Bernoulli **10**(5), 2004, 861–872

Dutch book against some 'objective' priors

MORRIS L. EATON[1] and DAVID A. FREEDMAN[2]

[1]*School of Statistics, University of Minnesota, 313 Ford Hall, 224 Church St. SE, Minneapolis MN 55455, USA. E-mail: eaton@stat.umn.edu*
[2]*Department of Statistics, University of California, Berkeley CA 94705-3860, USA. E-mail: census@stat.berkeley.edu*

'Dutch book' and 'strong inconsistency' are generally equivalent: there is a system of bets that makes money for the gambler, whatever the state of nature may be. As de Finetti showed, an odds-maker who is not a Bayesian is subject to a Dutch book, under certain highly stylized rules of play – a fact often used as an argument against frequentists. However, so-called 'objective' or 'uninformative' priors may also be subject to a Dutch book. This note explains, in a relatively simple and self-contained way, how to make Dutch book against a frequently recommended uninformative prior for covariance matrices.

Keywords: Dutch book; incoherence; Jeffreys prior; multivariate analysis; objective prior; prediction; strong inconsistency; uninformative prior

1. Introduction

Our aim here is to sketch a relatively simple and self-contained argument to show that Dutch book can be made against certain 'objective' priors, namely, invariant measures that have infinite mass and are used as 'priors' in formal Bayesian calculations. Such improper priors are often said to be 'uninformative'. Our main example involves a prior that was recommended by Jeffreys – and adopted by many Bayesians – for use in multivariate normal distributions where the covariance matrix is unknown.

To fix ideas, we begin with de Finetti's Dutch-book argument against the frequentists. Let Ω be a finite set. A bookie has to post finite, positive odds on every subset of Ω, apart from the empty set and Ω itself, accepting bets on each set at those odds. Bets can be laid in any amount (positive or negative) on any combination of sets. The Bayesian bookie will have a prior probability π on Ω, and will post odds $\pi(A)/[1-\pi(A)]$ on A. The non-Bayesian bookie posts odds compatible with no π. In this context, 'Dutch book' means a system of bets that – no matter what ω is chosen from Ω – yields a positive payoff to the bettor.

Plainly, Dutch book cannot be made against a Bayesian bookie, since the expected payoff is 0. On the other hand, as de Finetti showed, Dutch book can be made against any non-Bayesian bookie. The relevance to applied statistical work is not entirely clear, since few statisticians place bets when doing data analysis, and few bookies follow de Finetti's rules for accepting bets. However, the argument has often been deployed against the frequentists: for discussion, see Freedman (1995). The possibility of a Dutch book is sometimes referred

1350–7265

to as a ‘money pump’: if you can win a dollar, you make the poor bookie play the game over and over again, pumping money from him to you.

De Finetti’s example can be viewed as a prediction problem: nature will choose ω from Ω, and the odds are a stylized way of describing opinions about the future. Freedman and Purves (1969) modified the argument to cover a two-stage process with conditional bets. The pair (x, z) is chosen at random from P_θ, where θ is an unknown parameter; x is observed first, z second. (The parameter space and the observation space are required to be finite, as in de Finetti’s work.) The gambler is allowed to bet on z, and bets are allowed to depend on x. The concept of the Dutch book must be extended slightly: the clever gambler can arrange to have a positive expected payoff from a non-Bayesian bookie, simultaneously for all θ – but may have to take a loss for some combinations of θ, x and z. Some observers may view the passage from unconditional to conditional bets as a small variation on de Finetti’s set-up; others consider this generalization to be a major – and subversive – idea.

Our main example involves prediction, in a setting like that of Freedman and Purves (1969) – although the spaces are infinite. There are n independent observations from a common multivariate normal distribution, having mean 0 and (unknown) positive definite covariance matrix Σ. The observations are denoted $X_1, \ldots, X_n$; they are used to predict an $(n+1)$th observation, denoted Z. Indeed, having observed $X_1, \ldots, X_n$, the statistician is required to produce a ‘predictive distribution’ for Z. If the Jeffreys prior is used to generate this predictive distribution, the statistician is exposed to a Dutch book.

A similar example can be constructed for estimation (Section 2.2), but the argument is a little harder. Stone (1976) has the concept of ‘strong inconsistency’, defined below. In Section 2, we show that the Jeffreys prior leads to strong inconsistency. Section 3 demonstrates that strong inconsistency is equivalent to a Dutch book. Section 4 reviews the literature.

2. Main example

Let $X_1, \ldots, X_n$ be independent random $p \times 1$ vectors, with a common $N_p(0, \Sigma)$ distribution; the covariance matrix Σ is $p \times p$ and positive definite. Here, $n \geqslant p > 1$. An $(n+1)$th observation Z will be drawn independently from $N_p(0, \Sigma)$. How can the data $X = (X_1, \ldots, X_n)$ be used to predict Z? Since the problem is invariant under multiplication by a $p \times p$ non-singular matrix A,

$$X_i \to AX_i, \qquad Z \to AZ, \qquad \Sigma \to A\Sigma A',$$

an invariant ‘prior distribution’ might suggest itself. (Quote marks are used because the prior is improper, with infinite total mass.) The invariant prior is $\mathrm{d}\Sigma/|\Sigma|^{(p+1)/2}$. It is unique up to a positive constant, and can be recognized as the Jeffreys prior $\sqrt{|I|}\mathrm{d}\Sigma$, where I is the Fisher information matrix and $|M|$ is the determinant of M.

A predictive distribution for Z can be computed from the Jeffreys prior, by a formal application of Bayes’s rule. In more detail, let $\phi(x, z|\Sigma)$ be the multivariate normal density of (X, Z) given Σ; similarly, $\phi(x|\Sigma)$ is the multivariate normal density of X given Σ. Here,

$x \in (\mathbb{R}^p)^n$ and $z \in \mathbb{R}^p$. Formally, the 'predictive density' for X, Z – in advance of data collection – is obtained by integrating Σ against the Jeffreys prior:

$$\phi(x, z) = \int \phi(x, z|\Sigma)\mathrm{d}\Sigma/|\Sigma|^{(p+1)/2}. \tag{1}$$

Likewise, the 'predictive density' for X is

$$\phi(x) = \int \phi(x|\Sigma)\mathrm{d}\Sigma/|\Sigma|^{(p+1)/2}. \tag{2}$$

According to 'Bayes's rule', the 'predictive density' for Z when $X = x$ is

$$\phi(z|x) = \phi(x, z)/\phi(x). \tag{3}$$

We have quote marks because the prior is improper, so

$$\int\int \phi(x, z)\mathrm{d}x\mathrm{d}z = \int \phi(x)\mathrm{d}x = \infty.$$

On the other hand, $\phi(z|x)$ is a proper density for Z, because $\int \phi(x, z|\Sigma)\mathrm{d}z = \phi(x|\Sigma)$, so $\int \phi(x, z)\mathrm{d}z = \phi(x)$, and $\int \phi(z|x)\mathrm{d}z = 1$.

The main result of this section can now be stated: the predictive density $\phi(z|x)$ is strongly inconsistent, and Dutch book can be made against a statistician who uses it. (Strong inconsistency is defined below.)

Theorem 1. *Let* $n \geqslant p > 1$. *Suppose* $X_1, \ldots, X_n, Z$ *are independent* $N_p(0, \Sigma)$. *The predictive distribution for* Z *given* $X_1, \ldots, X_n$, *computed from the Jeffreys prior on the covariance matrix* Σ, *is strongly inconsistent.*

An outline of the proof follows, with details in Sections 2.1 and 3. Let $S = \sum_{i=1}^n X_i X_i'$. We take X_i and Z to be $p \times 1$ column vectors, so S is a $p \times p$ matrix. By eliminating a null set, we can take S to be positive definite. Write Z_1 for the first coordinate of Z, and S_{11} for the (1, 1)th element of S. Then, as is almost obvious, the sampling distribution D_0 of $T = Z_1/\sqrt{S_{11}}$ does not depend on Σ:

$$\sqrt{n}T \sim t_n. \tag{4}$$

Here '$\sim$' means 'is distributed as', and t_n is t with n degrees of freedom.

When $X = x$, the predictive distribution $Q(\mathrm{d}z|x)$ for the $(n+1)$th observation Z, computed formally by Bayes's rule from the Jeffreys prior, has a density on $\mathbb{R}^p$ given by

$$\phi(z|x) = C/\sqrt{|s|(1 + z's^{-1}z)^{n+1}}, \tag{5}$$

where C is a constant, and $s = \sum_{i=1}^n x_i x_i'$ is the value of S at the observed $X = x$. As before, we take x_i and z to be $p \times 1$, so x is $p \times n$ and s is $p \times p$. The proof of (5) is 'just' calculus (Section 2.1). The constant C depends on n and p, not on x or z.

Next, the predictive distribution D_1 of $T = Z_1/\sqrt{S_{11}}$ when $X = x$ does not depend on x; indeed,

$$\sqrt{n-p+1}T \sim t_{n-p+1}, \tag{6}$$

as will also be proved in Section 2.1. By (4) and (6),

$$D_1 \neq D_0. \tag{7}$$

Inequality (7) is the key point, and strong inconsistency will soon follow. (If $p = 1$ then $D_1 = D_0$; that is why we assumed $p > 1$.)

Let E_Σ denote expectation relative to our sampling model for $X_1, \ldots, X_n$, Z, and recall that $Q(\mathrm{d}z|x)$ is the predictive distribution for Z given $X = x$. 'Strong inconsistency' means there is a bounded measurable function f and an $\epsilon > 0$ with

$$\int f(x, z)Q(\mathrm{d}z|x) + \epsilon \leqslant \mathrm{E}_\Sigma\{f(X, Z)\} \tag{8}$$

for all x and Σ. Section 3 discusses the definition in a more general framework, but here is the point. The left-hand side of (8) depends on x not Σ; the right-hand side on Σ, not x. Thus,

$$\alpha = \sup_x \int f(x, z)Q(\mathrm{d}z|x) < \inf_\Sigma \mathrm{E}_\Sigma\{f(X, Z)\} = \beta.$$

Consider a 'lottery' that pays $f(x, z)$ when $X = x$ and $Z = z$. The expected value of the lottery, in advance of data collection, is at least β. Thus, a statistician (of the kind envisaged by de Finetti) should pay β to buy the lottery. On the other hand, after seeing X, a statistician who uses the predictive distribution Q should happily sell the lottery for α, no matter what X proves to be. The gap between α and β reflects an inconsistency in the pricing, and leaves room for Dutch book against Q.

Strong inconsistency is easily demonstrated, by constructing f and ϵ. In view of (7), there is a bounded measurable function h and $\epsilon > 0$ with

$$\int_{-\infty}^{\infty} h(v)D_1(\mathrm{d}v) + \epsilon \leqslant \int_{-\infty}^{\infty} h(v)D_0(\mathrm{d}v). \tag{9}$$

Let $f(x, z) = h(z_1/\sqrt{s_{11}})$, where z_1 is the first coordinate of the $p \times 1$ column vector z, and s_{11} is the (1, 1)th element of the $p \times p$ matrix $s = \sum_{i=1}^n x_i x_i'$. Now (9) boils down to (8), by the change-of-variables formula for integrals. On the left-hand side, D_1 is the $Q(\mathrm{d}z|x)$-distribution of $z_1/\sqrt{s_{11}}$. On the right-hand side, D_0 is the sampling distribution of $T = Z_1/\sqrt{S_{11}}$. With our f, both sides of (8) are constant, by (4) and (6). This completes a sketch of the argument for strong inconsistency. Some details are given next, and Dutch book is discussed in Section 3.

2.1. Some details

We begin with a two-step proof of (5).

Step 1. We first show that the predictive density $\phi(x)$ of $X = (X_1, \ldots, X_n)$, computed from the model and the Jeffreys prior in advance of data collection, is $C_{n,p}/|s|^{n/2}$, where $C_{n,p}$ is a constant, $s = \sum_{i=1}^n x_i x_i'$ is positive definite (almost everywhere), and $|s|$ is the determinant of s. As in (2),

$$\phi(x) = \int \phi(x|\Sigma)\mu(\mathrm{d}\Sigma),$$

where $\phi(x|\Sigma)$ is the multivariate normal density,

$$\phi(x|\Sigma) = (2\pi)^{-np/2}|\Sigma|^{-n/2} \exp\left(-\frac{1}{2}\sum_{i=1}^{n} x_i'\Sigma^{-1}x_i\right),$$

and $\mu(\mathrm{d}\Sigma) = \mathrm{d}\Sigma/|\Sigma|^{(p+1)/2}$ is the Jeffreys prior. As before, $x = (x_1, \ldots, x_n)$, with x_i a $p \times 1$ column vector, and Σ is a $p \times p$ positive definite matrix. Since trace(LM) = trace(ML) when both products are defined,

$$\sum_{i=1}^{n} x_i'\Sigma^{-1}x_i = \mathrm{trace}(\Sigma^{-1}s) = \mathrm{trace}(s^{1/2}\Sigma^{-1}s^{1/2}). \tag{10}$$

Since $|LM| = |ML|$ for $p \times p$ matrices,

$$|\Sigma|^{-n/2} = |s|^{-n/2}|s^{-1/2}\Sigma s^{-1/2}|^{-n/2}. \tag{11}$$

So

$$\begin{aligned}\phi(x) &= \int \phi(x|\Sigma)\mu(\mathrm{d}\Sigma) \qquad (12)\\ &= |s|^{-n/2}\int \psi(s|\Sigma)\mu(\mathrm{d}\Sigma)\\ &= |s|^{-n/2}\int \psi(I_{p\times p}|\Sigma)\mu(\mathrm{d}\Sigma),\end{aligned}$$

where $I_{p\times p}$ is the $p \times p$ identity matrix,

$$\psi(s|\Sigma) = D_{n,p}(s)\exp[-\tfrac{1}{2}\mathrm{trace}(s^{1/2}\Sigma^{-1}s^{1/2})],$$
$$D_{n,p}(s) = (2\pi)^{-np/2}|s^{-1/2}\Sigma s^{-1/2}|^{-n/2}.$$

The first line in (12) is just (2). The second line holds by (10) and (11), with a bit of algebraic juggling for the constants. The third line holds because the Jeffreys prior μ is invariant. Equation (12) is the required formula for the predictive density $\phi(x)$ of $X = (X_1, \ldots, X_n)$: the last integral in (12) is $C_{n,p}$, the mystery constant in that density. (Computing the integral is a task not lightly to be undertaken.) This completes step 1 in proving (5).

Step 2. The predictive density for Z when $X = x$ is obtained from (3), as the quotient of (1) and (2). In view of (12),

$$\phi(z|x) = (C_{n+1,p}/C_{n,p})|s|^{n/2}\,|s + zz'|^{-(n+1)/2}. \tag{13}$$

This simplifies to (5), because of the identity

$$|I_{p\times p} + ww'| = 1 + w'w, \tag{14}$$

with $w = s^{-1/2}z$, a $p \times 1$ column vector. To verify (14), let L rotate w into $(a, 0, \ldots, 0)'$, so $LL' = I_{p\times p}$ and $a^2 = w'w$. The left-hand side of (14) equals

$$|LL' + Lww'L'| = |I_{p\times p} + Lww'L'| = 1 + a^2.$$

This proves (14), completing the argument for (5).

We now turn to (6). Let z_j be the jth coordinate of the $p \times 1$ column vector z; recall that s_{11} is the (1, 1)th entry of the matrix s. What is the $Q(\mathrm{d}z|x)$ distribution of $z_1/\sqrt{s_{11}}$? When $s = I_{p\times p}$, the density is proportional to $1/(1 + z_1^2)^{(n-p+2)/2}$, as one shows by integrating out $z_2, \ldots, z_p$. To do the integral, set

$$z_j = \sqrt{1 + z_1^2}\, w_j \tag{15}$$

for $j = 2, \ldots, p$. This proves (6) when $s = I_{p\times p}$.

Next, $Q(\mathrm{d}z|x)$ is the distribution q_s on $\mathbb{R}^p$ whose density is

$$\phi(z|x) = C/\sqrt{|s|(1 + z's^{-1}z)^{n+1}}. \tag{16}$$

Dependence on x is only through $s = \sum_{i=1}^n x_i x_i'$: see (5). Abbreviate $I = I_{p\times p}$. Then q_s is the q_I-law of $s^{1/2}z$. What remains to be seen is that

$$\text{the } q_I\text{-law of } (s^{1/2}z)_1/\sqrt{s_{11}} \text{ does not depend on } s, \tag{17}$$

where $(s^{1/2}z)_1$ is the first coordinate of the $p \times 1$ column vector $s^{1/2}z$. Claim (17) is immediate from the invariance of q_I under rotation: $\sqrt{s_{11}}$ is the ℓ_2 norm of the first row of $s^{1/2}$. (Bear in mind that $s^{1/2}$ is symmetric.) Since q_s is the predictive distribution for Z when $X = x$, the argument for (6) is complete, and with it, the proof of Theorem 1.

2.2. Estimation

We have established strong inconsistency when the Jeffreys prior is used in a multivariate normal prediction problem. A parallel – but more technical – argument yields the same conclusion in an estimation context. Here is a brief sketch; readers can skip to Section 3 without loss of continuity. Let $X_1, \ldots, X_n$ be independent random $p \times 1$ vectors, with a common $N_p(0, \Sigma)$ distribution. Consider the Jeffreys prior distribution $\mu(\mathrm{d}\Sigma) = \mathrm{d}\Sigma/|\Sigma|^{(p+1)/2}$. Then the formal posterior distribution for $\theta = \Sigma^{-1}$ given $X_1, \ldots, X_n$ is Wishart:

$$\theta \sim W(S^{-1}, p, n), \tag{18}$$

with $S = \sum_{i=1}^n X_i X_i'$. This is non-trivial, but can be verified starting from the fact that the Jeffreys prior μ is invariant under the transformation $\Sigma \to \Sigma^{-1}$, by Jacobian trickery. Given S, the posterior density of $\theta = \Sigma^{-1}$ with respect to μ is proportional to

$$|\theta S|^{n/2} \exp[-\tfrac{1}{2}\,\mathrm{trace}(\theta S)],$$

which is the density of $W(S^{-1}, p, n)$, as in Eaton (1983, p. 240). One analogue of the statistic T is

$$U = \frac{\theta_{11}}{S^{11}},$$

where S^{11} is the (1, 1)th element of S^{-1}. When – given the data – θ has the distribution (18), then

$$U \sim \chi_n^2 \tag{19}$$

for each S, as is immediate from the definition of the Wishart: if ζ_i are independent $N_p(0, K)$, then $\sum_{i=1}^n \zeta_i \zeta_i' \sim W(K, p, n)$. However, under the sampling model, S is $W(\Sigma, p, n)$, and

$$U \sim \chi_{n-p+1}^2 \tag{20}$$

for any Σ, by Proposition (8.7) in Eaton (1983). The fact that (19) and (20) are different for $p > 1$ leads to strong inconsistency, as before.

3. Dutch book in the prediction problem

Our purpose here is to show that strong inconsistency is equivalent to a Dutch book. A 'measurable prediction problem' consists of: (i) $X \in \mathcal{X}$; (ii) $Z \in \mathcal{Z}$, which is to be predicted from X; (iii) a set of parametric models $\{P(\mathrm{d}x, \mathrm{d}z|\theta) : \theta \in \Theta\}$ specifying the joint distribution of X and Z. Here, $\mathcal{X}$ and $\mathcal{Z}$ are measurable spaces, while X and Z are measurable functions on some underlying probability space. A 'predictive distribution' $Q(\mathrm{d}z|x)$ is a distribution for Z that depends on the observed value $X = x$. If $x \to Q(A|x)$ is measurable on $\mathcal{X}$ for every measurable $A \subset \mathcal{Z}$, we will say that Q is measurable.

One way to evaluate Q involves gambling scenarios, as follows. Consider a measurable subset $C \subset \mathcal{X} \times \mathcal{Z}$ and let $C_x = \{z : (x, z) \in C\}$ be the x-section of C. Then a 'simple payoff function' is

$$\psi_C(x, z) = I_C(x, z) - Q(C_x|x). \tag{21}$$

If Q is measurable, then $(x, z) \to \psi_C(x, z)$ is measurable by the usual argument, starting from measurable rectangles.

By way of interpretation, $\psi_C(x, z)$ is the net payoff to a gambler who puts down $Q(C_x|x)$ dollars to get a dollar if $Z \in C_x$. This net payoff is $1 - Q(C_x|x)$ if $Z \in C_x$ and $-Q(C_x|x)$ if $Z \notin C_x$. No money changes hands if $C_x = \emptyset$ or $\mathcal{Z}$. (In particular, there are no interesting bets on x.)

Now, consider measurable subsets $C_1, \ldots, C_k$ of $\mathcal{X} \times \mathcal{Z}$. After $X = x$ is observed, allow the gambler to pay $b_i(x)Q(C_{i,x}|x)$ in order to get $b_i(x)$ dollars if $Z \in C_{i,x}$. The gambler is allowed to use any bounded measurable b_i; this is viewed as encouraging honesty on the part of the odds-maker. Bets are settled separately, and then summed. The net payoff to a gambler who uses the sets $\{C_1, \ldots, C_k\}$ and the betting functions $\{b_i : i = 1, \ldots, k\}$ is

$$\psi(x, z) = \sum_{i=1}^k b_i(x)[I_{C_i}(x, z) - Q(C_{i,x}|x)]. \tag{22}$$

Any such ψ is called a ‘payoff function’. Clearly,

$$\int \psi(x, z) Q(\mathrm{d}z|x) = 0 \tag{23}$$

for all x. Thus, if your predictive distribution for Z – after observing $X = x$ – is $Q(\mathrm{d}z|x)$, all these payoff functions seem fair.

Definition 1. Dutch book *can be made against the predictive distribution* $Q(\mathrm{d}z|x)$ *if there is a gambling system that provides a uniformly positive expected payoff to the gambler: in other words, there exists a payoff function* ψ *– as defined by* (22) *– and an* $\epsilon > 0$ *such that*

$$\epsilon \leqslant \int\int \psi(x, z) P(\mathrm{d}x, \mathrm{d}z|\theta) \qquad \textit{for all } \theta \in \Theta. \tag{24}$$

To paraphrase Freedman and Purves (1969), imagine a master of ceremonies who picks some $\theta \in \Theta$ and then draws (X, Z) from the model $P(\mathrm{d}x, \mathrm{d}z|\theta)$. The value of $X = x$ is revealed and the statistician announces the predictive distribution $Q(\mathrm{d}z|x)$. The gambler then lays bets with payoff function ψ. When (24) holds, the gambler expects to win at least ϵ no matter what the value of θ.

Of course, if (24) holds for some positive ϵ, any other positive ϵ can be obtained by rescaling the payoff function. We will say that Dutch book can be made against Q; more explicitly, Dutch book can made against a bookie who – after seeing that $X = x$ – sets odds on Z using $Q(\mathrm{d}z|x)$. Other language abounds: for instance, there is a Dutch book; or, if Q is computed from an improper prior π, Dutch book can be made against π.

In this paper, we allow only conditional bets, and obtain results on *expected* loss for a non-Bayesian bookie. De Finetti allowed unconditional bets (Section 1) and obtained results on *actual* loss. For more discussion, see Freedman and Purves (1969) or Sudderth (1994).

Definition 2. *The predictive distribution* $Q(\mathrm{d}z|x)$ *is* strongly inconsistent *if there exists a bounded measurable function* $f(x, z)$ *and an* $\epsilon > 0$ *such that*

$$\int f(x, z) Q(\mathrm{d}z|x) + \epsilon \leqslant \int\int f(x, z) P(\mathrm{d}x, \mathrm{d}z)|\theta) \tag{25}$$

for all $x \in \mathcal{X}$ *and* $\theta \in \Theta$.

Equation (8) was a special case, with $\theta = \Sigma$. We now show that $Q(\mathrm{d}z|x)$ is strongly inconsistent if and only if Dutch book can made against Q.

Theorem 2. *Let* $X \in \mathcal{X}$, $Z \in \mathcal{Z}$, *and* $\{P(\mathrm{d}x, \mathrm{d}z|\theta) : \theta \in \Theta\}$ *be a measurable prediction problem. Let* $Q(\mathrm{d}z|x)$ *be a measurable predictive distribution for* Z *when* $X = x$. *Then* $Q(\mathrm{d}z|x)$ *is strongly inconsistent if and only if Dutch book can be made against* Q.

Here is the argument. First, if (24) holds (Dutch book), then $f(x, z) \equiv \psi(x, z)$ is bounded and (25) holds because of (23), proving strong inconsistency. For the converse, assume (25) holds for some bounded measurable f and $\epsilon > 0$. The left-hand side of (25) is a function of x only; the right, of θ only. Thus,

$$\sup_x \int f(x, z)Q(\mathrm{d}z|x) + \epsilon \leqslant \inf_\theta \int\int f(x, z)P(\mathrm{d}x, \mathrm{d}z|\theta). \tag{26}$$

Since f is bounded, it can be uniformly approximated by a simple function f_0, and

$$\int f_0(x, \zeta)Q(\mathrm{d}\zeta|x) \leqslant \alpha < \beta \leqslant \int\int f_0(x, z)P(\mathrm{d}x, \mathrm{d}z|\theta) \tag{27}$$

for suitable real numbers α and β. The inequality holds for all $x \in \mathcal{X}$ and $\theta \in \Theta$; using ζ for the variable of integration may be helpful later. Let

$$f_1(x, z) = f_0(x, z) - \int_{\mathcal{Z}} f_0(x, \zeta)Q(d\zeta|x)$$

Plainly, f_1 is a payoff function in the sense of (22). Now

$$\int_{\mathcal{Z}}\int_{\mathcal{X}} f_1(x, z)P(\mathrm{d}x, \mathrm{d}z|\theta) = B - A,$$

where

$$B = \int_{\mathcal{Z}}\int_{\mathcal{X}} f_0(x, z)P(\mathrm{d}x, \mathrm{d}z|\theta) \geqslant \beta$$

and

$$A = \int_{\mathcal{Z}}\int_{\mathcal{X}}\int_{\mathcal{Z}} f_0(x, \zeta)Q(d\zeta|x)P(\mathrm{d}x, \mathrm{d}z|\theta) \leqslant \int_{\mathcal{Z}}\int_{\mathcal{X}} \alpha P(\mathrm{d}x, \mathrm{d}z|\theta) \leqslant \alpha;$$

the inequalities hold by (27), establishing the long-sought Dutch book, namely, inequality (24) with $\epsilon = \beta - \alpha$ and $\psi = f_1$. This completes the proof of Theorem 2.

The equivalence of strong inconsistency and Dutch book also holds for estimation. To see this, just take $\mathcal{Z} = \Theta$, and let Z be the identity map on Θ. Initially, $P(\mathrm{d}x|\theta)$ will be defined only on $\mathcal{X}$; we require some measurable structure on Θ and set $P(A \times B|\theta) = P(A|\theta)1_B(\theta)$ for measurable $A \subset \mathcal{X}$ and $B \subset \Theta$. A predictive distribution is a 'posterior' for θ given $X = x$. Quotes are needed unless Q is computed from a proper prior – but then, there are no paradoxes to discuss.

4. Literature review

Writing in 1926, Ramsey (1931) introduced the idea of betting odds as a means of assessing probability assignments. De Finetti (1931; 1937) used similar ideas in his discussion of what is now commonly known as coherence.

> di un individuo che debba tenere un banco di scommesse su dati eventi, accettando alle stesse conditizioni qualunque scommessa nell'uno o nell'altro senso. Vedremo che egli è costretto allora a rispettare certe restrizioni, che sono i teoremi del calcolo delle probabilità. Altrimenti egli pecca di

> *coerenza*, e perde *sicuramente*, purchè l'avversario sappia sfruttare il suo errore. (De Finetti 1931, p. 305)

In free translation:

> A person who is obliged to accept bets in any amount, positive or negative, on any finite combination of events, must fix prices according to the laws of probability theory. Otherwise, this person sins against coherence and loses money with certainty, provided the opponent knows how to exploit the mistake.

Some time afterward, the term 'Dutch book' entered the lexicon as a synonym for incoherence. The earliest citation we could find was Lehman (1955), although use of 'Dutch' as a pejorative dates back to seventeenth-century England.

Freedman and Purves (1969) gave the Dutch-book idea careful mathematical expression, for prediction and estimation, when all the spaces are finite. Rigorous treatments in a finitely additive setting for the estimation problem can be found in Heath and Sudderth (1978; 1989). Extensions to the prediction context appear in Lane and Sudderth (1984). Heath, Lane, and Sudderth allow infinite spaces.

Another foundational idea, 'strong inconsistency', was introduced by Stone (1976), and later adapted to the predictive setting by Lane and Sudderth (1984); also see Eaton and Sudderth (1993; 1999). The equivalence of strong inconsistency and incoherence is discussed in the finitely additive setting by Lane and Sudderth (1983). Proof in the countably additive setting is a little different (Section 3). At the risk of the obvious, strong inconsistency is an exact finite-sample property, rather than an asymptotic large-sample property.

That improper prior distributions can give rise to posterior distributions with disturbing properties has been known since at least the 1970s. Stone (1976) and the discussants of Stone's paper provide examples, including the Jeffreys prior. Eaton and Sudderth (1993; 1995; 1998; 1999; 2001; 2002) discuss invariant prediction problems, and show that in the multivariate linear model, fully invariant predictive distributions are strongly inconsistent: the 'principle of invariance' (Berger, 1985, p. 390) therefore leads to Dutch book. The Jeffreys prior can be viewed as a prototype where elementary arguments suffice (Section 2). In contrast, much of the Eaton–Sudderth work relies on separation theorems of the Hahn–Banach type, which makes the results less accessible. Eaton and Sudderth (1999, Section 8) show that if the transformation group is amenable, there will be an invariant predictive distribution immune to Dutch book – although other invariant predictive distributions will be vulnerable. If the transformation group is not amenable – like the non-singular linear transformations on $\mathbb{R}^p$ for $p > 1$ in Section 2 – all invariant predictive distributions may be subject to Dutch book (Eaton and Sudderth, 1998).

The uniqueness of the invariant prior in Section 2 above is demonstrated by Eaton (1983, Example 6.19). Use of this prior has been suggested by Jeffreys (1961, pp. 180–181), Box and Tiao (1973, p. 426; 1992, Section 8.2.2), Geisser (1993, Chapter 9), Schervish (1995, p. 122) and Keyes and Levy (1996). The latter also has a good survey of invariant predictive distributions in multivariate analysis of variance. For an interesting generalization of (14), see Eaton (1983, p. 43). Although incoherence is often a synonym for the possibility of a

Dutch book, other definitions have been suggested. See Regazzini (1987) as well as Berti *et al*. (1991). For more discussion, see Sudderth (1994, Section 7).

Acknowledgements

We thank Eugenio Regazzini and Philip Stark for helpful comments.

References

Berger, J.O. (1985) *Statistical Decision Theory and Bayesian Analysis*, 2nd edn. New York: Springer-Verlag.

Berti, P., Regazzini, E. and Rigo, P. (1991) Coherent statistical inference and Bayes theorem. *Ann. Statist.*, **19**, 366–381.

Box, G.E.P. and Tiao, G.C. (1973) *Bayesian Inference in Statistical Analysis*. Reading, MA: Addison-Wesley.

Box, G.E.P. and Tiao, G.C. (1992) *Bayesian Inference in Statistical Analysis*. New York: Wiley.

de Finetti, B. (1931) Sul significato soggestivo della probabilità. *Fund. Math.*, **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 and H.E. Smokler (eds), *Studies in Subjective Probability.* New York: Wiley, 1964.

Eaton, M.L. (1983) *Multivariate Statistics*. New York: Wiley.

Eaton, M.L. and Sudderth, W.D. (1993) Prediction in a multivariate normal setting: Coherence and incoherence. *Sankhyā Ser. A*, **55**, 481–493.

Eaton, M.L. and Sudderth, W.D. (1995) The formal posterior of a standard flat prior in MANOVA is incoherent. *J. Italian Statist. Soc.*, **2**, 251–270.

Eaton, M.L. and Sudderth, W.D. (1998) A new predictive distribution for normal multivariate linear models. *Sankhyā Ser. A*, **60**, 363–382.

Eaton, M.L. and Sudderth, W.D. (1999) Consistency and strong inconsistency of group-invariant predictive inferences. *Bernoulli*, **5**, 833–854.

Eaton, M.L. and Sudderth, W.D. (2001) Best invariant predictive distributions. In M. Viana and D. Richards (eds), *Algebraic Methods in Statistics and Probability*, Contemp. Math. 287, pp. 49–62. Providence, RI: American Mathematical Society.

Eaton, M.L. and Sudderth, W.D. (2002) Group invariant inference and right Haar measure. *J. Statist. Plann. Inference*, **103**, 87–99.

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.

Geisser, S. (1993) *Predictive Inference: An Introduction*. New York: Chapman & Hall.

Heath, D. and Sudderth, W.D. (1978) On finitely additive priors, coherence, and extended admissibility. *Ann. Statist.*, **2**, 333–345.

Heath, D. and Sudderth, W.D. (1989) Coherent inference from improper priors and from finitely additive priors. *Ann. Statist.*, **17**, 907–919.

Jeffreys, H. (1961) *Theory of Probability*, 3rd edn. Oxford: Oxford University Press.

Keyes, T.K. and Levy, M.S. (1996) Goodness of prediction fit for multivariate linear models. *J. Amer. Statist. Assoc.*, **91**, 191–197.

Lane, D. and Sudderth, W.D. (1983) Coherent and continuous inference. *Ann. Statist.*, **11**, 114–120.

Lane, D. and Sudderth, W.D. (1984) Coherent predictive inference. *Sankhyā Ser. A*, **46**, 166–185.

Lehman, R.S. (1955) On confirmation and rational betting. *J. Symbolic Logic*, **20**, 251–262.

Ramsey, F.P. (1931) Truth and probability. In F.P. Ramsey, *The Foundations of Mathematics and other Logical Essays*, edited by R.B. Braithwaite, ch. VII, pp. 156–198. London: Kegan, Paul, Trench, Trubner & Co.

Regazzini, E. (1987) De Finetti's coherence and statistical inference. *Ann. Statist.*, **15**, 845–864.

Schervish, M.J. (1995) *Theory of Statistics*. New York: Springer-Verlag.

Stone, M. (1976) Strong inconsistency from uniform priors. *J. Amer. Statist. Assoc.*, **71**, 114–125.

Sudderth, W.D. (1994) Coherent inference and prediction in statistics. In D. Prawitz, B. Skyrms and D. Westerstáhl (eds), *Logic. Methodology and Philosophy of Science IX*. Amsterdam: Elsevier Science.

Received March 2003 and revised March 2004