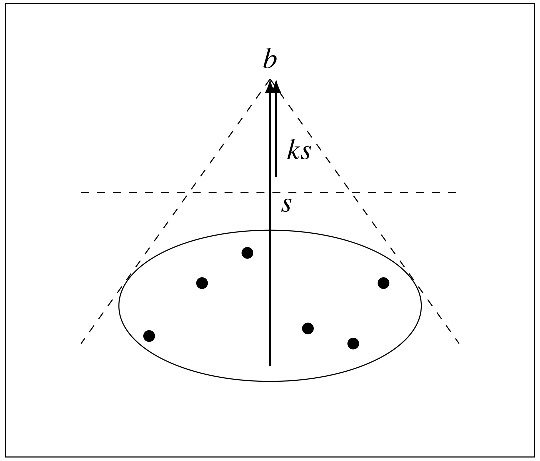

Figure 9: The horizontal line represents the closest world in the s-direction to b. We choose k so that the length of ks is small enough so that its endpoint is above the line. s and ks are both Dutch Books. $b+ks$ Accuracy-Dominates b, since the angle from it to an arbitrary world will be non-acute.

draw spheres around each world, such that b lies on the surface. A necessary and sufficient condition for x to be a point of Accuracy Domination is that it is in the interior of each sphere. Now consider the tangents to each sphere at b. These tangents partition the overall space. Consider that cell of the partition that contains the points of Accuracy Domination. Any vector that starts at a point in this cell and has b as its endpoint is a dutchbook.

Consider the simplest case of just two worlds, where b is equidistant from each. If the angle at b is over 60 degrees, then both the cell of Dutch Books and the region of Accuracy-Dominating points include every coherent point. If the angle at b is between 60 and 90, then the coherent points are within the cell of Dutch Books, but not every coherent point Accuracy-Dominates. If the angle at b is greater than 90 degrees, then there are coherent points out of the space of Dutch Books and which do not Accuracy-Dominate. Finally, when the angle is 180, the point itself is coherent and neither Dutch Booked nor Accuracy-Dominated.

Algebraically, we are given belief state b with Dutch Book s. Set $s_k = ks$. Paralleling our earlier notation, we let c_k be the endpoint of s_k from b (i.e. $s_k = b - c_k$) and for arbitrary $w \in W$, we let $u^w = b - w$ and $t_k^w = c_k - w$. Note that $t_k^w = u^w - s_k$ (see figure 10). That s is a Dutch Book gives us that $s \cdot u^w > 0$ for each $w \in W$.

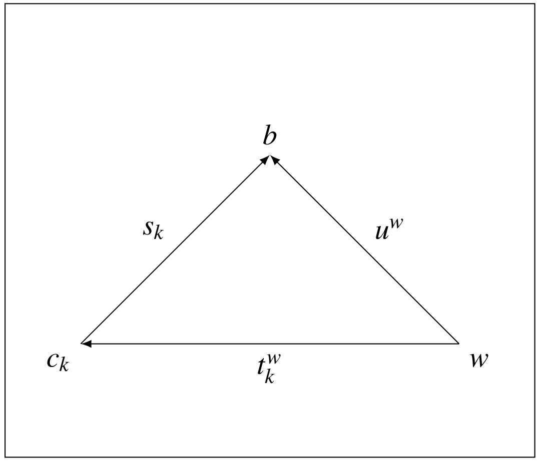

Figure 10:

We then argue:

$$\begin{aligned} t_k^w \cdot t_k^w &= (u^w - s_k) \cdot (u^w - s_k) \\ &= u^w \cdot u^w + s_k \cdot s_k - 2 s_k \cdot u \\ &= u^w \cdot u^w + k^2 (s \cdot s) - 2k(s \cdot u) \quad \text{(since } s_k = ks\text{)} \\ &= u^w \cdot u^w + k(k(s \cdot s) - 2(s \cdot u)) \end{aligned}$$

The condition for c_k to accuracy dominate b is for $(c_k - w) < (b - w)$ for each w, which is equivalent to $t_k^w \cdot t_k^w < u^w \cdot u^w$ holding for each w. We need to pick k to ensure this holds. By the above, it holds whenever $k(k(s \cdot s) - 2(s \cdot u^w)) < 0$. This is so for any choice of k meeting:

$$0 < k < 2\frac{s \cdot u^w}{s \cdot s}.$$

But this is satisfiable, since both $s \cdot s$ and $s \cdot u^w$ are greater than zero (the latter, recall, follows directly from s being a Dutch Book). To ensure that c_k Accuracy-Dominates b, we need this condition to be met for each world $w \in W$; since there are only finitely many w, this can be done.

7 Diachronic Dutch Books

The above completes our survey of the relation between Accuracy-Domination points and Dutch Books. We now turn to the question of whether the De Finetti approach can be extended to other results in the area. Defences of updating by conditionalization using a diachronic Dutch Book are well-known. This section formulates geometrical constraints that are required to avoid a diachronic Dutch Book argument in a potentially non-classical setting. Subsequent sections show this reduces to classical conditionalization and examines the analogues in non-classical settings.

Consider the familiar setup for a Dutch Strategy, due to David Lewis:

> Suppose that at time 0, you have a coherent belief function M. Let $E_1, \ldots, E_n$ be mutually exclusive and jointly exhaustive propositions that specify, in full detail, all the alternative courses of experience you might undergo between time 0 and time 1. Let M_i be the belief function you would have at time 1 if you had the experience specified by E_i—that is, if E_i were the true one of $E_1, \ldots, E_n$.
> (Lewis, 1999, p.405)[10]

Conditionalizers claim that in such a setting M_i should be the result

[10]This quote is from the Lewis paper, first published in 1999. Supplemented with a 1997 introduction, this reprints the original 1972 handout that was the source of the version of the Diachronic Dutch Book argument in Paul Teller (1973).

of conditionalizing M by E_i. Lewis asks: why would it be irrational to respond to experience in any other way?

Lewis's famous response is to describe a three-step strategy for extracting money from non-conditionalizers. Supposing that $M_i(P)$ is less than $M(P|E_i)$, he would (i) sell you a certain book of bets at time 0, at fair price; (ii) wait and see whether E_i is true; (iii) If E_i is true, sell another book of bets. In Lewis's setting, your net loss/gain is zero if E_i is not the case; and you lose money if E_i obtains. Thus Lewis says: "I can inflict on you a risk of loss uncompensated by any chance of gain" (we shall prove a slightly stronger result, whereby a sure loss is obtained). A dual pattern of bets can inflict similar results on anyone whose updated credences are higher than the result of conditionalization would be. Crucially, in implementing this strategy, the bookie needs no information beyond that which the subject possesses.[11]

The philosophical debate over what the argument shows takes off from this point. As Lewis notes in his 1992 introduction to the paper, the argument is 'addressed to a severely idealized, superhuman subject who runs no risks of mistaking his evidence'. If the subject cannot respond differentially to distinct 'course of experience' she undergoes between time 0 and 1 (perhaps they involve perceptions of colour patches whose wavelengths differ by a single nanometer), then we cannot legitimately put forward a strategy that relies on the bookie discriminating between the two courses of experience—again, there's no prospect of an argument that the subject is irrational on the basis that they can lose money when faced with a bookie who has strictly more information (Lewis mentions Jeffrey-conditionalization as an updating rule designed for less idealized settings). Another assumption of the case is that the subject and bookie are both certain how the subject reacts to the various pieces of evidence that come in—so that the bookie may legitimately build her strategy around exploiting this. Even granting this internalist assumption of perfect reflective grip on evidence and reactions to evidence, and just as with synchronic Dutch Books, one can also query whether the technical 'sure loss' results justify the normative conclusion that updating by anything other than conditionalization is irrational. But our concern here is with getting a grip on the technical result itself, so I set these philosophical concerns aside.

To build our geometrical picture of diachronic Dutch Books, we shall continue to assume that $E_1, E_2, \ldots, E_n$ represent an exhaustive, exclusive list of possible courses of experience. We let b be the point in belief-space that gives your current (coherent) credences; and c belief-state that results from implementing your updating strategy, on undergoing experience $E = E_i$. Our assumption is that both subject and bookie know the identities of b and c, but are uncertain whether E will come about, though they will recognize it as and when it occurs.

Two preliminary points should be noted. The first is that b itself must be a convex combination worlds in W. The second is that c must be a convex combination of worlds in E. We can argue for each by pointing to independent ways of constructing a diachronic Dutch Book if they fail. In the first case, we simply note that if b fails to meet the stated condition, we can find a synchronic Dutch Book for it. And this (with no further betting) will count as a limiting case of a diachronic Dutch Book. In the second case, if c fails to be a convex combination of worlds in E, from our earlier results once more we find a 'Dutch Book relative to the worlds in E' against c—i.e. a book of bets e that guarantees a loss at all $w \in E$. This isn't immediately something that gives the agent a *sure* loss, since if information other than E arrives the agent suffers no change in fortune whatsoever. But with a minor tweak we can create one. Suppose the guaranteed loss of e among E-worlds is at least m. Then consider the following Dutch Strategy: sell the agent a bet on E that pays out $m/2$ iff E obtains—for whatever its fair price is at b; and sell her e for its fair price in the eventuality that

[11] This last constraint crucial to understanding why Lewis sets up the argument in terms of the evidence partition in particular (if the argument went through with cells from an arbitrary partition, then we would be able to generate inconsistent updating advice). See his paper for detailed discussion.

E obtains. Two cases are possible: if we get information other than E, the agent faces a net loss of whatever she paid for d. If information other than E is received, then she'll get some gain from the previous bet $(m/2)$, but is guaranteed to lose m on the new bet. The net loss is then at least $m - m/2 > 0$. So we have our Dutch Strategy, and the only way to avoid it is for the updated point c to be a convex combination of the worlds in E, as required.

How, in general, are we to build a diachronic Dutch Book? Our strategy will be very simple: we will sell a book of bets d, and then, if the information E is received, we'll buy it back at the (possibly changed) fair price. If some other information is received (i.e. some world in the complement of E—$\bar{E}$—is actualized) we do nothing. All the action is in choosing the initial books carefully enough. The necessary and sufficient conditions for this to be a Dutch Strategy are:

1. $d \cdot w - c \cdot b < 0$ for each $w \in \bar{E}$.

2. $(d \cdot w' - d \cdot b) - (d \cdot w' - d \cdot c) < 0$ for each $w' \in E$.

The first condition corresponds to those cases where we receive information other than E. For sure loss, the return minus the price paid for d must be always negative. The second condition covers the E-cases, where the initial bet is bought back. Here, we add together the net contribution of the bet that is sold back to the agent and the net contribution of the original bet. Again, for Dutch Strategy we need this always to be negative.

The second condition simplifies, since the two occurrences of $d \cdot w'$ cancel. It becomes simply $(d \cdot c - d \cdot b) < 0$. Note that this doesn't depend in any way on which world is actual. So the condition for a buy-back strategy based on d to be diachronic Dutch Book is simply that $d \cdot x - c \cdot b < 0$, for every x in $\bar{E} \cup \{c\}$. In effect, this reduces the problem of finding a 'buyback' form diachronic Dutch Book, to the problem of finding a synchronic Dutch Book over a certain space of 'worlds' (note that in the highly abstract setting we are working within, it's quite legitimate to construe c as a 'world' in its own right—nothing in the formal results rested on *which* points in the space were identified with worlds). Because of this, our earlier results tell us that we can find a suitable Dutch Book d unless b is a convex combination of the relevant set of 'worlds', i.e. the worlds in $\{w_1, \ldots, w_n\} = \bar{E}$ together with c. This is represented in general in figure 12. Holding fixed E and b, this imposes an immediate geometrical constraint on what c can be—its implication in a special case is illustrated in figures 13 and 14.

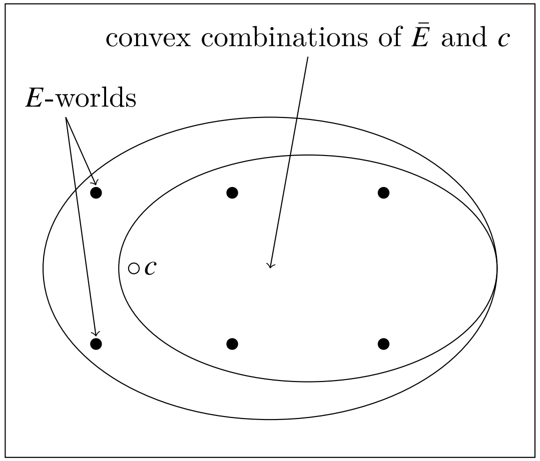

Figure 11: Any point within the outer ellipse is a coherent belief state (convex combination of worlds). Points within the inner ellipse are convex combinations of $\neg E$ worlds and c. With a sell-back strategy in place, b will be dutch-bookable unless it is in the latter set. c is also constrained to be a convex combination of the E worlds. So given b, we must pick a c such that both of these are met. This turns out to determine c uniquely.

The geometrical characterization of updating is now complete. The question is whether we can find some informative algebraic character-

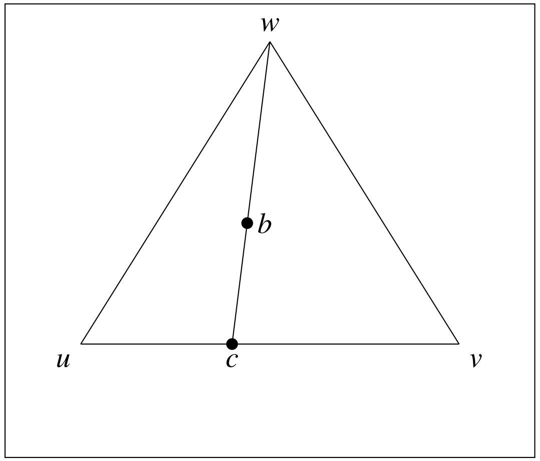

Figure 12: The geometrical constraint on updating for the special case of three worlds, updating on $E = \{u, v\}$. The constraint is that c be chosen so that b is a convex combination of w and c. In the special case illustrated, this amounts to 'projecting' b from w onto the convex combinations of u and v (note that b is a convex combination of the three worlds, and c a convex combination of u and v, as per our pair of preliminary results).

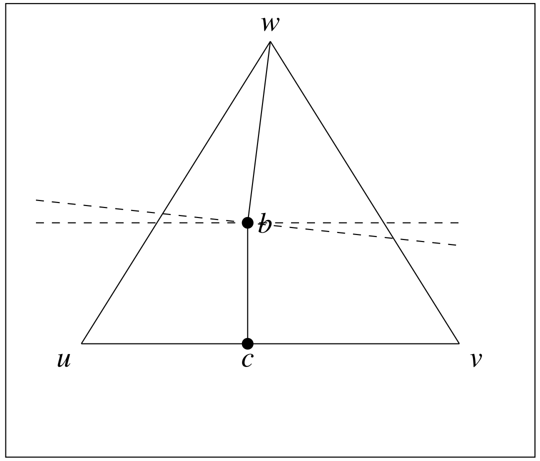

Figure 13: If b does not lie on the line between w and c, then (by our previous results) we can find a Dutch Book for b relative to the 'worlds' w and c; which we now know will count as a diachronic Dutch Book for updating from b to c. In the updating procedure sketched above, any vector directed to b and lying between the dashed lines on the right will be a diachronic Dutch Book.

ization—and whether it coincides with conditionalization in the standard cases. It turns out that we get exactly this result, by elementary calculation.

Our starting point for extracting an algebraic characterization of our geometrical constraint is that the requirement that b be a convex combination of points in $\{w_1, \ldots, w_n, c\}$ says that there are parameters $\alpha, \lambda_i \geq 0$, with $1 = \alpha + \sum_i \lambda_i$, such that:

$$b = \alpha c + \lambda_1 w_1 + \ldots + \lambda_n w_n$$

Rewriting (assuming $\alpha \neq 0$):

$$c = (1/\alpha)(b - \lambda_1 w_1 - \ldots - \lambda_n w_n)$$

This already shows the general form that updating b by E to get c must take—b is diminished by taking away components that live in $\bar{E}$, and then the result is scaled by α. Many updating rules would satisfy this abstract formulation, however, so we need to press on.

Since b is a convex combination of the original set of worlds, with weights μ'_i, μ_j, we write:

$$b = \mu'_1 w'_1 + \ldots + \mu'_m w'_m + \mu_1 w_1 + \ldots + \mu_n w_n$$

Write b^E for the component of b that is a convex combination of the worlds in E, i.e. $b^E = \mu_1 w'_1 + \ldots + \mu'_m w'_m$. Then we have:

$$b = b^E + \mu'_1 w_1 + \ldots + \mu_n w_n$$

Substituting this into the characterization of c above, and collecting the components of the respective w_i, we have:

$$c = (1/\alpha)(b^E + (\mu_1 - \lambda_1)w_1 + \ldots + (\mu_n - \lambda_n)w_n)$$

We also know that c must be a convex combination of the E-worlds, and hence the coefficient of each $\bar{E}$ world (the w_i) must be zero. Hence $\lambda_i = \mu_i$ and the w_i terms above disappear, leaving:

$$c = (1/\alpha)b^E$$

Note also that $\alpha = 1 - \sum_i \lambda_i = 1 - \sum_j \mu_j$ (by the identity we just derived). By construction, $1 - \sum_j \mu_j = \sum_i \mu'_i$, and hence $\alpha = \mu'_1 + \ldots \mu'_m$. Once we've specified what the 'worlds' are, this equation allows us to derive analogues of conditionalization.

First, suppose our worlds w_i, w'_i are truth value distributions across propositions. (To keep matters clear, it's important to distinguish between 'propositions' used in interpreting the formalism, and sets of worlds like E or $\bar{E}$.) Now consider the proposition that happens to be true (takes value 1) at all and only the worlds in E. Call this E_*.

Now consider $b(E_*)$—the entry in the vector b corresponding to the proposition E_*. We already know that $b = \mu_1 w'_1 + \ldots + \mu'_m w'_m + \mu_1 w_1 + \ldots + \mu_n w_n$, which means that focusing on the relevant component we have:

$$b(E_*) = \mu_1 w'_1(E_*) + \ldots + \mu'_m w'_m(E_*) + \mu_1(w_1)(E_*) + \ldots + \mu_n w_n(E_*)$$

But by construction $w'_i(E_*) = 1$ for each i, and $w_j(E_*) = 0$ for each j (this just reflects the fact that E is true at each of the w'_i and false at each of the w_j). Hence we obtain:

$$b(E_*) = \mu'_1 + \ldots + \mu'_m$$

which by an earlier result is just α.

Second, consider an arbitrary proposition X, and let $E_* \circ X$ be a proposition whose truth value at any w or w' is given by multiplying together the respective truth values of X and E_* at that world. In the classical setting, this recipe delivers the truth value of the conjunction $E \wedge X$—we'll discuss the non-classical interpretations (a la Paris) in the next section.

Recall that $b^E = \mu'_1 w'_1 + \ldots + \mu'_m w'_m$; and so $b^E(X) = \mu'_1 w'_1(X) + \ldots + \mu'_m w'_m(X)$. And recall once more than E takes truth value 1 at the w'_i and 0 at the w_j. Hence we can write:

$$b^E(X) = \sum_i \mu'_i(w'_i(E_*)w'_i(X)) + \sum_j \mu_j(w_j(E_*)w_j(X))$$

But notice that by construction, $w_i(E_* \circ X) = w_j(E_*)w_j(X)$, and likewise for the w'_i, so the above becomes:

$$b^E(X) = \sum_i \mu'_i(w'_i(E_* \circ X)) + \sum_j \mu_j(w_j(E_* \circ X))$$

which is simply $b(E_* \circ X)$.

We noted earlier that in the classical setting, this is simply conjunction, and hence:

$$c(X) = \frac{b^E(X)}{b(E_*)} = \frac{b(X \wedge E_*)}{b(E_*)}$$

This is the familiar ratio characterization of conditionalization.

8 Non-classical conditionalization

To recap: we described a very general diachronic Dutch Book argument, which led to a geometrical constraint on permissible updates of b by E: the updated credence c must be such that b is a convex combination of c and the $\bar{E}$ worlds. This this turns out to impose a constraint that we can most generally state as follows:

$$c(X) = b(X \circ E_*)/b(E_*)$$

If the propositions are classical $\circ$ is just conjunction, and we have the standard ratio characterization of conditionalization. But the advantage of the general setting is that the abstract characterization can be used as a constraint on updating generalized probabilities, where X can be an arbitrary random variable, or 'non-classical proposition' taking 'non-classical' truth values between 1 and 0. On this interpretation, $X \circ E_*$ is then a new random variable derived from X and E_* by multiplication. Relative to a given non-classical setting, $\circ$ is a possible truth-functional connective; and so we have the familiar ratio formula whenever the truth value distributions V meet the following constraint:

$$(\mathcal{T}p) \quad V(A) = r \wedge V(B) = s \iff V(A \wedge B) = rs$$

Note that $(\mathcal{T}p)$ is a strengthening of Paris's $(\mathcal{T}2)$. It is a feature of many of 'glut' and 'gap' semantics where propositions take truth values in $\{1,0\}$.[12]

There are well-known non-classical semantics for conjuntion where this fails, however. For example, the Lukasiewicz many-valued systems feature two kinds of conjunction: weak (where the value of the conjunction is the minimum of the value of the conjuncts) and strong (where the value of the conjunction is the sum of the values of the conjuncts, minus 1, when this is positive). Neither equals the product conjunction, which is studied in its own right in the 'product' t-norm logic (cf. Hájek, 1998).

This isn't a severe limitation for our purposes, however, due to a restriction that's been built into our Dutch Strategy construction from the beginning. Though we have a formulation of conditionalization that could can be used to update generalized probabilities *invested in* arbitrary non-classical propositions X, the argument is nevertheless restricted in scope. In particular, as currently formulated the proposition *updated upon* must be some E_* that is the characterizing function for some cell E within the 'evidence' partition of the set of non-classical worlds (i.e. E_* must take value 1 or 0 at each world, depending on whether that world is a member of E). Our belief states

[12]Potential applications include versions of Kleene and the paraconsistent logics (with an suitable underlying 'gap' and 'glut' semantics); likewise supervaluational and subvaluational systems. See (Author) for discussion.

are general; but in the argument as presently formulated, updates have to be on propositions that are 'classical' in this sense for our results to kick in. Because of this restriction in the values that E_* can take, less is needed to justify representing conditionalization (for the restricted range of propositions) by conjunction. The following will suffice:

$$(\mathcal{T}p') \quad \begin{array}{lcl} V(A) = r \wedge V(B) = 1 & \Longrightarrow & V(A \wedge B) = r \\ V(A) = r \wedge V(B) = 0 & \Longrightarrow & V(A \wedge B) = 0 \end{array}$$

This will ensure that *when the second conjunct is taking 'classical' values* the truth value of the whole equals the product of the two values. And this is satisfied by each of the fuzzy conjunctions we have mentioned.

The restriction just noted raises the question of what updates on genuinely 'nonclassical' propositions should be like. Of course, the constraint

$$c(X) = b(X \circ Y)/b(Y)$$

is well-defined for arbitrary Y,[13] and a natural question is whether this is the appropriate general formulation (at which point our caveats about the relation between $\circ$ and conjunction in various systems would kick in). But there are both local and general difficulties with evaluating this proposal. The local difficulty is that the very setup of the diachronic Dutch Book is in terms of a classical partition of worlds—so the argument would need to be rethought from the ground up. And the conceptual difficulty we hit in doing this is that it is not at all obvious how to conceptualize the 'evidence' gained by learning a non-classical proposition that may (for example) taken intermediate truth-value at some worlds (does learning some proposition that is 0.9 true at world w raise or lower our credence in w, all else equal?)

There are two very different cases to consider. The first is one where every proposition takes values from $\{1,0\}$. The second is where propositions can take a wider range of truth values. I am not sure how one could adapt the argument for updates on 'fuzzy' propositions in the latter case, so set it aside. But in the former case, the formal argument itself would not need changing. After all, the proposition E_* was brought in only at the end of the argument; and the assumption was simply that it took value 1 at all worlds in the cell E; and 0 everywhere else. This last constraint could be satisifed by a non-classical proposition of the kind we're considering (you might say that propositions of this kind are not *intrinsically* non-classical, but only relationally so—their non-classicality is only manifest in the behaviour of their negations, or more generally in the truth-values of compounds in which they are part.) The trouble comes, if anywhere, with the philosophical framing of the formal argument. Lewis started from a partition of worlds where each cell corresponds to a possible total course of experience. For the argument to go through, one has to accept that the information transmitted in learning that P is consistent with each world in where it takes value 1 and rules out each world where it takes value 0. The last condition is the contentious one: a weaker alternative is that the information transmitted is only inconsistent with worlds where $\neg P$ holds—leaving open the relation of the information to worlds in which P and $\neg P$ both take value 0.

To summarize: whenever we work with a backdrop (classical or non-classical) where the truth values are restricted to $\{1,0\}$, then a wholly general argument for updating-by-conditionalization may be propounded. However, the argument may be resisted on philosophical grounds if we can make the case that learning that P and learning that one is in a world where P has value 1 are importantly distinct. This is not a formal issue: it is a question of the philosophical interpretation of the semantic values we are working with, and several interpretations of the same formalism may be available (compare the epistemic and non-epistemic interpretations of partial value assignments mentioned in the first section). But even when the general argument fails, we still have a restricted rule of updating generalized probabilities, on

[13] Always supposing that $b(Y) > 0$.

receipt of fully classical information.[14]

Suppose that our argument is indeed restricted in scope even in the $\{1,0\}$ settings just considered. The constraint on updating then delivered cannot be assumed to point unambiguously to the generalization to the full setting we have hitherto been discussing. To illustrate this, we can consider a particular application of these methods—to Dempster-Shafer belief functions. There are two extant rival accounts of updating that are compatible with the restricted updating principle (where E_* is restricted to 'classical' propositions), each diverging from the generalized update rule considered above. Given a belief state b, define the 'plausibility' of a proposition A as $p(A) := 1 - b(\neg A)$. The two new rules, together with the generalization we've just been looking at, are:[15]

[14] Other diachronic constraints on conditioning may find purchase in the wholly general setting even if the above does not. Let $c_1, \ldots, c_n$ be the possible belief states one might end up in response to some information conveying the propositions $E_1, \ldots, E_n$ (set aside the issue of whether those propositions are classical or non-classical, and how to think about the update strategy). Now consider a universal sell-back strategy—that is, we sell book d, and then buy it back at a later time no matter what information has been received in the meantime. The condition under which this maximally simple strategy leaves the bettor with a sure loss is that the following hold:

- $(d \cdot w - d \cdot b) - (d \cdot w - d \cdot c_i) < 0$ for each $w \in W$

In each case, $d \cdot w$ cancels, so this reduces to:

- $(d \cdot c_i - d \cdot b) < 0$ for each i, $1 \leq i \leq n$.

Again, we can treat the c_i as 'worlds' for the purposes of applying our synchronic Dutch Book theorem, with the result that the condition for such a d to exist is that b fail to be a convex combination of the prospective belief states c_i (compare the 'fundamental theorem of arbitrage' discussed in (Skyrms, 2006)). This diachronic coherence constraint holds whether the worlds or credences are classical or non-classical (and as Skyrms notes, once we add in resources to explicitly represent credences in future credences, this is intimately related to 'reflection' principles).

[15] The first two are taken as well defined so long as $p(Y) > 0$; the last if $b(Y) > 0$. If $p(X \wedge \neg Y) = 0$, mush-conditionalization is stipulated to be equal to 1 (even if $b(Y \wedge X) = 0$; see (Halpern, 1995, p.93) for discussion).

Mush-Conditioning

$$b(X|Y) = \frac{b(Y \wedge X)}{b(Y \wedge X) + p(Y \wedge \neg X)}$$

DS-Conditioning

$$b(X||Y) = \frac{b(X \vee \neg Y) - b(\neg Y)}{p(Y)}$$

Naive-Conditioning

$$b(X|||Y) = \frac{b(X \wedge Y)}{b(Y)}$$

In the extant literature, the first two rules are candidate updates associated with distinct (epistemic) interpretations of DS belief functions. Generally, these three update rules will diverge; an example is given in the appendix to this paper. Furthermore, the appendix shows that all three coincide when $p(Y) = b(Y)$—that is, in order for them to recommend the same update, it is sufficient that the belief state that is to be updated behaves 'classically' *on the particular proposition conditioned upon*. This means that the diachronic Dutch Book result if restricted to 'classical information' enforces something that is the common core of all three update rules.

Conclusion

Accuracy-domination and Dutch Books are intimately linked. Whereever we have a point of Accuracy-Domination, there we find a Dutch Book. And wherever we have a Dutch Book, we need only scale it sufficiently to find a point of Accuracy Domination. The same fundamental geometrical result—the separating hyperplane theorem—underpins both.

In the final sections, we considered the conditions under which we find Dutch Strategies or diachronic Dutch Books. The key result we see here is that a 'buy-back' strategy allows us to turn the task of finding a diachronic Dutch Book into that of finding a synchronic one—albeit with respect to a modified range of 'worlds'. We've seen how to derive the standard characterization of conditionalization from this results in the classical setting. We get an analogue in the non-classical setting; though there are some subtleties with the formulation. The most important is the question of whether the argument applies to updating on genuinely 'non-classical' propositions. In some settings, this unrestricted result may be argued for. But even if the argument is restricted in scope, it limits the range of admissible update rules in interesting ways.

Given the close relationship between Dutch Books and Accuracy Domination in the synchronic case, one would have expected to find a parallel diachronic Accuracy-Domination argument. However, one does not find this in the literature. The nearest are the *expected* accuracy defenses of conditionalization in (Greaves & Wallace, 2006) and (Leitgeb & Pettigrew, 2010); but appeals to expectations (fixed by one's prior belief state) mean these are of a very different character from the Joyce-style synchronic argument—why should you trust an *outdated* belief state to tell you how to fix your beliefs now you have new information? However, it is tricky to see what a true domination argument constraining update would look like. After all, a rule that told us to jump to a belief state that matched the truth values at w would maximize accuracy of the posterior state at at least one world. I believe there is room for a more subtle domination argument for conditionalization, built on the geometrical characterization of conditionalization outlined above. But exploring this would require extensive discussion of the relation between norms on belief and update strategies, and this will not be pursued here.

Our discussion of Accuracy Domination focused on the special case where 'accuracy' is measured by the Brier Score (square Euclidean distance). This is clearly central to the results above: within our vector space we move easily from properties of acuteness/non-acuteness of angles that are central to the characterization of Dutch Books in particular, and which are expressed via the inner product of two vectors; to properties of the distance between points in the space, thought of in terms of the length of the vector joining the two (the inner product of the vector with itself). But while the Brier score is perhaps the most prominent candidate for measuring accuracy, much of the literature on Accuracy-Domination aims to prove its results in a more general setting.

Not all accuracy measures will correspond to an inner product on the space of vectors. But some will, and in those settings, the relationship to 'Dutch Book theorems' will be well-defined. Two natural projects suggest themselves: first, to investigate the relationship between Dutch Books (defined as above via the standard Euclidean inner product) and Accuracy-Domination characterized via a different inner product. Second, to investigate the relationship between Dutch Books characterized via the inner product relevant to Accuracy-Domination, and the Accuracy-Domination itself. In this setting, do the arguments above go through? And is there a betting-interpretation of the dutch-book theorem so produced (in particular, can we think of $z \cdot w$ as the 'returns' of a book of bets at w, and $z \cdot b$ as its fair price)?

A Appendix: Updating DS belief functions

Recall we had three candidate update methods for DS belief functions:

Mush-Conditioning

$$b(X|Y) = \frac{b(Y \wedge X)}{b(Y \wedge X) + p(Y \wedge \neg X)}$$

DS-Conditioning

$$b(X||Y) = \frac{b(X \vee \neg Y) - b(\neg Y)}{p(Y)}$$

Naive-Conditioning

$$b(X|||Y) = \frac{b(X \wedge Y)}{b(Y)}$$

The following chart shows describes the truth values of X and Y (and compounds thereof) at three worlds (the truth distributions on the three worlds can be thought about in a number of ways; one is by the rule that something gets value 1 in world α iff it is a classical consequene of $X \wedge Y$; gets value 1 in β iff it is a classical consequence of X; and gets value 1 in γ iff it is classical consequence of $\neg Y$. Cf. Paris's characterization of 'partially formed worlds' (2001, p.xx).) By Jaffray's results, any convex combination of these will be a DS function—the rightmost columns give the belief and plausibility levels appropriate to the particular weights written at the head of each column (recall that $p(A) = 1 - b(\neg A)$).

Proposition	$\frac{2}{10}$ value at α	$\frac{3}{10}$ value at β	$\frac{5}{10}$ value at γ	*belief*	*plausibility*
X	1	1	0	0.5	0.5
$\neg X$	0	0	1	0.5	0.5
Y	1	0	0	0.2	1
$\neg Y$	0	0	0	0	0.8
$Y \wedge X$	1	0	0	0.2	0.5
$\neg(Y \wedge X)$	0	0	1	0.5	0.8
$Y \wedge \neg X$	0	0	0	0	0.5
$\neg(Y \wedge \neg X)$ $= (\neg Y \vee X)$	1	1	0	0.5	1

We now calculate the belief in X under updating by Y in each of the three forms:

Mush-Conditioning

$$b(X|Y) = \frac{b(Y \wedge X)}{b(Y \wedge X) + p(Y \wedge \neg X)} = \frac{0.2}{0.2 + 0.5} = \frac{2}{7}$$

DS-Conditioning

$$b(X||Y) = \frac{b(X \vee \neg Y) - b(\neg Y)}{p(Y)} = \frac{0.5 - 0}{1} = \frac{1}{2}$$

Naive-Conditioning

$$b(X|||Y) = \frac{b(X \wedge Y)}{b(Y)} = \frac{0.2}{0.2} = 1$$

Hence the three update rules diverge in general. A quick note on their interpretation: the usual application of Dempster-Shafer belief functions is epistemic, with the gap between 'belief' and 'plausibility' representing the extent of uncertainty on the agent's part over whether the proposition in question is true. Mush-conditioning and DS-conditioning are taken to represent two different conceptions of in what this uncertainty consists (see (Halpern, 1995, §3.6) for discussion). A situation where A and $\neg A$ are both value 0 represents a situation concerning which you are uncertain whether A holds. Naive conditioning does not fit well with this application, since it effectively takes the information A to be sufficient to eliminate possibilities where A has value 0. A quite different application of DS belief functions is one where we interpret the truth values by naive extension from the classical case: A is neither true nor false in a situation in which both it and its negation are true (this is the interpretation that may be associated with a background supervaluational semantics—see (Author)). On that reading, eliminating possibilities where A is untrue, given information A, does not seem unreasonable. These observations fit neatly with the point made in the main text, that whether or not the strong reading of the Dutch Strategy argument is successful (showing that one must update via naive conditioning) will typically depend on the interpretation given to the belief states. The weaker form of the argument establishes a core form of updating that all the interpretations mentioned agree on, as we will now show.

What we will prove is that all three update methods will coincide if the proposition updated upon, X, is 'classical' by the lights of the belief function; i.e., if $b(Y) = p(Y)$. This will follow from the stronger assumption that at every world, either Y has value 1 or $\neg Y$ has value 1. The converse also holds if the belief function we are working with (which we know can be represented as a convex combination of appropriate truth value distributions) has a non-zero coefficient at each world. Since we can just throw away all worlds with zero coefficients for present purposes, we shall take it that 'bivalence' for Y holds of all the worlds in W.

The characterizations of belief functions appealed to above is in terms of 'partial' value assignments. There is an isomorphic representation in terms of *sets* of complete, classical truth value assignments. We let A^* be the set of all classical assignments that assign A value 1, and for an arbitrary set of classical assignments S, we let S_* be the strongest proposition that is true at each element of S. One can check that $(A^*)_* = A$ and $(S_*)^* = S$. Moreover A^* can be characterized as the 'classical completions' or 'sharpenings' of the partial assignment V_A—a classical assignment is in A^* iff it assigns 1 to X whenever V_A does.

By construction, Z is true at all members of a non-null set S iff $S \subset Z^*$. Also by construction, $V_A(Z) = 1$ iff Z is true at each member of A^*. Putting these together, $V_A(Z) = 1$ iff $A^* \subset Z^*$ Recall that b was a convex combination of the partial value assignments. So there are some λ_A such that $b(Z) = \sum \lambda_U V_U(Z)$. If we set $m(U^*) = \lambda_U$, then we have $b(Z) = \sum m(U^*)V_U(Z)$. But since $V_U(Z)$ is 1 only when $U^* \subset Z$ and otherwise 0, this means $b(Z) = \sum_{U^* \subset Z^*} m(U^*)$. This is the 'mass function' representation of a belief function over the *powerset* of classical truth value assignments. By similar reasoning, the plausibility level $p(A) = 1 - (b(\neg A)$ can be written as $p(Z) = \sum_{U:Z\cap U \neq \emptyset} m(U)$. See Halpern (1995, §2.4) for details.

For simplicity of notation, we will use roman capital A both to stand both for a proposition (a random variable taking truth values at worlds) and for the subset of classical truth value assignments A^*. Note that $\neg A = \bar{A}$, $A \vee B = A \cup B$, $A \wedge B = A \cap B$, and so on.

We first use this representation to show that DS conditioning collapses to naive conditioning, on the assumption that $b(Y) = p(Y)$. By the above we may write belief and plausibility values for Y in mass-function terms as follows:

$$b(Y) = \sum_{U \subset Y} m(U)$$

$$p(Y) = \sum_{U:Y\cap U \neq \emptyset} m(U)$$

Since $b(Y) = p(Y)$, the two sums must coincide. $U \subset Y$ entails that $Y \cap U \neq \emptyset$, so every element in the first sum features in the second. For the sums to coincide, no other term in the second sum can contribute any mass. Hence our assumption entails that when $m(Z) \neq 0$ and $Y \cap Z \neq \emptyset$, then $Z \subset Y$.

We want to calculate $b(X \vee \neg Y)$, or equivalently $b((X \wedge Y) \cup \bar{Y})$. Take any Z such that $Z \subset (X \cap Y) \cup \bar{Y}$. Either Z is entirely contained in $\bar{Y}$; or it has non-empty intersection with $X \cap Y$, and hence a non-empty intersection with Y itself. In the latter case, we appeal to the previous result, and conclude that unless Z has zero mass, then it is entirely contained in X which means its intersection with $\bar{X}$ is null. So it must be entirely contained in $Y \cap X$. Thus, any such Z with non-zero mass is either entirely contained in $\bar{Y}$, or entirely contained in $Y \cap X$.

Using the characterization of b by a mass function, we have:

$$b(X \vee \neg Y) = \sum_{U \subset (X \cup \neg Y)} m(U) = \sum_{U \subset (Y \cap X)} m(U) + \sum_{U \subset \bar{Y}} m(U) = b(Y \wedge X) + b(\neg Y)$$

The first line of the numerator for DS conditioning is $b(X \vee \neg Y) - b(\neg Y)$. What we've just seen is that so long as $b(Y) = p(Y)$, this becomes $b(X \wedge Y) + b(\neg Y) - b(XY) = b(X \wedge Y)$. This suffices to reduce DS conditioning to naive conditioning.

This result also allows us to show that mush-conditioning is equivalent to the other two under the same assumption. Note in particular that by the characterization of p and De Morgan's equivalences we have:

$$p(X \wedge \neg Y) = 1 - b(\neg(X \wedge \neg Y)) = 1 - b(X \vee \neg Y).$$

But our result above was exactly an alternative characterization of $b(X \vee \neg Y)$, and substituting this in and again using the definition of p, we get:

$$p(X \wedge \neg Y) = 1 - (b(Y \wedge X) + b(\neg Y)) = p(Y) - b(Y \wedge X).$$

Rearranging, this becomes:

$$p(Y) = p(Y \wedge \neg X) + b(Y \wedge X)$$

The right hand side here is the denominator of mush-conditioning. The left hand side, by the classicality of Y, is just $b(Y)$. But this is sufficient to reduce mush-conditioning to naive conditioning, as required.

References

Border, K. C. manuscript 2009. Separating hyperplane theorems. `www.hss.caltech.edu/~kcb/Notes/SeparatingHyperplane.pdf`. Accessed 14th September 2010.

Choquet, G. 1953. Theory of capacities. *Annales d'institute fourier*, **V**, 131–295.

De Finetti, Bruno. 1974. *Theory of probability: vol.1.* New York: Wiley.

Field, Hartry H. 2000. 'Indeterminacy, degree of belief, and excluded middle'. *Nous*, **34**, 1–30. Reprinted in Field, *Truth and the Absence of Fact* (Oxford University Press, 2001) pp. 278-311.

Fine, Kit. 1975. 'Vagueness, truth and logic'. *Synthese*, **30**, 265–300. Reprinted with corrections in Keefe and Smith (eds) *Vagueness: A reader* (MIT Press, Cambridge MA: 1997) pp.119-150.

Greaves, H., & Wallace, D. 2006. Justifying Conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, **115**(459).

Hájek, Petr. 1998. *Metamathematics of fuzzy logic.* Springer.

Halpern, J. Y. 1995. *Reasoning about uncertainty.* Revised edn. MIT Press. Revised paperback edition published 2003.

Jaffray, J-Y. 1989. Coherent bets under partially resolving uncertainty and belief functions. *Theory and decision*, **26**, 90–105.

Joyce, James M. 1998. 'A non-pragmatic vindication of probabilism'. *Philosophy of science*, **65**, 575–603.

Joyce, James M. 2009. 'Accuracy and coherence: prospects for an alethic epistemology of partial belief'. *Pages 263–297 of:* Huber, Franz, & Schmidt-Petri, Christoph (eds), *Degrees of belief*. Springer.

Keefe, Rosanna. 2000. *Theories of vagueness*. Cambridge: Cambridge University Press.

Leitgeb, H., & Pettigrew, R. 2010. An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy*. *Philosophy of science*, **77**(2), 236–272.

Lewis, David K. 1999. 'Why conditionalize?'. *Pages 403–7 of: Papers on metaphysics and epistemology*. Cambridge: Cambridge University Press. Written in 1972; introduction dated 1997. First published in this collection.

Paris, J.B. 2001. A note on the Dutch Book method. *Pages 301–306 of: Proceedings of the second international symposium on imprecise probabilities and their applications, isipta*. Ithaca, NY: Shaker.

Shafer, G. 1976. *A mathematical theory of evidence*. Princeton: Princeton University Press.

Skyrms, Brian. 2006. 'Diachronic coherence and radical probabilism'. *Philosophy of science*, **73**, 959–96.

Teller, P. 1973. Conditionalization and observation. *Synthese*, **26**(2), 218–258.

van Fraassen, Bas. 1966. 'Singular terms, truth-value gaps, and free logic'. *The journal of philosophy*, **63**(17), 481–495.