Probabilities as betting odds and the Dutch book

Carlton M. Caves
2000 June 24
Substantially modified 2001 September 6
Minor modifications 2001 November 30
Material on the lottery-ticket approach added 2003 May 24
Minor modifications on 2005 July 20

I learned about the ideas discussed here from the principal article by Bruno de Finetti on subjective probabilities: "Foresight: Its Logical Laws, Its Subjective Sources," in *Studies in Subjective Probability*, edited by Henry E. Kyburg, Jr., and Howard E. Smokler (Wiley, New York, 1964), pages 93–158 [Translation by Henry E. Kyburg, Jr., of original article: "La prévision: ses lois logiques, ses sources subjectives," *Annales de l'Institut Henri Poincaré* **7**, 1–68 (1937)]. There has been, however, considerable reworking of the ideas in the following, based mainly on discussions with Rüdiger Schack.

Introduction. A powerful method for justifying subjective probabilities is to regard the operational definition of a probability p as being that one is willing to make a bet at odds of $(p^{-1}-1):1$. One is willing to put up a stake of one dollar with the chance of winning S dollars [$(S-1):1$ odds], where $1\le S\le\infty$. The expected gain is $\bar{G}=pS-1$; requiring a fair bet, i.e., $\bar{G}=0$, gives $p=1/S$.

It is convenient in what follows to let both the stake and the payoff vary, the payoff being S dollars and the stake being pS [odds of $S(1-p):pS=(1-p):p=(p^{-1}-1):1$]. The bettor, willing to make a bet at odds $(p^{-1}-1):1$, approaches a bookmaker who is accepting stakes and making payoffs. The bookmaker is free to set the payoffs. The objective is to show that a bettor must use the standard rules of probability theory; otherwise, the bookie can set the payoffs so that the bettor always loses, i.e., has strictly negative gains for all outcomes that he believes can occur. The bookie can make the payoffs either positive or negative. A negative payoff reverses the roles of bettor and bookie: the bettor receives $-pS$ dollars and pays out $-S$ dollars.

The bookie in this scenario is conventionally called the *Dutch book*, and the requirements on betting behavior imposed by avoiding sure losses in dealing with a Dutch book are called *Dutch-book consistency*. A Dutch book doesn't act like an ordinary bookie, who takes bets and adjusts the odds on all outcomes in an attempt to balance his gains and losses no matter what the outcome and derives his sure income by imposing a fee on all bets. Instead a Dutch book takes advantage of the inconsistent behavior of a bettor to arrange to win outright no matter what the outcome.

Formally we want to show that avoiding all-negative gains implies the standard probability rules:

1. Probabilities lie between 0 and 1 inclusive.
2. Certainty implies unity probability.
3. Probabilities of exclusive events add.
4. Bayes's rule.

These rules are sufficient to determine the entire structure of probability theory. We also want to demonstrate the converse, that the standard probability rules imply that the bettor avoids being forced to accept all-negative gains. The converse is trivial, so we save it till the end.

1. Probabilities lie between 0 and 1 inclusive. Suppose one bets on an event E and its complement $\neg E$ (not E). These two events are mutually exclusive and exhaustive (one or the other occurs, but not both). The stakes are $p(E)S_E$ and $p(\neg E)S_{\neg E}$, and the payoffs are S_E and $S_{\neg E}$. The bettor's gain should event E occur is

$$G_E = S_E - p(E)S_E - p(\neg E)S_{\neg E} \ , \tag{1}$$

and the gain should event E not occur is

$$G_{\neg E} = S_{\neg E} - p(E)S_E - p(\neg E)S_{\neg E} \ . \tag{2}$$

The Dutch book can always set $S_{\neg E} = 0$, in which case the gains become

$$\begin{aligned} G_E &= \big(1 - p(E)\big)S_E \ , \\ G_{\neg E} &= -p(E)S_E \ . \end{aligned} \tag{3}$$

Avoiding both gains being negative requires that $1 - p(E)$ and $-p(E)$ have opposite signs, which implies that

$$0 \le p(E) \le 1 \ , \tag{4}$$

as desired.

What happens if $p(E) > 1$ is that the bettor is willing to put up a positive stake $p(E)S_E$, which is lost if E does not occur and which exceeds the payoff S_E if E does occur. In contrast, if $p(E) < 0$, the bettor is willing to put up a positive stake $p(E)S_E$, which is lost if E does not occur, but which leads to a negative payoff S_E if E does occur.

2. Certainty implies unity probability. Continuing with same scenario, suppose that the bettor believes that E is certain to occur and thus that $\neg E$ is certain not to occur. The bettor's gain on the occurrence of E is

$$G_E = S_E - p(E)S_E - p(\neg E)S_{\neg E} = \big(1 - p(E)\big)S_E - p(\neg E)S_{\neg E} \ . \tag{5}$$

The Dutch book can arrange that this gain have any value, unless $p(E) = 1$ and $p(\neg E) = 0$. Indeed, he can obtain any negative value by choosing payoffs $S_E < 0$ and $S_{\neg E} > 0$. Thus the requirement that there be no payoffs that yield all-negative gains requires that an event that the bettor believes is certain to occur have unity probability and, in addition, an event that he believes is certain not to occur have zero probability.

What if we want to show that unity probability implies certainty? For this purpose, suppose that the $P(E) = 1$. Then the gains of Eqs. (1) and (2) can be written in the matrix form

$$\begin{pmatrix} G_E \\ G_{\neg E} \end{pmatrix} = \begin{pmatrix} 0 & -p(\neg E) \\ -1 & 1 - p(\neg E) \end{pmatrix} \begin{pmatrix} S_E \\ S_{\neg E} \end{pmatrix} \ . \tag{6}$$

Given any set of gains, including all-negative gains, the Dutch book can find a set of payoffs that yields those gains, unless the determinant of the matrix in Eq. (6) vanishes, which implies that $p(\neg E) = 0$. This conclusion in hand, the gains read

$$\begin{aligned} G_E &= 0 \;, \\ G_{\neg E} &= -S_E \;, \end{aligned} \tag{7}$$

and we can go no further. To get the desired result, we must strengthen our consistency requirement to read that the bettor should never put himself in a situation where on all outcomes he deems possible, he never wins, but sometimes loses, i.e., has gains all of which are nonpositive and some of which are negative. To avoid this situation, the bettor must believe that E is certain to occur, so as to rule out the loss on occurrence of $\neg E$.

3. Probabilities of exclusive events add. Suppose now that there are two mutually exclusive, but not necessarily exhaustive events, E_1 and E_2, and that event E represents the occurrence of E_1 *or* E_2, written formally as $E = E_1 \vee E_2$. The relevant bets, on E, E_1, and E_2, have probabilities $p(E)$, $p(E_1)$, and $p(E_2)$, and the associated stakes are $p(E)S_E$, $p(E_1)S_{E_1}$, and $p(E_2)S_{E_2}$. There are three cases: (i) neither E_1 nor E_2 occurs, in which case the bettor's gain is

$$G_{\neg E} = -p(E)S_E - p(E_1)S_{E_1} - p(E_2)S_{E_2} \;; \tag{8}$$

(ii) E_1 occurs, in which case the gain is

$$G_{E_1} = \big(1 - p(E)\big)S_E + \big(1 - p(E_1)\big)S_{E_1} - p(E_2)S_{E_2} \;; \tag{9}$$

and (iii) E_2 occurs, in which case the gain is

$$G_{E_2} = \big(1 - p(E)\big)S_E - p(E_1)S_{E_1} + \big(1 - p(E_2)\big)S_{E_2} \;. \tag{10}$$

Summarizing these relations in a matrix equation, we get

$$\begin{pmatrix} G_{\neg E} \\ G_{E_1} \\ G_{E_2} \end{pmatrix} = \underbrace{\begin{pmatrix} -p(E) & -p(E_1) & -p(E_2) \\ 1 - p(E) & 1 - p(E_1) & -p(E_2) \\ 1 - p(E) & -p(E_1) & 1 - p(E_2) \end{pmatrix}}_{=\,\mathtt{P}} \begin{pmatrix} S_E \\ S_{E_1} \\ S_{E_2} \end{pmatrix} \;. \tag{11}$$

For any set of gains, including all-negative gains, the Dutch book can find corresponding payoffs unless the determinant of $\mathtt{P}$ vanishes. Thus we have

$$0 = \det \mathtt{P} = p(E) - p(E_1) - p(E_2) \;, \tag{12}$$

which is the desired result,

$$p(E_1 \vee E_2) = p(E) = p(E_1) + p(E_2) \;. \tag{13}$$

4. Bayes's rule. We now consider conditional bets. Let B be an event, and let A be an event that is conditioned on B. The relevant bets are on (i) the occurrence of B, (ii) the occurrence of A, given that B has occurred, and (iii) the joint occurrence of A and B, written formally as $A \wedge B$. The corresponding probabilities are $p(B)$, the conditional probability $p(A \mid B)$, and the joint probability $p(A \wedge B) = p(A, B)$. Bets are placed on all three situations, with the agreement that if B does not occur, the bet on $A \mid B$ is called off, with the stake returned. There are three cases: (i) B does not occur, in which case the bettor's gain is

$$G_{\neg B} = -p(B)S_B - p(A,B)S_{A,B}\ ; \tag{14}$$

(ii) B occurs, but not A, in which case the gain is

$$G_{\neg A \wedge B} = \big(1 - p(B)\big)S_B - p(A,B)S_{A,B} - p(A \mid B)S_{A|B}\ ; \tag{15}$$

and (iii) both A and B occur, in which case the gain is

$$G_{A \wedge B} = \big(1 - p(B)\big)S_B + \big(1 - p(A,B)\big)S_{A,B} + \big(1 - p(A \mid B)\big)S_{A|B}\ . \tag{16}$$

Again we write a matrix equation relating the gains to the payoffs,

$$\begin{pmatrix} G_{\neg B} \\ G_{\neg A \wedge B} \\ G_{A \wedge B} \end{pmatrix} = \underbrace{\begin{pmatrix} -p(B) & -p(A,B) & 0 \\ 1 - p(B) & -p(A,B) & -p(A \mid B) \\ 1 - p(B) & 1 - p(A,B) & 1 - p(A \mid B) \end{pmatrix}}_{= \mathtt{P}} \begin{pmatrix} S_B \\ S_{A,B} \\ S_{A|B} \end{pmatrix}, \tag{17}$$

and again in order to avoid having payoffs with all-negative gains, we require that the determinant of $\mathtt{P}$ vanish:

$$0 = \det \mathtt{P} = -p(A \mid B)p(B) + p(A,B)\ . \tag{18}$$

Requiring that there be no all-negative gains thus yields Bayes's rule relating joint and conditional probabilities:

$$p(A \wedge B) = p(A,B) = p(A \mid B)p(B)\ . \tag{19}$$

This concludes the Dutch-book derivation of the rules of probability theory.

Normalization of probabilities. The normalization of probabilities follows trivially from the first three rules: considering again events E and $\neg E$, we have that

$$1 = p(E \vee \neg E) = p(E) + p(\neg E)\ . \tag{20}$$

Not so obvious is that we can substitute normalization for additivity of probabilities, because we can derive the additivity of probabilities from the other three rules plus normalization. To see this, suppose again that there are two mutually exclusive, but not necessarily exhaustive events, E_1 and E_2, and that $E = E_1 \vee E_2$. Given E, E_1 and E_2 are mutually exclusive *and* exhaustive events, so normalization requires

$$p(E_1 \mid E) + p(E_2 \mid E) = 1 \; . \tag{21}$$

Furthermore, given E_1 (or E_2), E is certain, so we know that

$$p(E \mid E_1) = 1 \; , \qquad p(E \mid E_2) = 1 \; . \tag{22}$$

Bayes's rule then gives

$$\begin{aligned} p(E_1 \mid E)p(E) &= p(E \mid E_1)p(E_1) = p(E_1) \; , \\ p(E_2 \mid E)p(E) &= p(E \mid E_2)p(E_2) = p(E_2) \; . \end{aligned} \tag{23}$$

Adding these two equations gives the additivity of probabilities for mutually exclusive events:

$$p(E_1 \vee E_2) = p(E) = p(E_1) + p(E_2) \; . \tag{24}$$

It should come as no surprise that we can get the normalization of probabilities directly from a Dutch-book argument. Returning to the events E and $\neg E$, we can write the gains of Eqs. (1) and (2) as a matrix equation:

$$\begin{pmatrix} G_E \\ G_{\neg E} \end{pmatrix} = \begin{pmatrix} 1 - p(E) & -p(\neg E) \\ -p(E) & 1 - p(\neg E) \end{pmatrix} \begin{pmatrix} S_E \\ S_{\neg E} \end{pmatrix} \; . \tag{25}$$

Given any set of gains, including all-negative gains, the Dutch book can find a set of payoffs that yields those gains, unless the determinant of the matrix in Eq. (25) vanishes. Thus we require that

$$0 = \det \begin{pmatrix} 1 - p(E) & -p(\neg E) \\ -p(E) & 1 - p(\neg E) \end{pmatrix} = 1 - P \; , \qquad P \equiv p(E) + p(\neg E) \; . \tag{26}$$

The requirement that there be no all-negative gains means that the probabilities must be normalized to unity:

$$P = p(E) + p(\neg E) = 1 \; . \tag{27}$$

It is worth examining in detail how the Dutch book manages to make both gains negative when the probabilities aren't normalized. Let $\tilde{p}(E) = p(E)/P$ and $\tilde{p}(\neg E) = p(\neg E)/P$ be a pair of normalized probabilities. For specificity, assume that $S_E \le S_{\neg E}$. Since

$$S_E \le \tilde{p}(E) S_E + \tilde{p}(\neg E) S_{\neg E} \le S_{\neg E} \; , \tag{28}$$

i.e.,

$$P S_E \le p(E) S_E + p(\neg E) S_{\neg E} \le P S_{\neg E} \; , \tag{29}$$

we have

$$\begin{aligned} S_E - P S_{\neg E} \le G_E &\le (1 - P) S_E \; , \\ (1 - P) S_{\neg E} \le G_{\neg E} &\le S_{\neg E} - P S_E \; . \end{aligned} \tag{30}$$

If $P = 1$, we have $G_E \le 0$ and $G_{\neg E} \ge 0$. If $P < 1$, however, the Dutch book can choose $S_E \le S_{\neg E} < PS_E < 0$, which makes both gains negative. Similarly, if $P > 1$, the Dutch book can choose $0 < S_E \le S_{\neg E} < PS_E$, again making both gains negative.

If $P < 1$, we can understand what is going on by saying that the bettor has mistakenly concluded that E and $\neg E$ are exhaustive, when in fact, there must be another event, say D, with probability $p(D) = 1 - p(E) - p(\neg E)$. No bet is placed on D—the stake and payoff are zero—because the bettor doesn't recognize the existence of D. The bettor's putative gain on occurrence of D is

$$G_D = -p(E)S_E - p(\neg E)S_{\neg E} \ , \tag{31}$$

which is positive when G_E and $G_{\neg E}$ are negative. When $P < 1$, the Dutch book has chosen negative stakes and payoffs for both E and $\neg E$, so the bettor is holding the stake $-p(E)S_E - p(\neg E)S_{\neg E} = G_D$. What is supposed to happen when D occurs is that the bettor retains this stake as his gain; since the bettor doesn't recognize the existence of D, however, the Dutch book declares that no trial has occurred and tries again till D does not occur.

The preceding argument can easily be generalized directly to many events. Suppose one bets on N mutually exclusive, exhaustive events E_j, $j = 1, \ldots, N$, with stakes $p_j S_j$ and payoffs S_j. If event E_j occurs, the bettor's gain is

$$G_j = S_j - \sum_k p_k S_k = (1 - p_j)S_j - \sum_{k \ne j} p_k S_k \ . \tag{32}$$

The gains can be written as a matrix equation

$$\begin{pmatrix} G_1 \\ G_2 \\ \vdots \\ G_N \end{pmatrix} = \underbrace{\begin{pmatrix} 1-p_1 & -p_2 & \cdots & -p_N \\ -p_1 & 1-p_2 & \cdots & -p_N \\ \vdots & \vdots & \ddots & \vdots \\ -p_1 & -p_2 & \cdots & 1-p_N \end{pmatrix}}_{\equiv \mathtt{P}} \begin{pmatrix} S_1 \\ S_2 \\ \vdots \\ S_N \end{pmatrix} \ . \tag{33}$$

Given any set of gains, including all-negative gains, the Dutch book can find a set of payoffs that yields those gains unless the determinant of the matrix $\mathtt{P}$ is zero. Thus we require that

$$0 = \det \mathtt{P} = \sum_{j_1, \ldots, j_N} \epsilon_{j_1 \ldots j_N} P_{1 j_1} \cdots P_{N j_N} \ , \tag{34}$$

where $\epsilon_{j_1 \ldots j_N}$ is the completely antisymmetric symbol. Since $P_{jk} = \delta_{jk} - p_j$, we obtain

$$\det \mathtt{P} = 1 - P \ , \qquad P \equiv \sum_{j=1}^{N} p_j \ , \tag{35}$$

all the higher-order terms in the determinant vanishing. Thus the requirement that there be no payoffs that yield all-negative gains enforces normalization of the probabilities:

$$1 = P = \sum_{j=1}^{N} p_j \; . \tag{36}$$

Probability of nonexclusive events. Consider two events, A and B, which are not necessarily exclusive. The events $A \wedge \neg B$ and B are mutually exclusive, so we can use the additivity of probabilities for exclusive events to write

$$p(A \wedge \neg B) + p(B) = p\big((A \wedge \neg B) \vee B\big) \; . \tag{37}$$

Using the distributive rule, we have

$$(A \wedge \neg B) \vee B = (A \vee B) \wedge (\neg B \vee B) = A \vee B \; . \tag{38}$$

Bayes's rule gives

$$p(A \wedge \neg B) = p(\neg B|A)p(A) = \big(1 - p(B|A)\big)p(A) = p(A) - p(B|A)p(A) = p(A) - p(A \wedge B) \; . \tag{39}$$

Putting this together, we get

$$p(A \vee B) = p(A) + p(B) - p(A \wedge B) \; . \tag{40}$$

Discussion. The Dutch-book argument shows that to avoid being put in a situation where he loses on all outcomes he deems possible, the bettor must use the standard rules of probability. To show the additional result that probability 1 implies that the bettor believes an event is certain to occur (or, equivalently, that probability 0 implies than an event is certain not to occur), we needed to use the slightly stronger requirement that the bettor should avoid being put in a situation where on outcomes he deems possible, he never wins and sometimes loses.

In demonstrating the converse, we only need to consider the situations of betting on mutually exclusive and exhaustive events and betting on conditional events. In the first of these, betting on N mutually exclusive and exhaustive events, the expected gain is

$$\bar{G} = \sum_j p_j G_j = \sum_j p_j S_j - \sum_j p_j \sum_k p_k S_k = (1 - P) \sum_j p_j S_j \; . \tag{41}$$

If the probabilities are normalized, we have $\bar{G} = 0$, which means that it is impossible for all the gains to be negative. In the second situation, betting on conditional events as in the section on Bayes's rule, the expected gain,

$$\begin{aligned}\big(1 - p(B)\big)G_{\neg B} &+ \big(1 - p(A \mid B)\big)p(B)G_{\neg A \wedge B} + p(A,B)G_{A \wedge B} \\ &= \Big(p(A,B) - p(A \mid B)p(B)\Big) \\ &\quad \times \Big(S_B + S_{A|B} - p(B)S_B - p(A,B)S_{A,B} - p(A \mid B)S_{A|B}\Big) \; ,\end{aligned} \tag{42}$$

is zero when Bayes's rule is satisfied, making it impossible to arrange that the bettor always loses.

Another approach. Another approach to the Dutch book argument is phrased in terms of contracts, or lottery tickets, associated with some set of events. The price at which one is willing to buy or sell a ticket defines one's personal probability for the associated event. The argument proceeds by saying that consistency requires that two different ways of dividing up the same event in terms of contracts must have the same price. I learned this approach from Brian Skyrms ["Coherence," in *Scientific Inquiry in Philosophical Perspective*, edited by N. Rescher (Center for Philosophy of Science, Lanham, MD, 1987), pp. 225–242].

If one is willing to buy or sell the lottery ticket "Pay \$1 if E" for $\$q$, this *defines* one's probability for event E to be $p(E) = q$. One gets immediately that probabilities are between 0 and 1 inclusive: if one is willing to sell the ticket for a negative amount, this means that one is willing to pay someone to take it off one's hands, leading to a sure loss; likewise, willingness to buy the ticket for more than \$1 guarantees a loss.

If one believes E is certain to occur, then one must assign unit probability to E, for otherwise one would be willing to sell the ticket for less than \$1, guaranteeing a loss. To get the converse, we need again to assume the stronger version of consistency. If one assigns unity probability to E, then one is willing to buy the ticket for \$1. Though this doesn't lead to a sure loss, it does guarantee that one can't win and definitely loses if E does not occur; to eliminate the possible loss, one must believe that E is certain to occur.

Additivity on mutually exclusive events now comes from the following argument. Letting E_1 and E_2 be mutually exclusive, with $E = E_1 \vee E_2$, we imagine the following three tickets: (i) "Pay \$1 if E" at price $\$p(E)$; (ii) "Pay \$1 if E_1" at price $\$p(E_1)$; (iii) "Pay \$1 if E_2" at price $\$p(E_2)$. Under all circumstances, holding the second two tickets is equivalent to holding the first, so their prices must be the same, i.e., $p(E) = p(E_1) + p(E_2)$. Putting it differently, if $p(E) > p(E_1) + p(E_2)$, buying the first ticket while selling the second two guarantees a loss, whereas if $p(E) < p(E_1) + p(E_2)$, selling the first ticket while buying the second two guarantees a loss. Provided one is willing to contemplate a party buying and selling a countably infinite number of tickets, this method can obviously be extended to a countable infinity of exclusive events, thus giving the property of *countable additivity* for probabilities.

To get Bayes's rule, consider two events, D and E, not necessarily mutually exclusive, and the event $D \wedge E$. We imagine the following three tickets: (i) "Pay \$1 if $D \wedge E$; pay $\$p(E|D)$ if $\neg D$" at price $\$p(E|D)$ (the second clause in the ticket cancels the ticket at no cost if D does not occur); (ii) "Pay \$1 if $D \wedge E$" at price $\$p(D \wedge E)$; (iii) "Pay $\$p(E|D)$ if $\neg D$" at price $\$p(E|D)p(\neg D) = \$p(E|D)\big(1 - p(D)\big)$ [this is clearly equivalent to a ticket "Pay \$1 if $\neg D$" at price $p(\neg D)$]. Under all circumstances, holding the second two tickets is equivalent to holding the first, so their prices must be the same, i.e.,

$$p(E|D) = p(D \wedge E) + p(E|D)\big(1 - p(D)\big) \qquad \Longrightarrow \qquad p(D \wedge E) = p(E|D)p(D)\ . \quad (43)$$

This approach has the advantage of being conceptually simpler than the standard Dutch book argument. Moreover, it highlights the fact that the inconsistent betting behavior that drives the Dutch book argument is really a fundamental logical inconsistency that fails to recognize the equivalence of two different ways of stating the same thing. What one loses in this approach compared to the usual one is the ability to show that inconsistent betting behavior opens the way to arbitrarily large losses.

A final point. It is important to stress that the only player who assigns probabilities in the Dutch book argument is the bettor. The Dutch book doesn't need to know anything about the events and doesn't need to have his own probabilities for those events. All the Dutch book is interested in is whether the bettor is willing to place bets that are inconsistent with the rules of probability theory. If so, the Dutch book can always arrange to win, without knowing anything about the events on which the bets are placed. Thus the Dutch book argument is not about adversarial behavior of the bettor and the bookie; rather it is wholly about the internal consistency of the bettor.