Notes on the Dutch Book Argument

by D. A. Freedman, Statistics Department, U. C. Berkeley, CA 94720

The object here is to sketch the mathematics behind de Finetti's (1931, 1937) argument for the Bayesian position. Suppose a bookie sets odds on all subsets of a set, accepting bets in any amount (positive or negative) on any combination of subsets. Unless the odds are computed from a prior probability, dutch book can be made: for some system of bets, the clever gambler wins a dollar or more, no matter what the outcome may be. The extension by Freedman and Purves (1969) to statistical inference is also considered. Finally, there is a dutch-book argument for countable additivity.

1. De Finetti's argument

Let Ω be a finite set with card $(\Omega) > 1$. A bookie posts odds λ_A on A, for every proper $A \subset \Omega$; odds are positive and finite. If you bet b_A on A, and A occurs, then you win b_A/λ_A; if A does not occur, then you win $-b_A$. Your net payoff is

$$(1) \qquad \phi_A = 1_A \frac{b_A}{\lambda_A} - \Big(1 - 1_A\Big) b_A = b_A \frac{1 + \lambda_A}{\lambda_A}\Big(1_A - \frac{\lambda_A}{1 + \lambda_A}\Big).$$

The stakes b_A are finite, but may be positive, 0, or negative. Corresponding to each set of stakes there is a payoff function,

$$(2) \qquad \sum_A \phi_A,$$

where A runs over the proper subsets of Ω. For now, we take the odds as fixed, and consider various gamblers who bet against the bookie: each gambler generates a payoff function.

If $\lambda_A = \pi(A)/[1 - \pi(A)]$ for some probability π on Ω, i.e.,

$$\pi(A) = \lambda_A/(1 + \lambda_A),$$

we say the bookie is a Bayesian with prior π. Obviously, all payoff functions then have expectation 0 relative to π. In particular,

Proposition 1. *Dutch book cannot be made against a Bayesian bookie.*

Let V be the set of all real-valued functions on Ω, so V is a linear space of dimension card (Ω).

Theorem 1. *For each bookie, there are only two possibilities:* (i) *the set of payoff functions coincides with V; or* (ii) *the bookie is a Bayesian with prior π, where π is a probability on Ω assigning positive mass to every $\omega \in \Omega$.*

Proof. If case (i) does not obtain, there is a non-trivial function π on Ω such that for all ϕ_A,

$$\sum_{\omega \in \Omega} \pi(\omega)\phi_A(\omega) = 0,$$

which is to say, by (1),

$$\pi(A) = \pi(\Omega)\lambda_A/(1 + \lambda_A),$$

where

$$\pi(B) = \sum_{\omega \in B} \pi(\omega)$$

for all $B \subset \Omega$. If $\pi(\Omega) = 0$ then $\pi(\omega) = 0$ for all ω, as one sees by taking $A = \{\omega\}$. This is a contradiction, so $\pi(\Omega) \neq 0$. Renormalize π so that $\pi(\Omega) = 1$, i.e., replace π by $\pi/\pi(\Omega)$. Now $\pi(A) = \lambda_A/(1 + \lambda_A)$, so $\pi > 0$ and case (ii) obtains.

Theorem 2. *For the Bayesian bookie with prior π, the set of payoff functions coincides with the set of functions having expectation* 0 *relative to π.*

Proof. Otherwise, there is a non-trivial ν, not a multiple of π, such that all payoff functions are orthogonal to ν as well as π. Arguing as before, we find that $\nu(\Omega) \neq 0$. We renormalize so $\nu(\Omega) = 1$, and then $\nu(A) = \lambda_A/(1 + \lambda_A) = \pi(A)$, a contradiction.

Corollary 1. *The following are equivalent.*

(i) *The payoff functions are all of V.*
(ii) *The bookie is not a Bayesian.*
(iii) *Dutch book can be made against the bookie.*

Proof. (i) $\Rightarrow$ (iii): obvious. (iii) $\Rightarrow$ (ii): Proposition 1. (ii) $\Rightarrow$ (i): if the payoff functions were not all of V, the bookie would be a Bayesian (Theorem 1).

Corollary 2. *The following are equivalent.*

(i) *The payoff functions are not all of V.*
(ii) *The bookie is a Bayesian.*
(iii) *Dutch book cannot be made against the bookie.*

In this case, the payoff functions consist of all functions with expectation 0*, relative to the bookie's prior.*

Proof. (i) $\Rightarrow$ (ii): Theorem 1. (ii) $\Rightarrow$ (iii): Proposition 1. (iii) $\Rightarrow$ (i): obvious. The final assertion follows from Theorem 2.

Example 1. Negative stakes are needed for the theorems. Suppose $\Omega = \{0, 1\}$, the bookie puts 2:1 on 0 and 3:1 on 1, but accepts only positive bets on events. If you put $x \geq 0$ on 0 and $y \geq 0$ on 1, the payoff is $x/2 - y$ on 0 and $y/3 - x$ on 1. If both are positive, $x/2 > y > 3x$, so $x > 6x$, a contradiction. This bookie – although not a Bayesian – is immune from dutch book, being clever enough to set favorable odds and require positive stakes.

2. The extension to statistical inference

In this section, we discuss the principal result in Freedman and Purves (1969). In essence, a bookie has to post odds on subsets of a parameter space Θ after seeing an observation x drawn from $p(\bullet|\theta)$: dutch book can be made against the non-Bayesian bookie. In view of the previous section, the bookie will use an 'estimating probability' $q(\bullet|x)$ on Θ to set the odds. The only question is, how do these probabilities fit together?

A 'finite estimation problem' consists of (i) a finite set $\mathcal{X}$, and (ii) a finite set of parametric models $\{p(\bullet|\theta) : \theta \in \Theta\}$ specifying probability distributions on $\mathcal{X}$. Allowing $p(x|\theta) = 0$ creates one technical nuisance after another, so we assume

$$p(x|\theta) > 0 \text{ for all } x \in \mathcal{X} \text{ and } \theta \in \Theta. \tag{3}$$

An 'estimating probability' $q(\bullet|x)$ is a probability on Θ for each $x \in \mathcal{X}$. Consider subsets $C_1, \ldots, C_k$ of Θ. After x is observed, allow the gambler to pay $b_i(x)q(C_i|x)$ in order to get $b_i(x)$ dollars if $\theta \in C_i$. The gambler is allowed to use any bounded b_i. Bets are settled separately, and then summed.

The net payoff to a gambler who uses the sets $\{C_1, \ldots, C_k\}$ and the functions $\{b_i : i = 1, \ldots, k\}$ is

$$\phi(x, \theta) = \sum_{i=1}^{k} b_i(x) \left[I_{C_i}(\theta) - q(C_i|x) \right]. \tag{4}$$

Any such ϕ is called a 'payoff function'. Nothing requires q to be positive: if $q(C_i|x) = 0$, the gambler pays 0 to get $b_i(x)$ if $\theta \in C_i$. Similarly, C_i may be $\emptyset$ or Θ. In these respects, payoff functions work more smoothly than odds.

Corresponding to each payoff function ϕ on $\mathcal{X} \times \Theta$, there is an 'expected payoff function' on Θ:

$$E_\theta\{\phi\} = \sum_{x \in \mathcal{X}} \phi(x, \theta) p(x|\theta). \tag{5}$$

Definition 1. Dutch book *can be made against the estimating probability* $q(\bullet|x)$ *if there is a gambling system that provides a uniformly positive expected payoff to the gambler: in other words, there exists a payoff function* ϕ *– as defined by* (4) *– and an* $\epsilon > 0$ *such that*

$$E_\theta\{\phi\} > \epsilon \text{ for all } \theta \in \Theta. \tag{6}$$

To paraphrase Freedman and Purves (1969),

> Imagine a Master of Ceremonies who chooses some $\theta \in \Theta$ and then picks $x \in \mathcal{X}$ at random from $p(\bullet|\theta)$. The value of x is revealed and the bookie announces the estimating probability $q(\bullet|x)$. The gambler then constructs a system with payoff function ϕ. When (6) holds, the gambler expects to win at least ϵ no matter what the value of θ.

Of course, if (6) holds for some positive ϵ, any other positive ϵ can be obtained by rescaling the payoff function.

If π is a probability on Θ, we will say that the bookie is a Bayesian with prior π provided

$$m_\pi(x) q(\theta|x) = \pi(\theta) p(x|\theta) \tag{7}$$

where

$$m_\pi(x) = \sum_{\theta \in \Theta} \pi(\theta) p(x|\theta). \tag{8}$$

Here, $m_\pi(x)$ is the 'marginal' probability of x, integrated over the θ's.

Lemma 1. *For the Bayesian bookie with prior* π*, any expected payoff function – as a function of* θ *– integrates to* 0 *against* π.

Proof. In (4), take $k = 1$; write b for b_1 and C for C_1. Then $\sum_\theta E_\theta\{\psi\}\pi(\theta)$ is

$$\sum_{\theta\in\Theta}\sum_{x\in\mathcal{X}} b(x)\big[1_C(\theta) - q(C|x)\big]p(x|\theta)\pi(\theta) = \alpha - \beta$$

where

$$\alpha = \sum_{\theta\in\Theta}\sum_{x\in\mathcal{X}} b(x)1_C(\theta)p(x|\theta)\pi(\theta)$$

and

$$\beta = \sum_{\theta\in\Theta}\sum_{x\in\mathcal{X}} b(x)q(C|x)p(x|\theta)\pi(\theta) = \alpha$$

by 'the law of total probability'. In more detail,

$$\begin{aligned}
\beta &= \sum_{\theta\in\Theta}\sum_{x\in\mathcal{X}} b(x)q(C|x)p(x|\theta)\pi(\theta)\\
&= \sum_{x\in\mathcal{X}}\sum_{\theta\in\Theta} b(x)q(C|x)p(x|\theta)\pi(\theta)\\
&= \sum_{x\in\mathcal{X}} b(x)q(C|x)m_\pi(x)\\
&= \sum_{x\in\mathcal{X}}\sum_{\theta\in C} b(x)\pi(\theta)p(x|\theta) \quad \text{by (7)}\\
&= \sum_{\theta\in\Theta}\sum_{x\in\mathcal{X}} b(x)1_C(\theta)p(x|\theta)\pi(\theta) = \alpha.
\end{aligned}$$

Remark. Another way to say Lemma 1: the payoff functions ϕ in (4) all have expected value 0, relative to the probability on $\mathcal{X}\times\Theta$ that assigns mass $p(x|\theta)\pi(\theta)$ to the pair (x, θ).

Corollary 3. *Dutch book cannot be made against a Bayesian bookie.*

Let V be the set of all real-valued functions on Θ, so V is a linear space of dimension card (Θ).

Theorem 3. *For each bookie, there are only two possibilities:* (i) *the set of expected payoff functions coincides with V; or* (ii) *the bookie is a Bayesian with prior π, where π is a probability on Θ.*

Proof. If case (i) does not hold, there is a non-trivial function π on Θ, orthogonal to all expected payoff functions. For the moment, fix $x \in \mathcal{X}$ and $\theta \in \Theta$.

Specialize ϕ in (4), choosing $k = 1$, $b_1 = 1$ at x and $b_1 = 0$ elsewhere, $C_1 = \{\theta\}$. Then $E_\bullet\{\phi\} \perp \pi$ unpacks to

$$m_\pi(x)q(\theta|x) = \pi(\theta)p(x|\theta). \tag{9}$$

Equation (9) holds for all x and θ. In view of (3), it is not possible for π to be strictly positive at θ_1 and strictly negative at θ_2. Nor can π vanish identically. Hence, we can renormalize π to be a probability on Θ: then (9) says that the bookie is a Bayesian with prior π.

Theorem 4. *For the Bayesian bookie with prior π, the set of payoff functions coincides with the set of functions having expectation* 0 *relative to π.*

Proof. Otherwise, we find a non-trivial function ν on Θ, not a multiple of π, orthogonal to all expected payoff functions. Arguing as before, we show that (9) holds with ν in place of π. Hence, ν can be renormalized to a probability. Let $\mu = (\pi + \nu)/2$, and $\Theta_+ = \{\theta : \theta \in \Theta \;\&\; \mu(\theta) > 0\}$. If $\theta \in \Theta_+$, by (9),

$$\nu(\theta)/\mu(\theta) = m_\nu(x)/m_\mu(x). \tag{10}$$

The right hand side of (10) does not depend on θ. Therefore, $\nu = \mu$. Similarly, $\pi = \mu = \nu$, a contradiction proving Theorem 4.

The corollaries follow, as in the previous section.

Corollary 4. *The following are equivalent.*

(i) *The expected payoff functions are all of V.*
(ii) *The bookie is not a Bayesian.*
(iii) *Dutch book can be made against the bookie.*

Corollary 5. *The following are equivalent.*

(i) *The expected payoff functions are not all of V.*
(ii) *The bookie is a Bayesian.*
(iii) *Dutch book cannot be made against the bookie.*

In this case, the expected payoff functions consist of all functions with expectation 0*, relative to the bookie's prior.*

3. From finite to countable additivity

We modify de Finetti's argument (Section 1) to favor countable additivity rather than finite additivity. Let $(\Omega, \mathcal{A})$ be a probability space, finite or infinite; let ψ be a function from $\mathcal{A}$ to $[0, 1]$. In general, we require $\mathcal{A}$ to be a σ-field. A 'payoff function' is

$$\sum_A c_A[1_A - \psi(A)] \tag{11}$$

where $A \in \mathcal{A}$ and c_A is a finite real number: the sum is over a finite number of A's. Let $\mathcal{F}$ be the set of payoff functions. The normalization is a little different from (1): when $0 < \lambda_A < \infty$, $\psi(A) = \lambda_A/(1 + \lambda_A)$ and $c_A = b_A(1 + \lambda_A)/\lambda_A$. The imagery is different too. Informally, the gambler pays ψ_A for a lottery ticket 1_A, which yields \$1 if A occurs and 0 otherwise: the bet is just enough to get \$1 if A happens. This bet can be scaled by c_A, an arbitrary real number. Any finite number of bets can be placed. When λ_A is large, i.e., $\psi(A) \doteq 1$, then $c_A \doteq b_A$. When λ_A is small, i.e., $\psi(A) \doteq 0$, then $c_A \doteq b_A/\psi(A)$. The advantage of (11) is that the formula can be used even if $\psi(A) = 0$ or 1, corresponding to $\lambda_A = 0$ or ∞.

To avoid dutch book, ψ must be a finitely additive probability on $\mathcal{A}$. When Ω is finite, this can be proved just as before: the new proposition is a little more general, since the odds are unrestricted. As a matter of notation, the previous argument generates a (finitely additive) probability π, and identifies ψ with π. The result is easily extended to the case where Ω is infinite but $\mathcal{A}$ finite – and from there, to the general case.

Suppose next we allow the sum in (11) to be countably infinite, so long as there is uniformly bounded pointwise convergence. More exactly, let $\{A_i : i = 1, 2, \ldots\}$ be a sequence of subsets of $\mathcal{A}$ and $\{c_i\}$ a sequence of real numbers. Let

$$\phi_n = \sum_{i=1}^{n} c_i[1_{A_i} - \psi(A_i)].$$

Suppose

(i) $|\phi_n(\omega)| < L$ for all $n = 1, 2, \ldots$ and all $\omega \in \Omega$, where L is a real number;

(ii) $\lim_n \phi_n(\omega) = \phi(\omega)$ for each $\omega \in \Omega$.

Then ϕ is also considered to be a payoff function. Said another way, the set of payoff functions is enlarged by taking uniformly bounded pointwise limits. Let $\overline{\mathcal{F}}$ denote this larger set of payoff functions.

Now, unless ψ is countably additive, a dutch book can be made. Indeed, let $A_i \in \mathcal{A}$ be pairwise disjoint and $A = \bigcup_i A_i$. Of course, $\psi(A) \geq \sum_{i=1}^{\infty} \psi(A_i)$. Suppose that the inequality is strict. Consider a payoff function that has a bet on each A_i, and a bet against A, to win \$1 in each case:

$$[\psi(A) - 1_A] + \sum_{i=1}^{\infty} [1_{A_i} - \psi(A_i)] = \lim_n \left\{ [\psi(A) - 1_A] + \sum_{i=1}^{n} [1_{A_i} - \psi(A_i)] \right\}$$
$$= \psi(A) - \sum_{i=1}^{\infty} \psi(A_i) > 0.$$

In short, finite additivity exposes the odds-maker to a dutch book – if the gambler can use payoff functions in $\overline{\mathcal{F}}$. Of course, if the odds are computed from a countably additive π, each payoff function has expectation 0 relative to π, by the dominated convergence theorem, and dutch book is impossible.

Theorem 5. *Suppose the set $\mathcal{F}$ of payoff functions is initially defined by* (11), *with $A \in \mathcal{A}$, a σ-field of subsets of Ω; only finitely many A's are allowed. The set of payoff functions is then enlarged to $\overline{\mathcal{F}}$ by allowing uniformly bounded pointwise limits. A gambler can make dutch book unless the odds are compatible with a countably additive prior on* $(\Omega, \mathcal{A})$.

Example 2. Suppose ψ is a remote (finitely additive) probability on the integers. The gambler pays 0 to win \$1 if integer j materializes, for all integers. The odds-maker will in the end have to fork over \$1.

A primer on odds

If the odds on A are 3:1, and we bet \$1 on A:
- if A occurs we win \$1/3;
- if A does not occur, we lose our \$1.

If the odds on A are 3:1, and we bet \$1 against A:
- if A occurs we lose our \$1;
- if A does not occur, we win \$3.

If the odds against B are 3:1, and we bet \$1 on B:
- if B occurs we win \$3;
- if B does not occur, we lose our \$1.

If the odds against B are 3:1, and we bet \$1 against B:

if B occurs we lose our \$1;

if B does not occur, we win \$1/3.

If $P(A) = 3/4$—

the odds on A are 3:1

the odds against A are 1:3

the odds on A^c are 1:3

the odds against A^c are 3:1

References

de Finetti, B. (1931). Sul significato soggestivo della probabilità. *Fundamenta Mathematicae* 17: 298–329.

de Finetti, B. (1937). La prévision: ses lois logiques ses sources subjectives. *Ann. Inst. H. Poincaré* 7: 1–68. English translation in H. E. Kyburg, H. E. Smokler (eds), *Studies in Subjective Probability*. New York: Wiley, 1964.

Eaton, M. L. and Freedman, D. A. (2003). Dutch book against 'objective' priors. Technical Report #642, Department of Statistics, U. C. Berkeley. http://www.stat.berkeley.edu/users/census/642.pdf

Freedman, D. A. (1995). Some issues in the foundations of statistics. *Foundations of Science* 1: 19–83, with discussion. Reprinted in B. van Fraassen (ed), *Topics in the Foundation of Statistics*. Dordrecht: Kluwer, 1996.

Freedman, D. A. and Purves, R. A. (1969). Bayes method for bookies. *Ann. Math Statist.* 40: 1177–1186.

Ramsey, F. (1926). Truth and Probability. Reprinted in H. E. Kyburg and H. E. Smokler (eds), *Studies in Subjective Probability*, pp. 25–52. Huntington, New York: Krieger Publishing, 1980.

15 July 2003