DICE HAVE NO MEMORIES, BUT I DO: A DEFENCE OF THE REVERSE GAMBLER'S BELIEF

B. O'NEILL[*] AND B. D. PUZA[**], *Australian National University*[***]

WRITTEN 10 OCTOBER 2004, REVISED 12 NOVEMBER 2004

Abstract

We investigate the problem of predicting the outcomes of a sequence of discrete random variables that are *almost uniform,* in the sense that they are generated from a random process that is designed to produce independent uniform outcomes but may not do so exactly. Using assumptions based around this notion we derive a useful stochastic ordering. We reject the *gambler's belief* as unsound and find that the *reverse gambler's belief* is the optimal prediction method. This method arises under a wide class of Bayesian models. One of the main contributions of this paper is that it uses only weak and intuitive prior assumptions and should therefore be more palatable to sceptics than existing Bayesian models.

BAYESIAN STATISTICS; EXCHANGEABILITY; STOCHASTIC DOMINANCE; GAMBLER'S FALLACY; REVERSE GAMBLER'S FALLACY; FREQUENT OUTCOME APPROACH.

1. The gambler's fallacy

If an abundance of heads come up on a coin, observers may be heard to assert that a tail is due; that it is more likely to come up than another head. This kind of belief or assertion is often called the *gambler's fallacy* (or the *Monte Carlo fallacy*), though it should more properly be called the *gambler's belief* when it is not accompanied by any proposed justification. The term is often applied outside a gambling context and, at its widest, the term can be ascribed to any belief that deviations from expected behaviour are likely to eventually be evened out by opposite deviations.

In statistical literature this belief has been attributed to a failure to understand statistical independence or to a misunderstanding of informal principles akin to Kolmogorov's strong law of large numbers (that assert the almost sure convergence of expectation and sample mean for independent random variables). Kahneman, Slovic and Tversky (1982) attribute the belief to a heuristic cognitive principle that they call the *representativeness heuristic*, where the probabilities of outcomes are estimated according to how they resemble quintessentially random-looking outcomes.

An abridged version of this paper is published as O'Neill and Puza (2005) In Defence of the Reverse Gambler's Belief, *The Mathematical Scientist* **30**,

[*] E-mail address: ben.oneill@anu.edu.au.
[**] E-mail address: borek.puza@anu.edu.au.
[***] School of Finance and Applied Statistics, Australian National University, ACT, 0200, Australia.

However, as Cowan (1969) rightly warns, the reasoning behind the gambler's belief is rarely explicated and is therefore imputed by the logician:

> Thus the argument, whether valid or fallacious, and if valid, whether sound or unsound, is for the most part an artificial creature, a construction the logician makes to help us decide what evidence to look for next, for esthetic purposes, or whatever.

In the absence of explicit reasoning for the belief it is rather presumptuous to say that the argument behind the gambler's belief is unsound, and it is outright incorrect to say that the argument is a logical fallacy (since the latter is an error in logical argument that is independent of the truth of its premises). Quite rightly, Cowan finds that the argument (if there is one), while often unsound, is not fallacious:

> The gambler's argument may very well be valid in the sense that *if* the premises are true, the conclusion is also true, but in many cases the premises required are simply not true. … [I]t is the logician's job to find out what is necessary to get the conclusion. Then we can see whether what would have to be true for the conclusion to be evidenced is in fact true.

In fact, the usual informality of the assertion of the gambler's belief and the lack of explicit reasoning sometimes makes it difficult to establish what is even being asserted. The assertion of the belief is usually framed or timed so as to suggest that it is at least a denial of the exchangeability of observations, but even this may be uncertain. Moreover, it is not always clear whether the gambler is looking to all previous tosses of the coin or just the last run of tosses as evidence for the conclusion that a tail is due. In this paper we take the former view as our conception of the gambler's belief: that is, we take the belief as being that the most likely next outcome is one of the outcomes that have occurred the least so far.

Following Cowan's suggestion we will demonstrate sufficient conditions that preclude the gambler's belief. We thereby find that the denial of these conditions is necessary for the belief. This gives some power to the analysis by attributing to the gambler a minimal belief that the gambler cannot escape by failure to properly explicate their reasoning. We endeavour to use only conditions that we believe are reasonable in the case of random processes used in gambling; this provides a reasonable argument against the gambler's argument however it may be framed.

2. The reverse gambler's fallacy

Continuing the coin tossing example, where we have observed an abundance of heads, some observers may also be heard to assert, in opposition to the gambler's belief, that another head is more likely than a tail. This opposite assertion has been called the *reverse gambler's fallacy* though again it should more properly be called the *reverse gambler's belief* when it is not accompanied by any proposed justification. Again, the term is often applied outside a gambling context and, at its widest, can be ascribed to any belief that deviations from expected behaviour are likely to continue.

In statistical literature this belief has also been attributed to a failure to understand statistical independence so that, to its detractors, the belief is generally considered to be a manifestation of the same kind of unsound thinking as the gambler's fallacy. However, such criticisms often themselves fail to understand statistical independence; many confuse causal independence with statistical independence and incorrectly use the former to assert the latter with simplistic (and flawed) arguments along the lines that "dice have no memories".

3. The Bayesian approach

The notion that successive outcomes are produced in an identical manner is, quite literally, described by the assumption of exchangeability. In random processes with a finite number of possible outcomes this gives rise to a multinomial model with unknown long-run proportions of outcomes. Under standard Bayesian multinomial models using either the reference prior or symmetric conjugate prior (both of which are special cases of the Dirichlet prior) it can be shown that the reverse gambler's belief arises as the correct posterior conclusion. The implications of these models are well known. What is not known is, in Cowan's words: "what is necessary to get the conclusion". This turns out not to be particularly instructive and so we instead find weaker sufficient conditions that also imply the reverse gambler's belief. Again we endeavour to use only conditions that we believe are reasonable in the case of random processes used in gambling. One of the main contributions of this paper is that it uses only weak and intuitive prior assumptions and should therefore be more palatable to sceptics than existing Bayesian models.

A note of caution is needed here. We do not intend to assert that advocates of the reverse gambler's belief are necessarily reasoning correctly; their reasoning may indeed be unsound or even fallacious. What we do intend is to show that their beliefs may be justifiable by the reasoning that *we* present. Again we stress that in the absence of explicit reasoning it is rather presumptuous to label a belief as unsound or fallacious.

4. Modelling almost uniform sequences

Throughout our analysis we will follow the notation in Bernardo and Smith (1994); in particular, we will not make any notational distinction between known values (constants) and unknown values (random variables) and we will not make any distinction between mass functions and density functions. This will not lead us into any trouble since the arguments that we present will not rely on continuity and can be considered totally in terms of general probability measures. The processes of interest can be described formally as follows:

DEFINITION 1 (Preliminary definitions): Let $\mathbf{x} \equiv (x_1, x_2, x_3, ...)$ be a sequence of values each with the same finite range $1, 2, ..., m$ and let $\mathbf{x}_k \equiv (x_1, x_2, ..., x_k)$ be the *observed outcomes*. To facilitate easier discussion we denote the *observed counts* by $\mathbf{n}(\mathbf{x}_k) \equiv (n_1, n_2, ..., n_m)$ with $n_i \equiv n_i(\mathbf{x}_k) \equiv \sum_{j=1}^{k} I(x_j = i)$, and we denote the *long-run proportions* by $\boldsymbol{\theta} \equiv \boldsymbol{\theta}(\mathbf{x}) \equiv (\theta_1, \theta_2, ..., \theta_m)$ with $\theta_i \equiv \theta_i(\mathbf{x}) \equiv \lim_{k\to\infty} n_i(\mathbf{x}_k)/k$. Finally, we let $p(x) \equiv p(x_{k+1} = x \mid \mathbf{x}_k)$ be the predictive probability of interest.

In modelling sequences where each outcome is produced in the same way it may be reasonable to assume that the order of the outcomes is always uninformative so that $\mathbf{x}$ is exchangeable. We note immediately that the exchangeability of $\mathbf{x}$ is fundamental to the results of the paper and that finite exchangeability will be insufficient; this may be taken as a possible criticism of the model, though we note that this assumption is weaker than, and is implied by the assumptions underlying existing models.

5. The common approach to modelling almost uniform sequences

When we are modelling processes that are designed to produce independent random outcomes it is common to assume that these processes generate exactly that. This is equivalent to assuming that $\boldsymbol{\theta} = \left(1/m, \ldots, 1/m\right)$ almost surely. In this case, if $\mathbf{x}$ is exchangeable then the elements of $\mathbf{x}$ are independent with $p\left(x\right) = 1/m$ so that, as is widely known, prediction is arbitrary.

6. An alternative approach to modelling almost uniform sequences

Contrary to the common approach, it is the authors' belief that some processes, such as the rolling of a die or the tossing of a coin produce outcomes that may not be perfectly uniform, in the sense that the long-run proportions of outcomes may differ. After all, in order to distinguish between the faces of a die or coin, these faces must be made to be non-identical so that the items themselves must be non-symmetric. This non-symmetry gives us plausible reason to believe that the long run proportions of outcomes may not be equal. We therefore diverge from the common approach by instead assuming that $\boldsymbol{\theta}$ is unknown and may therefore be biased towards some outcome. Of course, in the absence of information regarding which way the process is likely to be biased we may reasonably assume that $\boldsymbol{\theta}$ is exchangeable. Again it is the authors' belief that this assumption is reasonable in the context of processes such as the rolling of a die or the tossing of a coin. Of course, this assumption does not preclude the common approach. Rather, we diverge from the common approach by making some assumption that corresponds to our contemplation of the possibility of bias in the random process. Either of the following captures this notion:

DEFINITION 2 (Prior non-degeneracy): If *a priori* there is some positive probability that the elements of $\boldsymbol{\theta}$ are all positive and not all equal we say that $\boldsymbol{\theta}$ is *non-degenerate*.

DEFINITION 3 (Posterior non-degeneracy): If *a posteriori* there is some positive probability that the positive elements of $\boldsymbol{\theta} | \mathbf{x}_k$ are not all equal we say that $\boldsymbol{\theta} | \mathbf{x}_k$ is *non-degenerate*.

7. A useful stochastic ordering

Following Fishburn (1980) and suppressing notation of the order of domination we have the following definition of first order stochastic dominance.

DEFINITION 4 (First order stochastic dominance): If $P(x>t) \geq P(y>t)$ for all t we say that x *stochastically dominates* y (written as $x \succeq y$). Moreover, if $x \succeq y$ and additionally $P(x>t') > P(y>t')$ for some t' we say that x *strictly stochastically dominates* y (written as $x \succ y$).

THEOREM 1: If $\mathbf{x}$ and $\boldsymbol{\theta}$ are both exchangeable then:

(a) if $n_a \geq n_b$ then $\theta_a | \mathbf{x}_k \succeq \theta_b | \mathbf{x}_k$; and

(b) if $n_a > n_b$ and either $\boldsymbol{\theta}$ or $\boldsymbol{\theta} | \mathbf{x}_k$ are non-degenerate then $\theta_a | \mathbf{x}_k \succ \theta_b | \mathbf{x}_k$.

PROOF: Since $\mathbf{x}$ is exchangeable, it follows from the representation theorem of de Finetti (1937/1964) that:

$$p(\mathbf{x}_k | \boldsymbol{\theta}) = \prod_{i=1}^{m} \theta_i^{n_i(\mathbf{x}_k)} .$$

Let $\hat{\boldsymbol{\theta}}$ be the permutation of $\boldsymbol{\theta}$ with θ_u swapped with θ_v. Since $\boldsymbol{\theta}$ is exchangeable, for all $\theta_a > 0$ we have:

$$p(\hat{\boldsymbol{\theta}} | \mathbf{x}_k) = \frac{p(\mathbf{x}_k | \hat{\boldsymbol{\theta}}) p(\hat{\boldsymbol{\theta}})}{p(\mathbf{x}_k)} = \left(\frac{\theta_b}{\theta_a}\right)^{n_a - n_b} \frac{p(\mathbf{x}_k | \boldsymbol{\theta}) p(\boldsymbol{\theta})}{p(\mathbf{x}_k)} = \left(\frac{\theta_b}{\theta_a}\right)^{n_a - n_b} p(\boldsymbol{\theta} | \mathbf{x}_k) .$$

It follows that for all t:

$$\begin{aligned} P(\theta_a > t | \mathbf{x}_k) - P(\theta_b > t | \mathbf{x}_k) &= P(\theta_a > t \geq \theta_b | \mathbf{x}_k) - P(\theta_b > t \geq \theta_a | \mathbf{x}_k) \\ &= \mathrm{E}\left(I(\theta_a > t \geq \theta_b) \middle| \mathbf{x}_k\right) - \mathrm{E}\left(I(\theta_b > t \geq \theta_a) \middle| \mathbf{x}_k\right) \\ &= \mathrm{E}\left(I(\theta_a > t \geq \theta_b) \middle| \mathbf{x}_k\right) - \mathrm{E}\left(I(\theta_a > t \geq \theta_b) \left(\frac{\theta_b}{\theta_a}\right)^{n_a - n_b} \middle| \mathbf{x}_k \right) \\ &= \mathrm{E}\left(I(\theta_a > t \geq \theta_b) \left(1 - \left(\frac{\theta_b}{\theta_a}\right)^{n_a - n_b}\right) \middle| \mathbf{x}_k \right). \qquad (1) \end{aligned}$$

(a) If $n_a \geq n_b$ then the integrand in (1) is non-negative over the range of the integral (that is, the range of the indicator). It follows that $P(\theta_a > t | \mathbf{x}_k) \geq P(\theta_b > t | \mathbf{x}_k)$ for all t so that $\theta_a | \mathbf{x}_k \succeq \theta_b | \mathbf{x}_k$ which was to be shown.

(b) If $n_a > n_b$ then the integrand in (1) is strictly positive over the range of the integral (that is, the range of the indicator). Also, from Lemma 1 (see Appendix) we have $P(\theta_a > \theta_b | \mathbf{x}_k) > 0$ so that $P(\theta_a > t' \geq \theta_b | \mathbf{x}_k) > 0$ for some t'. It follows that $P(\theta_a > t' | \mathbf{x}_k) > P(\theta_b > t' | \mathbf{x}_k)$ for some t' so that $\theta_a | \mathbf{x}_k \succ \theta_b | \mathbf{x}_k$ which was to be shown. ■

COROLLARY 1 (Reverse gambler's belief): If $\mathbf{x}$ and $\boldsymbol{\theta}$ are both exchangeable then:

(a) if $n_a \geq n_b$ then $p(a) \geq p(b)$; and

(b) if $n_a > n_b$ and either $\boldsymbol{\theta}$ or $\boldsymbol{\theta} | \mathbf{x}_k$ are non-degenerate then $p(a) > p(b)$.

PROOF: Since $\mathbf{x}$ is exchangeable, it follows from the representation theorem that $x_{k+1} | \mathbf{x}_k, \boldsymbol{\theta} \sim \mathrm{Mu}(1, \boldsymbol{\theta})$ so that $p(a) - p(b) = \mathrm{E}(\theta_a - \theta_b | \mathbf{x}_k)$. The proof then follows from the stochastic dominance findings of Theorem 1 using Fishburn (1980). ■

Thus the exchangeability of $\mathbf{x}$ and $\boldsymbol{\theta}$ is sufficient to preclude the gambler's belief. Moreover, the addition of either non-degeneracy assumption then necessitates the reverse gambler's belief. This immediately gives us an optimal prediction method: namely, that we should predict one of the outcomes that has occurred the most in our observations. We call this method the *frequent outcome approach*.

8. The frequent outcome approach

It is easy to show —under the assumptions of Theorem 1(b)— that under the frequent outcome approach, unless we have observed all outcomes the same number of times (including none, as is the case *a priori*), the probability of successful prediction is strictly greater than $1/m$. This result calls into question the fairness of some so-called fair bets that are predicated on an assumption of independence between outcomes. The above results suggest that from a purely monetary perspective, we should accept so-called fair bets that are predicated on the outcomes of almost uniform sequences.

Of course, in commercial gambling situations, we are rarely presented with a fair bet[1]; rather, there is usually a profit loading. Whether or not we should accept such a bet depends upon weighing the degree to which we believe that the long-run proportions may be biased against the magnitude of the profit loading. Such an analysis requires further assumptions. Since the random processes in commercial gambling are designed to be very close to uniform, we would generally expect that the profit loading will outweigh any advantage gained by the frequent outcome approach in all but exceptional cases; thus the above analysis may be of little value. However, there are a vast number of statistical endeavors that involve processes that are designed to be uniform, not all of which involve adverse profit or utility weightings. The real value of the frequent outcome approach lies in these situations.

9. Conclusion

In scrutinising probabilistic beliefs, as with any application of logic, it is rather unfair to impute to the believer an unsound argument (which they have not explicated) as a basis for rejecting their beliefs. It may very well be that gamblers rely on the gambler's belief or the reverse under erroneous logical arguments. However, what we have shown is that in some situations the reverse gambler's belief arises as the optimal rational behaviour from reasonable assumptions about random processes. In particular, if we contemplate the possibility of bias in a random process *designed* to produce independent uniform random outcomes then we should reject the common approach that prediction is arbitrary in favour of the frequent outcome approach.

[1] …though this is not unheard of. In Australia for example, casinos and other gaming establishments offer a gambling game called *two-up* at fair odds (that is, the odds are fair under the common approach) every ANZAC day (25 April).

Appendix: Non-degeneracy conditions for Theorem 1

Theorem 1(b) requires some assumption that ensures that $P\left(\theta_a > \theta_b \middle| \mathbf{x}_k\right) > 0$. In this Appendix we show that the addition of either non-degeneracy assumption is sufficient for this condition.

LEMMA 1: If $\mathbf{x}$ and $\boldsymbol{\theta}$ are both exchangeable and either $\boldsymbol{\theta}$ or $\boldsymbol{\theta}|\mathbf{x}_k$ is non-degenerate, then $P\left(\theta_a > \theta_b \middle| \mathbf{x}_k\right) > 0$ for all $a \neq b$ such that $n_a > 0$.

PROOF: It is easily shown that the non-degeneracy of $\boldsymbol{\theta}$ implies the non-degeneracy of $\boldsymbol{\theta}|\mathbf{x}_k$; it is therefore sufficient to proceed for the latter case. Suppose —in contradiction to the Lemma— that $P\left(\theta_a > \theta_b \middle| \mathbf{x}_k\right) = 0$. It can be shown that, since $\mathbf{x}$ and $\boldsymbol{\theta}$ are exchangeable we then have $P\left(\theta_a \neq \theta_j, \theta_a > 0, \theta_j > 0 \middle| \mathbf{x}_k\right) = 0$ for all $j \neq a$. But since $n_a > 0$ this contradicts the non-degeneracy of $\boldsymbol{\theta}|\mathbf{x}_k$. By contradiction it follows that $P\left(\theta_a > \theta_b \middle| \mathbf{x}_k\right) > 0$ which was to be shown. ■

References

BERNARDO, J.M. AND SMITH, A.F.M. (1994) *Bayesian Theory*. John Wiley, Chichester.

COWAN, J. L. (1969) The gambler's fallacy. *Philosophy and Phenomenological Research*, **30(2)**, pp. 238-251.

DE FINETTI, B. (1937/1964). La prévision: ses lois logiques, ses sources subjectives. *Ann. Inst. H. Poincaré* **7,** 1-68. Translated and reprinted as DE FINETTI, B. (1980) Foresight; its logical laws, its subjective sources. *Studies in Subjective Probability* (eds KYBERG, H. E. & SMOKLER, H. E.), Dover, New York, pp. 93-158.

FISHBURN, P.C. (1980) Stochastic dominance and moments of distributions. *Mathematics of Operations Research* **5**, pp. 94-100.

FORTINI, S., LADELLI, L. AND REGAZZINI, E. (2000) Exchangeability, predictive distributions and parametric models. *Sankhyā: the Indian Journal of Statistics* **62A**, pp. 86-109.

KAHNEMAN, D., SLOVIC, P. AND TVERSKY, A. (1982) *Judgement under uncertainty: Heuristics and biases*. Cambridge University Press, Cambridge, pp. 3-22.