STUDI STATISTICI

N. 71

Reinforcement and Finite Exchangeability

Stephen G. Walker – Pietro Muliere

Istituto di Metodi Quantitativi – Università "L. Bocconi" – Milano, Febbraio 2003

Reinforcement and Finite Exchangeability

STEPHEN G WALKER AND PIETRO MULIERE

Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY
Istituto di Metodi Quantitativi, Bocconi University, Milano, Italy.

Abstract

This paper considers a finite sequence of Bernoulli experiments. A subjective modeler is interested in the notion of *reinforcing* observations from the past to predict new observations. She is also interested in the notion of the (finite) exchangeability of the sequence. The consequences of these two assumptions are investigated.

Key words: Bernoulli random variables, Finite exchangeability, Reinforcement.

1. Introduction. This paper is based on two ideas: reinforcement and finite exchangeability. The idea of reinforcement is the basis of many physical processes, such as spatial exploration, learned behavior, evolution of organisms, feedbacks in economics, among others. See, for example, Coppersmith and Diaconis (1986), Diaconis (1988), Pemantle (1988), Iosifescu and Theodorescu (1969). Reinforcement is a model that works to predict phenomena in different fields, such as biology, genetics, economics, physics and psychology. Reinforcement also plays a role in Bayesian nonparametric inference. In this respect the paper by Muliere, Secchi and Walker (2000) shows how the notion of reinforcement is the key to the understanding of many of the currently used Bayesian nonparametric prior distributions, such as the Dirichlet process (Ferguson, 1973), neutral to the right processes (Doksum, 1974) and Pólya trees (Mauldin et al., 1992). The idea of reinforcement is also the key to the paper of Walker (1998) and the paper of Muliere, Secchi and Walker (2003).

The second idea is that of finite exchangeability. It is well known that de Finetti's representation theorem (de Finetti, 1938) requires infinite exchangeable sequences and it easily seen to be false if the sequence is finite. Finite exchangeability does not guarantee the existence of a prior, see Diaconis (1977), Diaconis and Freedman (1980) and Feller (1971, pp. 228–230).

In this paper the connection between reinforcement and *finite* exchangeability is explored.

To pin down the idea of reinforcement and to establish the context of this paper, consider an infinite sequence of Bernoulli random variables, say $X_1, X_2, \ldots$. Observations are reinforced via the stipulation that, in the first place,

$$P\left(X_2 = 1 | X_1 = 1\right) > P(X_1 = 1)$$

and

$$P(X_2 = 0 | X_1 = 0) > P(X_1 = 0).$$

That is, having witnessed $\{X_1 = 1\}$, the subjective modeler of the sequence hypothesizes that the probability of seeing the outcome $\{X_2 = 1\}$ is greater than the probability for the event $\{X_1 = 1\}$. In general, if Σ_n is any sequence of 1's and 0's of length n, then the reinforcement plan is

$$P\left(X_{n+1} = 1 | X_n = 1, \Sigma_{n-1}\right) > P\left(X_n = 1 | \Sigma_{n-1}\right).$$

In fact it is straightforward to show that this condition implies

$$P(X_{n+1} = 1|X_n = 0, \Sigma_{n-1}) < P(X_n = 1|\Sigma_{n-1})$$

and

$$P(X_{n+1} = 0|X_n = 0, \Sigma_{n-1}) > P(X_n = 0|\Sigma_{n-1}).$$

A standard assumption of a modeler of a sequence is that the order in which the observations arise does not change the probability of the joint observations. That is, for any permutation σ on $\{1, \ldots, n\}$,

$$P(X_1, \ldots, X_n) = P\left(X_{\sigma(1)}, \ldots, X_{\sigma(n)}\right).$$

The celebrated de Finetti Representation theorem then guarantees the existence and the uniqueness of a prior distribution on the closed interval $[0, 1]$ such that

$$P(X_1, \ldots, X_n) = \int p^m (1-p)^{n-m}\, \pi(\mathrm{d}p),$$

for all n, where $m = \sum_{1 \le i \le n} 1(X_i = 1)$.

Suppose a beta prior is employed; $\pi(p) \propto p^{\alpha-1}(1-p)^{\beta-1}$, for $\alpha, \beta > 0$. For such a prior it is easily seen that

$$P(X_1 = 1) = \alpha/(\alpha + \beta) = \mu_1$$

and

$$P(X_2 = 1|X_1 = 1) = (\alpha + 1)/(\alpha + \beta + 1) = \mu_2.$$

Modelling the first two observations actually pins down the beta prior exactly so let us first only consider the reinforcement associated with the first two observations. It is easy to check that $\mu_2 > \mu_1$ and so Bayes updates automatically provides reinforcement. Allocating a (μ_1, μ_2) pair is equivalent to allocating an (α, β) pair via

$$\alpha = \frac{\mu_1(1-\mu_2)}{\mu_2 - \mu_1}$$

and

$$\beta = \frac{(1-\mu_2)(1-\mu_1)}{\mu_2 - \mu_1}.$$

2. Infinite exchangeability and reinforcement. In general, it is possible to see that the existence of a prior distribution (i.e. infinite exchangeability) implies reinforcement, which is established in the following lemma. In what follows we write, for example, $P(X_{n+1} = 1|X_n = 1, \Sigma)$ as $P(1|1, \Sigma)$.

LEMMA. For any sequence Σ of 1's and 0's,

$$P(1|1, \Sigma) - P(1|\Sigma) = \mathrm{Var}(p|\Sigma)/\mathrm{E}(p|\Sigma) > 0.$$

PROOF.

$$\begin{aligned}\mathrm{Var}(p|\Sigma) &= \textstyle\int p^2\, \pi(\mathrm{d}p|\Sigma) - \{\int p\, \pi(\mathrm{d}p|\Sigma)\}^2 \\ &= P(1, 1|\Sigma) - P(1|\Sigma)^2 \\ &= P(1|1, \Sigma)P(1|\Sigma) - P(1|\Sigma)^2 \\ &= P(1|\Sigma)\{P(1|1, \Sigma) - P(1|\Sigma)\},\end{aligned}$$

which completes the proof.

The amount of reinforcement can be understood via the conditional mean and the conditional variance; that is,

$$P(1|1,\Sigma) - P(1|\Sigma) = \mathrm{Var}(p|\Sigma)/\mathrm{E}(p|\Sigma).$$

If the sequence of observations are infinite then the assumption of exchangeability implies a prior. In this situation reinforcement must be constructed in such a way as to ensure a prior does exist. In the case of a finite exchangeable sequence there is more freedom for the type of reinforcement which can be done, since it only has to be compatible with exchangeability, not with the existence of a prior distribution.

3. Finite exchangeability and reinforcement. The aim in this section is to see how far reinforcement and finite exchangeability imply the existence of a prior. Consider a finite population of size N. It is shown that reinforcement and finite exchangeability imply the existence of a prior for $N \leq 3$.

N=2. Suppose we model two observations via the ideas of exchangeability and reinforcement. So, define

$$Q(i,j) = P(X_1 = i, X_2 = j)$$

for $i, j \in \{0,1\}$ and $Q(0,1) = Q(1,0)$ to ensure exchangeability. Reinforcement is characterized by

$$Q(1,1) > Q(1)^2,$$

where $Q(1,1) + Q(0,1) = Q(1)$. Thus $Q(0,1) < Q(0)\,Q(1)$, where $Q(0) = 1 - Q(1)$, and hence we can also show that $Q(0,0) > Q(0)^2$. We can now find a beta prior such that

$$Q(i,j) = \int p^{i+j}(1-p)^{2-i-j}\,\pi(\mathrm{d}p).$$

The parameters are given by $\mu_1 = Q(1)$ and $\mu_2 = Q(1,1)/Q(1)$ from which α and β can be calculated.

N=3. It is well known that we can write the probability Q as

$$Q = w_3^3 H_{000} + w_0^3 H_{111} + w_2^3 H_{001} + w_1^3 H_{011},$$

where H_{ijk} is an urn, sampled without replacement, consisting of the elements $\{ijk\}$. For this representation, see for example Crisma (1971), de Finetti (1969), Kendall (1967) and Diaconis and Freedman (1980).

So, for example, $Q(0,1,1) = Q(1,1,0) = w_1^3/3$, $Q(0,1) = Q(1,0) = w_2^3/3 + w_1^3/3$, and so on. If we retain $\mu_1 = P(1)$, $\mu_2 = P(1|1)$ and let $\mu_3 = P(1|1,1)$ then we can solve for $(w_0^3, w_1^3, w_2^3, w_3^3)$:

$$\begin{aligned} w_3^3 =&\ 1 - 3\mu_1(1-\mu_2) - \mu_1\mu_2\mu_3 \\ w_0^3 =&\ \mu_1\mu_2\mu_3 \\ w_2^3 =&\ 3\mu_1(1 + \mu_2\mu_3 - 2\mu_2) \\ w_1^3 =&\ 3\mu_1\mu_2(1-\mu_3). \end{aligned}$$

It is easy to show that these weights are all non-negative given that $\mu_3 > \mu_2 > \mu_1$.

The constraints from reinforcement can now be built in, having guaranteed exchangeability. Clearly, we have $\mu_3 > \mu_2 > \mu_1$. There is also obtain an upper bound for μ_3.

LEMMA. $\mu_3 < 1 - \{\mu_1(1-\mu_2)^2\}/\{\mu_2(1-\mu_1)\}$.

PROOF. From reinforcement, we have $P(1|1,0) > P(1|0)$. Hence

$$Q(1,1,0)/Q(1,0) > Q(1,0)/Q(0)$$

so

$$\frac{\mu_1\mu_2(1-\mu_3)}{\mu_1(1-\mu_2)} > \frac{\mu_1(1-\mu_2)}{1-\mu_1},$$

from which the result follows.

We now prove that we can find moments $c_i = \mathrm{E}(p^i)$ such that $c_3 = \mu_1\mu_2\mu_3$, $c_2 = \mu_1\mu_2$ and $c_1 = \mu_1$. The plan is to show that

$$(\mu_1, \mu_1\mu_2, \mu_1\mu_2\mu_3) \in M_3,$$

where

$$M_3 = \left\{(c_1, c_2, c_3) : c_i = \int p^i\,\pi(\mathrm{d}p), \quad \pi \in \mathcal{P}\right\}$$

and $\mathcal{P}$ is the set of probability measures on the Borel subsets of $[0,1]$. The following is taken from Skibinsky (1967). Define A to be the matrix

$$A = \begin{pmatrix} c_1 & c_2 \\ c_2 & c_3 \end{pmatrix}$$

and B to be the matrix

$$B = \begin{pmatrix} d_1 & d_2 \\ d_2 & d_3 \end{pmatrix}$$

where $d_1 = 1 - c_1$, $d_2 = c_1 - c_2$ and $d_3 = c_2 - c_3$. Then (c_1, c_2, c_3) is an interior point of M_3 if and only if $|A| > 0$ and $|B| > 0$. Some elementary algebra leads to

$$|A| = \mu_1^2\mu_2(\mu_3 - \mu_2)$$

and

$$|B| = (1-\mu_1)\mu_2\mu_1(1-\mu_3) - \mu_1^2(1-\mu_2)^2.$$

Establishing the positiveness of $|A|$ and $|B|$ is straightforward and is seen to coincide with the reinforcement constraints. Consequently, reinforcement and exchangeability imply, for $N \leq 3$, the existence of a prior distribution.

N=4. Let us add $\mu_4 = P(1|1,1,1)$. Again, using the same representation for Q and following Skibinsky (1967), the constraint on μ_4 to ensure the existence of a prior is given by

$$\mu_3 + \frac{\mu_1(\mu_3-\mu_2)}{\mu_3(\mu_2-\mu_1)} < \mu_4 < 1 - \frac{\mu_2(1-\mu_3)^2}{\mu_3(1-\mu_2)}.$$

The right hand-side inequality follows from the reinforcement constraint; via consideration of $P(1|1,0,1) > P(1|0,1)$. The left hand-side inequality does not follow from the reinforcement constraint. The three inequalities for μ_4 come from consideration of

1. $P(1|1,1,1) > P(1|1,1)$

2. $P(1|1,1,0) > P(1|1,0)$

3. $P(1|1,0,0) > P(1|0,0)$

Inequality 1 gives $\mu_4 > \mu_3$. Inequality 2 gives $\mu_4 < 1 - \{\mu_2(1-\mu_3)^2\}/\{\mu_3(1-\mu_2)\}$ and inequality 3 gives

$$\mu_4 > 2 - \frac{1}{\mu_3} + \frac{\mu_1(\mu_2\mu_3 - 2\mu_2 + 1)^2}{\mu_2\mu_3(1 - 2\mu_1 + \mu_2\mu_1)}.$$

So the the lower bound for μ_4 is the maximum of the terms from inequalities 1 and 3, which does not necessarily coincide with constraint for the existence of a prior.

4. The general case. Suppose we wish to undertake inference for a finite population with the assumptions of exchangeability and reinforcement. Let the population size be N. Thus we have

$$Q = \sum_{r=0}^{N} w_r^N H_{r,N-r},$$

where $H_{r,N-r}$ is the urn with r 0's and $N-r$ 1's. To complete the model we need to assign values to the probabilities $\{w_r^N\}_{r=1}^N$ such that $\sum_{r=0}^N w_r^N = 1$. A Bayesian would take

$$w_r^N = \binom{N}{r} \int p^{N-r}(1-p)^r \, \pi(\mathrm{d}p),$$

for some prior distribution π. If instead we only wish to assume reinforcement, then we assign values to $\{\mu_1, \ldots, \mu_N\}$, where $\mu_n = P(1|\mathbf{1}_{n-1})$, satisfying, it has to be said, some rather complicated constraints. Thinking about Q, the w_r^Ns can be found from the μ_ns via

$$w_r^N = \binom{N}{r} \sum_{k=0}^{r} \binom{r}{k} (-1)^k \lambda_{N-r+k},$$

where

$$\lambda_j = P(\mathbf{1}_j) = \prod_{i=1}^{j} \mu_i.$$

Now let $Q(k,l)$, for $k + l \le N$, be the probability of seeing k 0's and l 1's. Then

$$Q(k,l) = \sum_{r=k}^{N-l} w_r^N \binom{N-k-l}{r-k} \Big/ \binom{N}{r}.$$

See, for example, Cifarelli and Regazzini (1996). The constraints of reinforcement are then given by

$$Q(k,l+2)\,Q(k,l) > \{Q(k,l+1)\}^2$$

for all $k + l + 2 \le N$. We can find the implications of these constraints on the λ_j's sequentially. Note that we can write

$$Q(k,l) = \sum_{r=k}^{N-l} \sum_{j=0}^{r} \delta_{rj}\, c_{rkl}\, \lambda_{N-r+j},$$

where

$$\delta_{rj} = \binom{r}{j}(-1)^j$$

and

$$c_{rkl} = \binom{N-k-l}{r-k} \Big/ \binom{N}{r}.$$

From this it is possible to see that $Q(k,l)$ depends only on $\{\lambda_l, \ldots, \lambda_N\}$ and so let us write

$$Q(k,l) = \alpha_l^{-1}\lambda_l + \beta_{kl},$$

where

$$\alpha_l = \binom{N}{N-l}$$

and β_{kl} depends only on $\{\lambda_{l+1}, \ldots, \lambda_N\}$. So, given $\{\lambda_{l+1}, \ldots, \lambda_N\}$, the constraints on λ_l are

$$\lambda_l > \alpha_l \left\{ Q^2(k,l+1)/Q(k,l+2) - \beta_{kl} \right\}$$

for $\{k : k+l+2 \leq N\}$.

We also need to ensure that $w_r^N \geq 0$. This is straightforward to also develop sequentially and the constraint is

$$\lambda_{N-r} \geq \sum_{k=1}^{r} (-1)^{k+1} \binom{r}{k} \lambda_{N-r+k}$$

for all $r = 1, 2, \ldots, N-1$, which is of the form $\lambda_l \geq \gamma_l$ and γ_l depends on $\{\lambda_{l+1}, \ldots, \lambda_N\}$.

5. Conclusions. We have considered and discussed the assumptions of finite exchangeability and reinforcement starting from the fact that finite exchangeability alone does not guarantee the existence of a prior. We show that finite exchangeability and reinforcement imply the existence of a prior for $N \leq 3$. In order to ensure the existence of a prior for $N \geq 4$ we need finite exchangeability, reinforcement and some further constraints on the sequence $\{\mu_1, \ldots, \mu_N\}$.

Acknowledgments. The research of S. G. Walker is funded by a UK EPSRC Advanced Research Fellowship. The work of P. Muliere was partially supported by a travel grant from Bocconi University.

References

CIFARELLI, D.M. and REGAZZINI, E. (1996). de Finetti's contribution to probability and statistics. *Statistical Science* **11**, 253–282.

COPPERSMITH, D. and DIACONIS, P. (1986). Random walk with reinforcement. Unpublished Manuscript.

CRISMA, L. (1971). Sulla proseguibilita' di processi scambiabili. *Rend. Ist. Mat. Univ. Trieste* **3**, 96–124.

DE FINETTI, B. (1969). Sulla proseguibilita' di processi aleatori scambiabili. *Rend. Ist. Mat. Univ. Trieste* **1**, 53–67.

DIACONIS, P.(1988). Recent progress on de Finetti's notion of exchangeability. In *Bayesian Statistics 6*. Bernardo, J.M., De Groot. M.H., Lindley, D.V. and Smith A.F.M. (eds), Oxford University Press.

DIACONIS, P. and FREEDMAN, D. (1980). Finite exchangeable sequences. *Ann. Statist.* **8**, 745–764.

DOKSUM, K.A. (1974). Tailfree and neutral random probabilities and their posterior distributions. *Ann. Probab.* **2**, 183–201.

FELLER, W. (1971). *An Introduction to Probability Theory and its Application.* Vol II, John Wiley & Sons, New York.

FERGUSON, T.S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.* **1**, 209–230.

IOSIFESCU, M. and THEODORESCU, R. (1969). *Random Processes and Learning.* Springer-Verlag. New York

KENDALL, D.G. (1967). On finite and infinite sequences of exchangeable events. *Studia Sci. Math. Hungar.* **2**, 319–327.

MAULDIN, R.D., SUDDERTH, W.D. and WILLIAMS, S.C. (1992). Pólya trees and random distributions. *Ann. Statist.* **20**, 1203–1221.

MULIERE, P., SECCHI, P. and WALKER, S.G. (2000). Urn schemes and reinforced random walks. *Stoch. Proc. Appl.* **88**, 59–78.

MULIERE, P., SECCHI, P. and WALKER, S.G. (2003). Reinforced random processes in continuous time. *Stoch. Proc. Appl.* **104**, 117–130.

PEMANTLE, R. (1988). Phase transition in reinforced walk and RWRE on trees. *Ann. Probab.* **14**, 1229–1241.

SKIBINSKY, M. (1967). The range of the $(n+1)$th moment for distributions on $[0,1]$. *J. Appl. Prob.* **4**, 543–552.

WALKER, S.G. (1998). A nonparametric approach to a survival study with surrogate endpoints. *Biometrics* **54**, 662–672.