

Dutch Book Arguments and Consistency

Colin Howson

The London School of Economics and Political Science

1. Bayesian confirmation theory

Classical Bayesian methodology is based on the following three principles:

(i) individuals have degrees of belief which, measured in the closed unit interval, and subject to a mild consistency constraint, are formally probabilities.

(ii) belief functions are updated with the acquisition of new evidence by Bayesian conditionalisation. In other words, if B is learned to be true, then your new probability function P' takes the value $P'(A) = P(A|B)$ on every A in domain P' , where P is your probability function prior to learning B .

(iii) where H_i is a statistical hypothesis and E sample data, the terms $P(E|H_i)$ in Bayes' Theorem calculations are set equal to the probability assigned E by H_i .

More recent work has shown how some very strong assumptions implicit or explicit in (i)-(iii) can be relaxed while maintaining many of the distinctive elements of the Bayesian approach (Jeffrey 1987 gives a clear introduction to this neoBayesian theory, for much of which he is himself responsible). Most people regard (i)-(iii) as a reasonably good working theory nevertheless, and judging by the literature many still believe that at any rate (i) and (ii) are backed by a type of argument known as a Dutch Book argument, which is fundamentally sound. A Dutch Book argument shows that if certain constraints, e.g., those imposed by the probability axioms, are not met by a set of betting quotients then it is possible to find a betting strategy which will generate a certain loss or gain. Following Christensen (1991), I shall compare a popular Dutch Book argument for (ii) unfavourably with that for (i). I shall conclude by suggesting one for (iii).

2. Personal Probabilities

A Dutch Book argument for (i) typically commences with the following preamble. A natural measure of your degree of belief in a proposition A is the quantity $R/(R+Q)$ which you think makes the odds $R:Q$ fair in a bet on A . If S is the sum of what the bettor-on wins if A is true (kQ for some k) and loses if A is false (kR), then the bet is representable in the following form which has become standard in these discussions since it was introduced by de Finetti (1937):

PSA 1992, Volume 2, pp. 161-168

Copyright © 1993 by the Philosophy of Science Association

A	Bettor-on gets
T	$S(1-p)$
F	$-Sp$

where $p = R/(R+Q)$ is called the *betting quotient on A*, and S is the *stake*. Your 'fair' value of p is taken as the measure of belief rather than the odds R/Q because the latter is a semi-infinite scale not symmetrical about the point of indifference between A and $\neg A$.

To say that a particular value p_0 is your fair betting quotient on A is usually taken to mean that you believe that p_0 equalises the advantage in the bet, whatever the stake, and that any other value of p gives a definite advantage to one or other side. That for any given A there is such a value of p is of course a very strong assumption, and almost certainly false outside rather narrow betting situations. A more realistic complication allows you an infimum of values which you think give the advantage to the bettor against, and a supremum of values which you believe give the advantage to the bettor on, and these are called your upper and lower probabilities of A : in effect this makes your probability function interval-valued, where the intermediate values are those about which you are agnostic in the advantage they might give to either side. I shall not pursue this line further here, because my purpose is to evaluate the arguments which are alleged to justify (i) and (ii) *given* that your beliefs are point-valued.

Now we move to *sets* of degrees of belief p_1, \dots, p_n in propositions A_1, \dots, A_n , and considerations of their joint consistency. p_1, \dots, p_n are said to be consistent (or *coherent*) if there is no betting strategy, represented by a set of *signed* stakes $\pm S_1, \dots, \pm S_n$, which results in a loss or gain independently of the truth-values of A_1, \dots, A_n . For given a plausible additivity assumption, a positive gain or loss independently of the truth-values of the A_i contradicts the assumption that all the p_i are fair.

Thus we arrive at the fundamental theorem: p_1, \dots, p_n are consistent if and only if they are the values of a probability function. The 'only if' is known as the Dutch Book theorem; the converse, the 'if', is easily shown by considering expectations. This result is one of the standard Bayesian justifications of (i), and granted its assumptions, which if idealisation is regarded as permissible are not outlandish, it is an impressive one.

3. Bayesian conditionalisation

But does that justification extend to principle (ii)? At first sight not. As Ian Hacking was the first to point out (Hacking 1967), even if my updated belief function P' is not obtained by conditionalisation from P when some proposition B is learned, as long as P' itself is a probability function no betting strategy based on its values is vulnerable to a Dutch Book. However, such was—and is—the popularity of Dutch Book arguments that ways were sought round the apparent impasse. One which commands much support is as follows. True, there is no vulnerability to a Dutch Book in selecting P' as your updated function once you have learned B , however you do it. But suppose you update according to a rule *specified in advance*: in other words, suppose you have a rule for updating which dictates how $P(\cdot)$ will change in the event of any member of some partition $\{B_i\}$ of future possibilities turning out true (this is what Brian Skyrms (1993) refers to as an epistemic strategy).

Given that you update according to a rule, it is straightforward to show that if the rule is not that of conditionalisation, there is a betting strategy based on your current fair betting rates and your updating rule which leads to an inevitable loss. Paul Teller (1973) gave the first demonstration, attributing it (and the idea of an updating rule) to David Lewis. Here is a simplified version of their Dutch strategy. Let B be some member of the partition of future possibilities such that $P'(A)$ is the fair betting quotient your updating rule tells you you will have if B is true, and where $P'(A) \neq P(A|B)$;

suppose in fact that $P(A|B) > P'(A)$. Let $p = P(A|B)$ and $q = P(B)$. Suppose now you were to (I) buy a conditional bet paying 1 on $A|B$ at your current fair price p , and (II) buy a bet on B paying $p - P'(A)$ for your fair price $q(p - P'(A))$. If B is false you lose $q(p - P'(A))$. Suppose you learn that B is true, and you now (III) sell a bet paying 1 at your planned fair price $P'(A)$. Totting up, we see that whatever the truth-values of A and B , you lose $q(p - P'(A))$. If $p < P'(A)$ reverse all the bets. Hence there is a betting strategy based only on your own current fair betting rates and your updating rule which guarantees you a certain loss if $P'(A) \neq P(A|B)$. QED

If that seems to you to clinch conditionalisation as a consistency constraint on a par with the synchronic probability axioms, consider the following story.

4. Conditionalisation-breaking

You are certain that A is true, but you have grounds for believing that you are suffering from an incipient brain lesion which will make you doubt what you now feel to be certain (or you're Descartes thinking the demon's on the way). Since in practice you can distinguish your degrees of belief only to within an interval, we can suppose that there are k numbers r_1, \dots, r_k such that the set $\{P_t(A) = r_i: i=1, \dots, k\}$ represents your discriminable degrees of belief in A at some future time t —say tomorrow. Let B be the proposition $P_t(A) = r$, where r is one of the r_i less than 1. So, given the assumptions, $P(B) > 0$ and $P(A) = 1$. It follows that if you are consistent $P(A|B) = 1$. Also if you are consistent we can find out your updating rule. For consider what happens if B is true, i.e., if $P_t(A)$ actually is equal to r . You presumably can only know this at t . But since B is true this means that your probability function P_t at t must set $P_t(A) = r$. So consistency demands that your updating rule on learning B sets your updated probability $P'(A)$ of A equal to r . But $P(A|B) = 1$, and so your updating rule infringes conditionalisation.

Now there is nothing at all incoherent about that updating rule. On the contrary, it is forced on you by impeccable *deductive* reasoning. So either there is something wrong with the conclusion of the Lewis-Teller dynamic Dutch Book argument, that premediated nonconditionalisation displays inconsistency, or there is indeed an inconsistency in your beliefs, but one inherited from your initial assignment $P(A) = 1, P(B) > 0$.

Despite the fact that that initial assignment appears to represent—and, I claim, does represent—a perfectly reasonable state of belief, it has been claimed to be inconsistent by van Fraassen, since its consequence $P(A|B) = 1$ infringes the condition $\forall t \geq 0 [P(A|P_t(A) = r) = r, P_0 = P]$. He shows (1984) that any assignment infringing this condition, which he calls the Reflection Principle, is vulnerable to the same Dutch strategy used by Lewis-Teller. For suppose you were to make $p = P(A|P_t(A) = r)$ different from r . Simply replace $P'(A)$ in bet (II) above by r , and note that in bet (III) $P'(A)$ is also r , since if $P_t(A) = r$ is true then r is your fair betting rate at t on A . So you lose come what may.

But to conclude that your assignment is inconsistent is to assume what is being disputed, that any assignment which is vulnerable to the Lewis-Teller Dutch strategy is *eo ipso* inconsistent. And there is no independent reason to think there is anything wrong with the Reflection-Infringing assignment; on the contrary. But if the Lewis-Teller Dutch Book doesn't show the victim to be inconsistent, whereas the Dutch Books against those who infringe the 'synchronous' probability axioms do, what is wrong with the Lewis-Teller one?

5. Dynamic Dutch Book arguments are unsound

David Christensen (1991) seems, remarkably, to have been the first to give the correct answer to the question. I say 'remarkably', since the point is really an elementary

one. Suppose today you believe that A is true, but tomorrow $\neg A$. Is your set of beliefs inconsistent? Not if by this you understand the set of statements you believe true at one and the same time, for both sets $\{A\}$, $\{\neg A\}$ of synchronously accepted statements are clearly consistent if A is.

Similar observations apply to degrees of belief cashed out as fair betting rates. If my current fair betting quotients are such that combined with a certain betting strategy they would incur loss or gain independently of the truth-values of the propositions bet on, then I have reason to believe that I have erred in supposing them all fair. For this reason the Dutch Book arguments for the synchronic probability axioms are sound and do in my opinion establish those principles as genuine consistency constraints. In the Lewis-Teller argument, however, the dynamic Dutch Book shows that a set of betting rates you now accept as fair *plus one which you will accept as fair if B is true*, together generate certain loss. But there is no reason at all to suppose this to indicate incoherence in your beliefs unless the predicted new probability on A in the event of B 's being true is also one which you would currently accept now—and in our counterexample to conditionalisation you certainly would not accept $P_t(A)$, where that is less than one, as describing your current beliefs about A .

Can we salvage the Lewis-Teller argument by restricting it to updating rules in which your updated probability of $P'(A)$ of A , should B be true, is equal to your *current* fair betting rate on A given B ? Yes, but at the cost of making it trivial, for your current fair betting rate on A given B is by *definition* equal to your conditional probability $P(A|B)$, and you don't need to prove a definition. It is, of course, *not* trivial that coherence requires that $P(A|B)$ should be equal to $P(A \& B)/P(B)$, which de Finetti showed in his famous 1937 paper. Interestingly, the form of the Lewis-Teller argument as presented in Teller 1973 is to show that the quotient $P(A \& B)/P(B)$ is equal to $P'(A)$, which raises the question whether Lewis and Teller really were just independently re-proving de Finetti's result. Be that as it may, the fact remains that either we trivialise the Lewis-Teller argument, or we concede that it is unsound. Either way, it does not and cannot show that failure to conditionalise necessarily signifies incoherence in your beliefs.

The same verdict applies, *mutatis mutandis*, to dynamic Dutch Book arguments for Jeffrey's rule of conditionalisation. Jeffrey's rule says that if your probability $P(B)$ of B shifts exogenously to any other value $P'(B)$, not just to 1 as in Bayesian conditionalisation, then your updated function P' must be such that

$$P'(A) = P(A|B)P'(B) + P(A|\neg B)P'(\neg B)$$

for all A in the domain of P (this clearly subsumes Bayesian conditionalisation as the special case corresponding to $P'(B) = 1$). Here again a Dutch Book strategy is possible only against a rule specifying in advance how the probability function will be updated in the eventuality of an exogenous change from P to P' on some partition $\{B_j\}$. A necessary and sufficient condition for Jeffrey's rule is that $P(A|B) = P'(A|B)$ and $P(A|\neg B) = P'(A|\neg B)$ for all A , so the nonconditionaliser's updating rule will violate one of those identities. Brad Armendt showed (1980) that a straightforward modification of the Lewis-Teller betting strategy would also penalise a Jeffrey nonconditionaliser short-sighted enough to accept the relevant bets.

And of course, you would be short-sighted, or merely foolish, to engage in the Lewis-Teller betting strategy if you were not a conditionaliser. But that is only the same sort of folly as advertising a willingness to bet as directed on a proposition at two different betting rates, one you hold fair now and another you will, for whatever reason, hold fair tomorrow. And while that would indeed lose you money, it is no reflection on the consistency of your beliefs: you are certainly not inconsistent in hold-

ing one betting rate fair today and another fair tomorrow: you will have just changed your set of accepted fair betting rates, not incorporated two different ones into the set which you at any one time believe fair.

6. Is conditionalisation valid in typical Bayesian applications?

Having an updating rule which conflicts with conditionalisation, either of the Bayesian or the Jeffrey kind, is not, therefore, to be necessarily inconsistent. We have seen that there are sensible probability assignments where conditionalising would itself create incoherence. On the other hand, these are assignments to propositions which describe your future states of belief. This suggests that conditionalisation might have some validity where such propositions are not involved—in standard scientific contexts, for example. I think this is true. We understand by your unconditional personal probability of A your idea of the fair betting quotient on A relative to the stock of information K you possess at the time. Your conditional probability $P(A|B)$ on A relative to B is your idea of the fair betting quotient on A relative to $K \cup \{B\}$. It sounds very sensible to say that when you come to know B and no more then you should adopt $P(A|B)$ as your new betting quotient. In the examples we have looked at you are prohibited from doing this because B deductively entailed that your new degree of belief in A differs from $P(A|B)$.

Suppose however that neither B nor the statement 'I know B' has deductive consequences about your own states of belief. Now the only reason why you should fail to conditionalise is that on learning B you change your mind about the fair betting rate on A relative to $K \cup \{B\}$. You may well do so: nobody can always foresee exactly how they will or should react to new information. We are not perfect reasoners, and cannot see all the consequences of adopting an intellectual position: we discover new ones as we go on, particularly under the stimulus of fresh information. But if you were an ideal reasoner then conditionalisation becomes a trivial consequence of your status, given the proviso that the prospective information is not about your own states of belief.

As Jeffrey has pointed out, the conditions in which Bayesian conditionalisation is valid can be formulated purely probabilistic terms (1987, p.80). It is a simple consequence of the probability axioms that if $P'(B) = 1$ (you learn that B is true), and $P'(A|B) = P(A|B)$ (learning that B is true doesn't change your assessment of how probable A is given B), then $P'(A) = P(A|B)$. Now when $P(A|B) = 1$, and B is the proposition $P_t(A) = r$, $r < 1$, and $P(B) > 0$, as in our example earlier, we can actually demonstrate that the equality $P(A|B) = P'(A|B)$ fails. For we know that $P(A|P_t(A) = r) = 1$, while the synchronic Reflection principle $P_t(A|P_t(A) = r) = r$ will also hold if the agent is consistent (the synchronic Reflection principle, unlike the diachronic one earlier, is established by the existence of a Lewis-Teller Dutch strategy against any infringement, since all the betting rates on which that strategy is based are all accepted as fair at one and the same time (t)).

To sum up: I believe we should view the rule of Bayesian conditionalisation as describing how an ideal reasoner responds to new information, once it has formulated its conditional probabilities. The synchronic rules of the probability calculus and the diachronic rule of conditionalisation, properly restricted, should in other words be viewed as no less than a programme for an idealised inference machine. This seems to be how the pioneers of the Bayesian theory viewed it. How else are we to account for the fact that no adverse comment was ever made, until recently, on the apparently superhuman demands made by the theory for its correct application, and that conditionalisation had apparently never been regarded as some sort of independent postulate requiring separate justification?

7. Jeffrey's rule.

Jeffrey conditionalisation is not so easy to argue for as Bayesian conditionalisation, even in the restricted conditions set out. Where the change from P to P' originates in B , we know that a necessary and sufficient condition for the Jeffrey rule is that for all A in the domain of P , $P(A|B) = P'(A|B)$ and $P(A|\neg B) = P'(A|\neg B)$. While it seems plausible that if you've done your homework properly the mere change in B 's probability shouldn't cause the conditional probabilities to change, by analogy with what we know about conditional propositions, that is hardly a knock-down drag-out argument—especially as we know from David Lewis's justly celebrated result (1976) that probabilities of conditionals are conditional probabilities only in the most exceptional circumstances.

8. The Principal Principle

What about (iii)? Following David Lewis I shall call it the Principal Principle. It can be expressed in the equation

$$P(A|p(A) = r) = r,$$

where P is your subjective probability function and $p(A)$ is the objective probability of an event of type A . It is stipulated that were you to learn that $p(A) = r$ then that is the most information you would have relevant to the truth of A .

Clearly, before you can even start asking about the justification of the principle, you have to have some notion of what an objective probability is. I know of only one theory of objective probabilities which yields an acceptable justification for the principle. The theory is von Mises's, and the justification is a type of Dutch Book argument.

The virtues of von Mises's theory are in general too little appreciated, especially inside the circle of the Bayesian elect. That theory supplies, in its central notion of a collective, an elegant purely set-theoretical model of random sampling: that is to say, of families of independent, identically distributed random variables. Suppose, for example, that an experiment E generates collectives with common attribute space $\{0,1\}$ and limiting relative frequency p . Then any sequence of outcomes of the experiment E_n , consisting of repeating E n times, is representable by partitioning some sequence generated by E into successive segments of length n . It is a straightforward consequence of the Randomness and Convergence axioms defining a collective that if C is a collective generated by E , then the n -fold partitioning C_n of C is also a collective with attribute space $\{0,1\}^n$, in which the random variables X_1^n, \dots, X_n^n are independent with constant probability p of taking the value 1, where $X_i^n(s_n) = 1, \dots, n$, $s \in \{0,1\}^n$, is equal to 1 if the i th coordinate of s_n is 1, and 0 if not.

The explicit definition of probabilities as limits in collectives is usually held up as an objectionable feature of von Mises's theory, but in fact it is precisely that feature which distinguishes his theory as one in which inductive inferences about probabilities can be demonstrated to converge on the truth. Let Ω be the set $2^{\mathbb{N}}$ of all denumerably infinite sequences of 0s and 1s, and \mathcal{F} the Borel field of subsets of Ω generated by unions of cylinders. Let P be a countably additive subjective probability function on \mathcal{F} . Finally, let $H(\omega)$ state that $\omega \in \Omega$ is a collective with a specified limiting relative frequency of the attribute 1. It follows from a well-known theorem (Halmos 1950 Theorem B p. 213) that the probability that H is true, conditional on an initial segment $e_n(\omega)$ of ω , tends to 1 if $H(\omega)$ is true, and 0 if not, except on a subset of Ω of probability 0. In other words, there is a non-empty subset of Ω —indeed, a subset of P -measure 1—on which you will in the limit correctly identify whether or not you are being presented with an initial segment from a specified collective.

This result does not of course tell you what your posterior probability distribution will look like at any finite stage in the gathering of evidence. In order to perform a Bayes's Theorem calculation of such a posterior distribution, the likelihood $P(e_n(\omega)|H)$ must be evaluated. This is, of course, the role of the Principal Principle. If $e_n(\omega)$ is a particular sequence s_n of 0s and 1s, with r 1s, the Principal Principle equates $P(e_n(\omega)|H)$ to $p^r(1-p)^{n-r}$, i.e., to the probability of the attribute s_n in a collective C_n derived by partitioning from one with probability p of 1. The Principal Principle therefore does indeed generate the sorts of posterior distribution over a binomial parameter which Bayes himself discovered.

The principle is justified in the context of von Mises collectives by a type of Dutch Book argument (Mellor 1971 employs a very similar argument, though not for von Mises probabilities). Suppose all you know about E is that it is generated by a collective-generating experiment, and that within such a collective the limiting relative frequency of E is q . If you were to claim that in such conditions a betting rate smaller than q on E were not favourable then this should be true at any and therefore each repetition of those same conditions. So your claim will entail that after a finite number of identical bets, none of which is favourable by assumption, there will necessarily be a positive gain which will persist thereafter. There is therefore only one betting rate on E which can in the conditions specified be consistently claimed to confer no advantage to either party, and that is p .

It is at first sight surprising that this argument makes no use of the randomness property of collectives, only their convergence. But the randomness property is used, to compute the physical probabilities of events of the form 'a one appears r times out of n ', which are then identified by the Principal Principle with the appropriate subjective probabilities.

9. Conclusion

Dutch Book arguments have proved their worth in establishing core principles of the Bayesian theory. Even when the assumption of point-valued beliefs is relaxed, Dutch Book arguments are still used to establish the characteristic properties of the probabilistic machinery which takes their place.

But there are Dutch Book arguments and Dutch Book arguments. Not all reveal inconsistency, and as we have seen the Lewis-Teller one does not. Discarding it, we can (thankfully) discard the Reflection principle, but (equally thankfully) we need not discard conditionalisation. I have no space to discuss any of the other attempts to ground conditionalisation, like van Fraassen's symmetry argument (1989, pp. 322-324), or that based on the principle of minimum cross-entropy, though some reasons for doubting their efficacy appear in Howson and Urbach (1993). But perhaps the sort of limited justification I have offered here, which seems to me be substantially that offered by Jeffrey himself, will do.

References

- Armendt, B. (1980), "Is there a Dutch Book Argument for Probability Kinematics?", *Philosophy of Science*, 47: 583-588
- Christensen, D. (1991), "Clever Bookies and Coherent Beliefs", *Philosophical Review*, 229-247.

- de Finetti, B. (1937), "Foresight: its Logical Laws, its Subjective Sources", (English translation of "*La prévision: ses lois logiques, ses sources subjectives*") *Studies in Subjective Probability*, H.E. Kyburg and H. Smokler (eds), New York: Krieger, 1980.
- Hacking, I. (1967), "Slightly More Realistic Personal Probability", *Philosophy of Science*, 34: 311-325.
- Halmos, P. (1950), *Measure Theory*, New York: van Nostrand.
- Howson, C. and Urbach, P. (1993), *Scientific Reasoning: the Bayesian Approach* (second edition). Chicago: Open Court.
- Jeffrey, R.C. (1987), *Probability and the Art of Reasoning*. Cambridge: Cambridge University Press.
- Lewis, D. (1976), "Probabilities of Conditionals and Conditional Probabilities", *Philosophical Review*, 85: 297-315.
- Mellor, D.H. (1971), *The Matter of Chance*. Cambridge: Cambridge University Press.
- Skyrms, B. (forthcoming), "A Mistake in Dynamic Coherence Arguments"
- Teller, P. (1973), "Conditionalisation and Observation", *Synthese*, 26: 218-258.
- van Fraassen, B.C. (1984), "Belief and the Will", *Journal of Philosophy*, 81: 235-256.
- (1989), *Laws and Symmetry*, Oxford: the Clarendon Press.