What has gone wrong?

7. The gambler's fallacy is to regard it as less probable that there will be an A at the next trial if there has just been a long run of A's. Is this a fallacy? Is regarding another A as *more* probable equally fallacious?

CHAPTER 5 Subjective Probability

a DEGREES OF BELIEF AND THE PROBABILITY CALCULUS

a.1 Betting Quotients and Degrees of Belief

Our point of departure is the theory of betting odds. In actual betting practice, odds are non-zero numbers k which are offered by one party (the bookmaker) to be accepted or not by another (the punter). The odds are offered usually against the occurrence of some event E, and the punter nominates a sum Q such that he or she will contract to receive from the bookmaker the sum Qk if E occurs and forfeit Q if it does not.

In what follows we shall talk of the truth and falsity of hypotheses rather than the occurrence and non-occurrence of events. Our particular interest is going to be in those odds on a hypothesis h which you believe confer no positive advantage or disadvantage to either side of a bet on h at those odds, in the ideal world in which the bet is immediately and veraciously settled after the bet. We shall also suppose that these advantage-equilibrating odds are unique: values above or below would, you believe, confer advantage to one or other side. This is a strong idealising assumption; we shall consider what happens, in section **b**, when it is relaxed. For the time being suppose it holds. Such odds, if you can determine them, we shall call your *subjectively fair* odds on h.

This definition of subjectively fair odds does not presuppose that any odds are fair *in fact*. We shall discuss later the question of whether any odds are actually fair. We assume only that people do, rightly or wrongly, *think* that some odds are fair, and we believe this assumption to be borne out in the fact that people frequently bet. This is not, of course, to say that the odds they bet at are the ones they find fair. Usually this will not be the case, for most people bet only when they think the odds advantageous to them. But this does mean that they have a notion of advantageous and disadvantageous odds, and indeed in certain cases are capable of narrowing down the band between the odds they deem advantageous and those they think disadvantageous to a number correct to so many places of decimals. One way this quantity can be elicited is by asking people which they would prefer: a reward if the event in question occurs or that same reward if another event with agreed odds occurs, where the latter can be manipulated at will (we give an example, due to Lindley, in section **c.3** below).

We are less concerned with elicitation, however, than with the fact that there are subjectively fair odds there to be elicited, for it is this fact we shall exploit to provide a convenient measure of people's degrees of belief. For the odds you take to be fair on h will clearly reflect the extent to which you believe it likely that h will turn out to be true. Indeed, we would make your assessment of the fair odds on h the measure of your belief in h but for the inconvenient fact that on the odds scale, length of interval will not measure the *difference* between degrees of belief. The odds scale goes from 0 to plus infinity, with 1 as the point of indifference; hence the difference between being cognitively indifferent between h and $\sim h$ and being certain that $\sim h$ is true is 1, whereas the difference between being certain that h is true and being cognitively indifferent between h and $\sim h$ is infinite. The standard solution to the problem is to transform the semi-infinite odds scale, with ∞ appended, into the closed unit interval by means of the one-to-one mapping $p = \frac{k}{(1+k)}$. Odds of $\frac{1}{1}$, that is to say, even money odds, go to $\frac{1}{2}$ under this mapping; 0 goes to 0; and ∞ goes to 1, giving the desired symmetry about the point of indifference between h and $\sim h$.

The quantity $p = \frac{k}{(1+k)}$, where k are the odds on h you believe fair, *will therefore be taken as the numerical measure of your degree of belief in h.* p is called the *betting-quotient* associated with the odds k. Odds can be recovered uniquely from betting-quotients by means of the reverse transformation $k = \frac{p}{(1-p)}$.

Characterising degrees of belief in terms of characteristic odds or betting-quotients commenced with Ramsey (1931), and most authors have since followed him in making a willingness actually to bet in suitable circumstances the criterion of strength of belief (Ramsey, 1931, p. 79). This leads, as we shall see in section **c.1** below, to severe if not intractable problems when the question is posed why these behaviourally elicited quantities should satisfy the probability calculus.

We emphasise that we are *not* assuming that the intellectual judgment that odds are fair commits the judge to any behavioural display whatever. To believe odds fair is tantamount to believing the price of a gamble a fair one. But you can certainly believe a price fair without buying the good in question, and only in special circumstances would you actually do so. This may sound an obvious point, but it has been traditional in the literature to measure the strength of belief in terms of a willingness to bet at all odds up to some maximum value. Kyburg, for example, writes (1983, p. 64) that "The time-honored way of finding out how seriously someone believes what he says he believes is to invite him to put his money where his mouth is". But even equipped with enough capital to withstand betting losses, backing up judgments with financial commitment is not to everyone's taste, and declining to do so is no necessary indicator of belief.

Attempts to measure the values of options in terms of utilities are traditionally the way people have sought to forge a link between belief and action, and much contemporary Bayesian literature takes this as its starting point. We do not want to deny that beliefs have behavioural consequences in appropriate conditions, they clearly do, but stating what those conditions are with any precision is a task fraught with difficulty, if not impossible. Our view is that the fewer special—and questionable—assumptions that have to be made, the better, and the more secure the conclusions that one draws. Fortunately, we can derive our desired conclusion without assuming, or presuming, anything at all about the nature of the link between belief and action. For the conclusion we want to derive, that beliefs infringing a certain condition are inconsistent, can be drawn merely by looking at the consequences of what *would* happen if anyone *were* to bet in the manner and in the conditions specified.

a.2 Why Should Degrees of Belief Obey the Probability Calculus?

Following de Finetti (1937), we are going to assume a canonical form for bets between two individuals A and B as a contract whereby A pays the sum pS (dollars, pounds, or whatever) to B in exchange for the payment of the sum S if the hypothesis bet on is true, and 0 if it is not (we shall assume that S is arbitrarily finely divisible). The payoff conditions therefore look like this

h	*Payoff to A*
T	$S - pS$
F	$-pS$

where T stands for 'true' and F stands for 'false'. A is clearly betting *on* h at odds $pS{:}S - pS = p{:}1 - p$, and B is betting *against* h at the reciprocal odds $1 - p{:}p$; p can therefore be identified as the betting-quotient on h. In future when we refer to *a bet on h with betting-quotient p* we shall mean a contract of the above form. S is often called the *stake*. We can also speak of A buying from B a bet on a paying S for the price pS. Clearly, B strictly speaking needs no separate name; he or she is merely the other side of the bet.

Such bets can be brought into the traditional form described at the beginning of the chapter, given by the payoff table

h	*Payoff to A*
T	Qk
F	$-Q$

where $k = \frac{(1-p)}{p}$, by writing $Q = pS$. We use the de Finetti (S,p) representation for bets rather than the (Q,k) one since our focus of interest is p rather than k, and the constraints to be imposed on p emerge more simply in that formalism.

Now define a *betting strategy* with respect to a set of hypotheses $\{h_1, h_2 \ldots\}$ to be a set of instructions of the form 'bet on (against) h_i', for each i. Suppose that $p_1, p_2, \ldots$ is a

set of betting-quotients on the h_i. A celebrated theorem, proved independently by F. P. Ramsey and B. de Finetti, shows that

> If the p_i do not satisfy the probability axioms, then there is a betting strategy and a set S_i of stakes such that whoever follows this betting strategy will lose a finite sum whatever the truth-values of the hypotheses turn out to be.

The Ramsey-de Finetti theorem is often also called the *Dutch Book Theorem,* because a Dutch Book is a system of stakes which ensures a net loss.

The significance of the theorem lies in its corollary that betting-quotients which do not satisfy the probability axioms cannot consistently be regarded as fair. For (i) fair odds have been characterised as odds which offer zero advantage to either side of a bet; (ii) the sum of finitely (or even denumerably) many zeros is zero; hence the net advantage of a *set* of bets at fair odds is zero; and, finally, (iii) if a particular betting strategy is *assured* of a positive net gain or loss for whoever adopts it, then the net advantage in betting at the odds involved cannot be zero. We conclude that the assurance of a net gain or loss from finitely many simultaneous bets implies that they cannot all be fair. It follows immediately that *if your degrees of belief are measured by the betting-quotients you think fair, then consistency demands that they satisfy the probability axioms.* Thus agreement with the probability axioms is a necessary condition of consistency; in section **a.6** below we shall show that it is also sufficient.

a.3 The Ramsey–de Finetti Theorem

Ramsey's and de Finetti's theorem involves only elementary algebra and is very simple to prove, as we shall now show (the proof we give here owes much to Skyrms [1977, Ch. VI]). For each axiom of the probability calculus we shall show how its infraction entails the existence of a betting strategy leading to a necessary loss for one of the bettors.

(i) Axiom 1. Suppose that $p < 0$ and that you buy a bet on a proposition a paying one dollar, for the price p. Clearly, you will make a sure gain of $1 + |p|$ if a is true, and $|p|$ if a is false. Hence your fair betting-quotient on a must be non-negative.

(ii) Axiom 2. Suppose that you buy a bet on a tautology t paying one dollar for a price p. If $p < 1$, then you will make a certain gain of $1 - p$; if $p > 1$, then you will make a certain loss of $p - 1$. So the only fair betting-quotient on t is 1.

(iii) Axiom 3. Suppose that you buy bets on two mutually exclusive propositions a and b, each bet paying one dollar, for the prices p and q respectively. Then your net gain is as below (remember that a and b cannot both be true):

a	b	net gain
T	F	$1 - p - q = 1 - (p + q)$
F	T	$-p + 1 - q = 1 - (p + q)$
F	F	$-p - q = -(p + q)$

This diagram is clearly equivalent to the following:

a **v** b	net gain
T	$1 - (p + q)$
F	$-(p + q)$

Thus your separate bets on a and b determine a bet on the disjunction a **v** b paying one dollar and with betting-quotient $p + q$. Were you now also to bet against that disjunction with a betting-quotient r not to equal $p + q$, where the stake is also one dollar, then you will have a net gain of $r - (p + q)$ (positive or negative) *whatever the truth-values of* a *and* b. For if the first two bets are labelled (i) and (ii), and the bet against the disjunction is (iii), then the net gain from (i) + (ii) + (iii) is as below:

a **v** b	(i)+(ii) + (iii)	
T	$1-(p + q) - (1 - r)$	$= r - (p + q)$
F	$-p + q + r$	$= r - (p + q)$

Hence if your fair betting-quotient on a is p and on b is q, your fair betting-quotient on the disjunction can only be $p+q$, and we have proved the additivity axiom.

Before we turn to the remaining axiom, that of conditional probability, we shall show that the same type of argument requires not merely finite but also countable additivity. Consider a class of mutually exclusive hypotheses h_i, $i = 1,2,3,\ldots$. Suppose that a unit stake is placed on each of the 'even-number' hypotheses h_{2i} and that you bet on all these hypotheses simultaneously, with betting-quotients p_2, p_4, etc. If the infinite 'disjunction' of those hypotheses is true, then exactly one of them, h_{2j} say, is true, and the net gain is $-p_2 - p_4 \ldots + (1-p_{2j}) - \ldots = 1-(p_2+ \ldots +p_{2n}+ \ldots)$, which is independent of j. Hence if h is true, the net gain from all these bets is $1 - (p_2+ \ldots +p_{2n}+ \ldots)$. If h is false, then you lose the quantity $p_2+ \ldots +p_{2n}+\ldots$. So a set of simultaneous bets on all the h_{2i} with the same stake on each is equivalent to a bet on h with betting-quotient $(p_2+p_4+ \ \ldots)$. The fair betting-quotient on h must equal $(p_2 + p_4+ \ \ldots)$. QED.

There are, however, vigorous critics of the thesis that subjective probabilities are countably additive. De Finetti, for example, has produced many counter-arguments. To reassure the reader that we are not dismissing out of hand these objections from someone whose authority is certainly not to be considered lightly, let us consider briefly one of the most seductive of these counter-arguments.

This considers the example of a positive integer chosen 'at random'. It might seem natural in these circumstances to require a uniform, zero, degree of belief in each integer being selected. This is quite consistent with finite additivity, but not countable additivity, as we saw in Chapter 2, section **h.** But, as Spielman (1977) points out, it is not at all clear what selecting an integer at random could possibly amount to: any actual process would inevitably be biased toward the 'front end' of the sequence of positive integers, and so there is in reality little force in de Finetti's counter-example. Let us now move on to consider the remaining probability axiom, axiom 4.

a.4 Conditional Betting-Quotients

Axiom 4 we shall take to impose a condition on so-called *conditional betting-quotients*. A conditional betting-quotient is a betting quotient for a conditional bet, where a conditional bet on a given b is a bet on a which is to proceed in the event of b's turning out true and is called off if b is false. We imagine a scenario in which the truth-value of b is announced as soon as the contract has been made. The payoff conditions for the bettor-on in such a bet, with conditional betting quotient p and stake S, are therefore:

a	b	payoff
T	T	$S(1-p)$
F	T	$-pS$
	F	0

We shall define your conditional degree of belief in a given b to be the betting rate you think fair in a conditional bet of this type. We can interpret this in a possibly more illuminating way as follows. Your degree of belief in a proposition c is what you believe the fair betting-quotient on c to be. This is less a personal statement about yourself than a claim about which betting-quotient you believe to be fair relative to the information which you happen to possess. So we can gloss your conditional degree in a given b to be what you believe the fair betting-rate on a *would be* relative to the same information stock augmented by the additional information consisting of the statement that b is true. Note that this is not the same as saying that your conditional degree of belief in a given b is what you now believe the fair betting-quotient on a would be were you to come to know b in addition to what you already know, and no more (as we erroneously stated in the first edition of this book).

It is tempting to think of a conditional degree of belief in a given b as a degree of belief in a conditional 'proposition' $a \mid b$. The temptation should be resisted. We shall show in this section that consistent conditional degrees of belief, as we have defined them, are formally conditional probabilities, and David Lewis (1976) has shown that the usual rules of the probability calculus will not permit an interpretation of a conditional probability as the probability of a conditional sentence, even a non-truth-functional one.

We shall now proceed to show that axiom 4 of the probability calculus is a consistency condition for conditional degrees of belief as defined. In particular, we shall show that if axiom 4 is not satisfied, then there is a betting strategy involving conditional bets which will lead to an inevitable loss for one party.

The proof proceeds by showing that bets on a suitable combination of hypotheses determine some other bet, in this case a conditional bet. To be precise, we shall show that by setting appropriate stakes on b and a & b, simultaneous bets on those two statements are equivalent to a bet on a conditional on b, and that any odds placed on b and a & b can therefore be made to determine the odds for a bet on a conditional on b.

Suppose your fair betting-quotients on a & b and b are q and r respectively, where $r>0$. Suppose you were to bet at these rates on a & b with stake r and against b with stake q. Your net payoff is as follows:

a & b	b	net payoff
T	T	$r(1-q) - q(1-r) = r(1-\frac{q}{r})$
F	T	$-rq - q(1-r) = -q = -r(\frac{q}{r})$
	F	$-rq + qr = 0$

But this is clearly the payoff matrix of a bet on a conditional on b, with stake r and conditional betting-quotient $\frac{q}{r}$, *i.e., the ratio of the betting quotients q on a & b and r on b.* As with two mutually exclusive hypotheses, therefore, simultaneous bets with appropriate stakes also determine a further bet—in this case, a conditional one. Hence, if you were to state a fair conditional betting-quotient which differed from $\frac{q}{r}$, you would implicitly be assigning different conditional betting-quotients to the same hypothesis.

It does not follow, however, that you would necessarily make a positive net loss by buying a bet-on at your dearer price and selling one at your cheaper, with the same stake. For b may turn out to be false, whereupon the net gain from all the bets would be zero; the net gain is only non-zero if b is true. Nevertheless, anyone who believes that the betting-quotients q,r, and the conditional betting-quotient $p \neq \frac{q}{r}$, are all fair is no less inconsistent in that belief; indeed, it is quite easy to show that by suitably extending A's bets, a non-zero (positive or negative) net gain is assured whether b turns out to be true or false.

For suppose you were to bet (i) on a & b with stake r, (ii) against b with stake q, (iii) conditionally against a given b, with stake r, and finally (iv) on b with stake $q - pr$. As before, suppose your fair betting-quotient on a & b is q, on b is r, and on a given b is p. Bets (i) and (ii) above determine, as we saw, a conditional bet on a given b with stake r and betting-quotient $\frac{q}{r}$. Taking on bet (iii) simultaneously with (i) and (ii) guarantees, as we also saw, a net gain of $pr - q$ if b is true, with zero gain if not. It is straightforward to work out that making bet (iv) simultaneously with all the others guarantees an overall net gain (positive or negative) equal to $r(pr - q)$ whatever the truth-values of a and b. Given $r > 0$, this will be zero if and only if $p = \frac{q}{r}$, i.e., if and only if $P(a \mid b) = \frac{P(a \,\&\, b)}{P(b)}$. (Note, if you have not already done so, that a positive net gain can be turned into a positive net loss of the same magnitude by reversing the direction of all the bets.)

This completes the proof that if a set of betting-quotients does not satisfy the probability calculus, then they cannot all be fair (in **a.6** below we shall prove a form of converse to this). As we pointed out earlier, this result is independent of any formal characterisation of fairness of odds beyond the stipulation that they confer no advantage to either side of a bet at those odds. To round off the discussion, we shall now consider a particular method, used since the eighteenth century, of *computing* the advantage to taking a particular side in a bet.

a.5 Fair Odds and Zero Expectations

Laplace (1820, p.20) defined the *advantage* to taking a given side in a wager to be the expected value of the bet. Thus advantage, so defined, is calculated in the same units as the stake S, and so can be subjected to straightforward arithmetical operations, like taking sums of separate advantages. Carnap, Laplace's twentieth-century successor, calls that same expected value the *"estimated gain"* (1950, p. 170), and a bet *fair* just when the "estimated gain" is zero, where the expectation is computed relative to an appropriate Carnapian c-function.

Assuming your fair betting-quotients are consistent, a bet on a with stake S is formally a random variable X_a which takes the value $S(1 - p)$ if a is true and $-pS$ if not, where p is your fair betting-quotient on a. A bet against a with the same stake is $-X_a$. Simultaneously making bets on or against n hypotheses is the arithmetical sum of the corresponding random variables. Thus if we explicitly define the *advantage* of the bet represented by X_a to be its *expected value*, relative to your subjective probability distribution, then we deduce as theorems (i) that the advantage, *as you see it*, of betting at odds determined by your degree of belief is zero, and (ii) that the advantage attached to a betting strategy, as we defined it in the previous section, is the sum of the advantages of each of the bets separately which comprise that strategy (because the expectation of a sum of random variables is equal to the sum of their expected values). (i) is very easily seen, since $E(X_a) = S(1 - p)p - pS(1 - p) = 0$.

We have, in other words, found a mathematical representation of the informal notion of advantage which yields as a consequence the results that degrees of belief are subjectively fair betting-quotients and that the net advantage to placing n bets is the sum of each separately. These results do not of course prove anything substantially new. They merely show that the informal notion of subjective fairness can, to use Carnapian terminology, be given a formal explication which preserves all the desired consequences.

a.6 Fairness and Consistency

We have laid a foundation for a theory of consistent degrees of belief, characterised as subjectively fair odds, whose methodological consequences we shall explore in the subsequent chapters. A natural question to arise at this point is whether there are any odds other than those on tautologies and contradictions which are in some clear and objective sense fair. One candidate for a criterion of objective fairness was, as we have seen, having zero expectation relative to a 'logical' probability distribution of the type Laplace, Keynes, and Carnap tried to define. We have seen that their attempts foundered on the rock of pure arbitrariness. However, there is famously an alternative criterion: *odds are fair when they are determined by the real physical probabilities of the events concerned, where* *those probabilities exist.* We believe that, with certain qualifications, this claim is true, and indeed we shall base our theory of statistical inference on it. But any argument for that thesis must await a discussion of the notion of physical probability itself, a notion which, as we shall see, is fraught with difficulties. We shall take up that discussion again in Chapter 13.

Ramsey (1931) used the term "consistent" to characterise degrees of belief having the formal structure of probabilities. We have shown in sections **a.2** and **a.3** that your system of beliefs is inconsistent if the betting quotients you believe fair do not satisfy the probability axioms. But what about the converse—are we justified in claiming that your belief system is consistent if the betting-quotients you believe fair do satisfy the probability axioms? This would amount to the claim that if P is a set of betting-quotients over a set H of n hypotheses, and if P satisfies the probability axioms 1–4, then for any betting strategy and any system of stakes, it is not the case that for every truth-value distribution over the members of H the net gain is uniformly negative (or positive: remember that a negative net gain can be transformed into a positive one by reversing the directions of the bets). For consider any set of bets with arbitrary stakes on or against each of the hypotheses in H. These, as we know from the previous section, are random variables $X_1, \ldots, X_n$ and if the value of their sum Y were always negative, say, then the expected value of Y would clearly be negative also. But as we also know, $E(Y) = \Sigma E(X_i)$, and $E(X_i) = 0$ for each i, by (i) of the previous section. Hence, if any set of betting quotients satisfies the probability calculus, then no betting strategy can generate a positive or negative gain come what may.

So, consistency for partial beliefs is equivalent to their being formally probabilities. Today it is usual, following de Finetti, to use the adjective 'coherent' to mean that partial beliefs satisfy the probability axioms. This seems to us to direct the attention away from the all-important logical fact that the probability calculus is a complete axiomatisation of consistent partial belief. The probability axioms, as Ramsey emphasised, do have therefore a purely logical interpretation; not, as Keynes and Carnap believed, as a calculus of partial entailment, but as the logic of consistent partial belief.

■ b UPPER AND LOWER PROBABILITIES

We promised that we would return to discuss those hypotheses and data sets where it might seem to be unrealistic to suppose that one would have point-valued degrees of belief. To borrow an example from Suppes (1981, p. 41): if we consider the question of whether it will rain at some specified time in Fiji, we can certainly suggest a value k_1 such that odds less than k_1 on that hypothesis are, in our opinion, unrealistically low, and we can also suggest odds k_2, such that odds greater than k_2 are unrealistically high. But we might also say that there is an intermediate interval of odds between which we feel quite unable to discriminate. The typical indefiniteness of one's knowledge would, it seems, be more faithfully reflected by an interval-valued function which only in certain cases takes degenerate intervals, or points, as values.

We believe that this suggestion reveals a confusion as to what subjective probabilities actually are. The whole point of introducing the apparatus of subjective probability is precisely because one's knowledge is typically indefinite: subjective probabilities *express* that indefiniteness by taking non-extreme values. Nevertheless, we have defined your subjective probability of h as the betting-quotient on h you believe to be fair in the present circumstances, and this does leave open the possibility that you may feel unable to specify an exact value. Indeed, the occasions on which you feel that you can specify a unique number with confidence may well turn out to be exceptions rather than the rule.

It turns out that very little is lost in conceding that what we have supposed to be point-valued degrees of belief are actually interval-valued, so long as the intervals are small. Suppose that $P_*(a)$ is the least upper bound (supremum) of all the betting quotients on a at which you definitely think a bet on a advantageous to the bettor-on, and $P^*(a)$ is the greatest lower bound (infimum) of betting-quotients at which you definitely think a bet on a would be advantageous to the bettor-against. For all intermediate values you have no opinion at all about the relative advantages of either side of the bet. $P_*(a)$ is called your *lower probability* of a and $P^*(a)$ is your *upper probability* of a.

We can define consistency for upper and lower probabilities analogously to consistency for point-probabilities. We then

SCIENTIFIC REASONING: THE BAYESIAN APPROACH

COLIN HOWSON

AND

PETER URBACH

Second Edition

. . . if this [probability] calculus be condemned, then the whole of the sciences must also be condemned.
—Henri Poincaré

Our assent ought to be regulated by the grounds of probability.
—John Locke

Open Court
Chicago and La Salle, Illinois

1996