

Some Properties of Markovian Designs¹

Alcune Proprietà degli Esperimenti Markoviani

Alessandra Giovagnoli²

Department of Statistical Sciences, University of Bologna

giovagno@stat.unibo.it

Riassunto: Si possono definire markoviani quegli esperimenti sequenziali in cui la decisione sul come effettuare la prossima osservazione dipende solo dall'osservazione più recente. Molti dei disegni sequenziali per il confronto tra trattamenti presenti in letteratura sono markoviani. Qui si mostra come sfruttare questa osservazione per derivarne in modo unificato le proprietà, con particolare interesse per quelle asintotiche.

Keywords: Biased coin. Markov chains. Sequential experiments. Stationary distribution.

1. Markovian experiments

Sequentially designed experiments occur in biomedical or pharmaceutical practice and are deemed to be more efficient than non-sequential ones for any given sample size. The experiment may be *response-adaptive*, which means that the choice of the next design point is driven by previous responses as well as by previous allocations. Mathematically, a sequential experiment is a sequence $\{(X_n, Y_n), n \in \mathbb{N}\}$, where the X 's denote the “*design points*”, namely combinations of experimental conditions under the control of the experimenter, and the Y 's the responses. A probabilistic model for the responses is assumed and the design consists in an assignment rule of the point X_{n+1} given the σ -algebra generated by the “past”, $\mathfrak{S}_n = \sigma\{X_1, Y_1, \dots, X_n, Y_n\}$. This rule may be deterministic or, more likely, probabilistic, since in general a certain amount of randomization is desired. A frequent assumption is conditional independence of the observations, namely, with obvious notation,

$$Pr(Y_{n+1} = y \mid X_{n+1}, \mathfrak{S}_n) = Pr(Y_{n+1} = y \mid X_{n+1}) \quad \forall n. \quad (1)$$

An experiment can be defined to be *Markovian* if the allocation of the next design point depends only on the most recent outcome, i.e.

$$Pr(X_{n+1} = x \mid \mathfrak{S}_n) = Pr(X_{n+1} = x \mid X_n, Y_n) \quad \forall n. \quad (2)$$

Several sequential designs for treatment comparison satisfy condition (2). We show how this can lead to a unified approach to investigating their properties, especially asymptotic ones. These results complement other recent ones (Baldi Antognini and Giovagnoli, 2004; Bortot and Giovagnoli, 2004). The proofs are omitted or just sketched.

Fundamental Lemma Under (1) and (2), given a starting point x_0 , both $\{(X_n, Y_n)\}$ and $\{X_n\}$ are Markov chains, whether or not the experiment is response-driven.

¹This research was supported by the 2003 Research Project: “Asymptotic Properties of Stochastic Processes: methodological tools and physical, medical and engineering applications” of the Department of Statistical Sciences, University of Bologna.

²Address for correspondence: Dipartimento di Scienze Statistiche, via Belle Arti 41, 40126 Bologna, Italy

If the experiment is not response-adaptive there is nothing to prove; if it is, the proof is trivial for $\{(X_n, Y_n)\}$ and only a little more elaborate for $\{X_n\}$. The importance of this lemma lies in its generality, since it holds for any pair of sequences $\{X_n\}$ and $\{Y_n\}$ whatever the meaning of the random variables X and Y . In other words, Y_n could be replaced by an observable quantity which *depends only on the past history of the experiment*, and is not an actual physical “response”. The urn example in Section 3 will help clarify this. Definition (2) can be extended by replacing the sequence of design points $\{X_n\}$ by a sequence of design “statistics” $S_n = s(X_1, \dots, X_n)$ such that $\sigma(X_1, \dots, X_n) = \sigma(S_1, \dots, S_n) \forall n$. Then the Markov condition (2) becomes

$$Pr(S_{n+1} = s \mid \mathfrak{S}_n) = Pr(S_{n+1} = s \mid S_n, Y_n) \quad \forall n \quad (3)$$

and it can be shown that under (1) and (3) the results of the Fundamental Lemma still hold, in particular that $\{S_n\}$ is a Markov chain.

2. Asymptotic properties of Markovian experiments

If the transition probabilities (2) do not depend on n , not too severe conditions on the experiment usually ensure the existence of the stationary distribution π of the chain. On a discrete design space $\mathcal{D} = \{d_k, k \in \mathbb{Z}\}$, if $N_n(d_k)$ is the number of times the design point d_k has been employed in n steps, then for all k

$$\bar{N}_n(d_k) = N_n(d_k)/n \rightarrow \pi(d_k) \text{ as } n \rightarrow \infty \text{ almost surely} \quad (4)$$

by the Law of Large Numbers for Markov chains. Thus, the asymptotic behaviour of the experiment is described by π . Properties of the “*asymptotic experiment*” π are often easier to establish than small sample ones. By (4), for large samples the asymptotic experiment is approximated by the design point frequencies $\bar{N}_n(x) \forall x \in \mathcal{D}$ and the approximation is the more efficient the smaller the asymptotic variance of $\bar{N}_n(x)$, for all x .

In each of the examples of the next Section the chain is a random walk, with transition probabilities $p_k = \Pr(X_{n+1} = d_{k+1} \mid X_n = d_k) > 0$, $q_k = \Pr(X_{n+1} = d_{k-1} \mid X_n = d_k) > 0$ (some readjustments are needed for barriers). The following are useful results which can be easily proved.

Proposition 1 a) *The asymptotic design $\pi = \{\pi_k = \pi(d_k)\}$ is given by the equilibrium equations*

$$\pi_k = \pi_{k-1} \lambda_k \quad \forall k, \quad \text{with } \pi_0 = \left[1 + \sum_{i=1}^{\infty} \prod_{k=1}^i \lambda_k + \sum_{i=1}^{\infty} \prod_{k=1}^i 1/\lambda_{-k+1} \right]^{-1} \quad (5)$$

where $\lambda_k = p_k/q_k, \forall k$ if and only if $\sum_{i=1}^{\infty} \prod_{k=1}^i \lambda_k + \sum_{i=1}^{\infty} \prod_{k=1}^i 1/\lambda_{-k+1} < \infty$.

b) *If $\{\lambda_k\}$ is non-increasing, in particular if $\{p_k\}$ is non-increasing and $\{q_k\}$ is non-decreasing, then π is unimodal so there is a design point - or a value of the design statistics - which tends to occur most frequently.*

3. Some examples of Markovian experiments

Biased coin designs Efron (1971) proposed a design for the comparison of two treatments basing randomization on the hypothetical tossing of a biased coin with probability $p \geq 1/2$. An extension has been suggested by Baldi Antognini and Giovagnoli

(2004), the Adjustable Biased Coin Design (*ABCD*). Let $\delta_n = 1$ if the n -observation is allocated to treatment T_1 , $\delta_n = 0$ otherwise and D_n be the difference between the treatment allocations after n assignments, and put

$$\Pr(\delta_{n+1} = 1) = F(D_n), \quad (6)$$

where $F(\cdot)$ is a non-increasing function $F : \mathbb{R} \rightarrow [0, 1]$ such that $F(-x) = 1 - F(x)$. The *ABCD* generalises several other biased coin designs to be found in the literature, as well as Efron's. If we take $\{D_n\}$ to be the design statistic of interest for the *ABCD*, the sequence $\{D_n, n = 0, 1, \dots\}$ is Markovian. More precisely it is a random walk over the integers $\mathbb{Z}$, with transition probabilities $p_j = F(j)$, $q_j = F(-j) \forall j \in \mathbb{Z}$. Applying Section 2 we see that the asymptotic distribution $\pi_D(\cdot)$ of $\{D_n\}$ exists, $\pi_D(j) = \pi_D(-j)$, and, from the viewpoint of balance, the asymptotic behaviour of the *ABCD* is identical to the completely randomized design, namely $D_n/n \rightarrow 0$ almost surely as $n \rightarrow \infty$.

As regards randomness, an (inverse) indicator is predictability, i.e. the expected proportion G_n of correct guesses at stage n using an optimal guessing strategy. For the *ABCD*, $G_n = \frac{1}{n} \sum_{h=0}^{n-1} F(-|D_h|)$ and, by (4), $\lim_{n \rightarrow \infty} G_n = \sum_{h=0}^{\infty} F(-h) \pi_{|D|}(h)$ almost surely, where $\pi_{|D|}$ is the stationary distribution of $\{|D_n|\} : \pi_{|D|}(0) = \pi_D(0)$, $\pi_{|D|}(h) = 2\pi_D(h) \forall h > 0$.

Some biased coin designs, e.g. that of Wei (1978), give rise to a non-stationary Markov chain $\{D_n\}$, so the above results do not hold.

Urn Designs

Urns are a classical tool as a randomization device in sequential experiments to compare several treatments. Assume there are v urns $U_1, \dots, U_v$, standing for the v treatments, containing a total number of w balls; at each trial, eligible subjects arrive sequentially and, for each subject, a ball is drawn at random among all the balls of all the urns: if the ball is drawn from urn U_i ($i = 1, 2, \dots, v$), treatment T_i is allocated to the subject. Assume for simplicity only two treatments; define $\delta_k = 1$ if the k -th subject is allocated to T_1 at step k , 0 otherwise, and let $W_1(k)$ be the number of balls in urn U_1 after k steps. Thus $\Pr(\delta_{n+1} = 1) = \frac{W_1(n)}{w}$. At each step the selected ball is returned to either urn - so that the total number w of balls is constant - according to a prescribed probability rule. When the replacement rule is of the form $\Pr(W_1(n+1) = x | W_1(n), \delta_n)$, by letting $\delta_k = Y_k$ and $W_1(k) = X_k$ in the Fundamental Lemma it follows that $\{W_1(n)\}_{n \in \mathbb{N}}$ is a stationary random walk on $\{0, 1, \dots, w\}$. Ehrenfest urn designs are discussed in Baldi Antognini (2004), where it is shown that the convergence of an urn design towards a target is given by the convergence of the time average of the Markov chain $\{W_1(n)\}_{n \in \mathbb{N}}$ to the ergodic mean E_π , where π is the stationary distribution. The speeds of convergence of different types of Ehrenfeld designs can be compared.

Up-and-Down designs

Let $\mathcal{D} = \{d_k, k = 0, 1, \dots, M\}$ denote a finite set of possible treatments in increasing order. Let X_n be the treatment applied at the n -th step and Y_n be the corresponding observation taking values 0 and 1 for negative and positive responses respectively. Assume the probability of positive response to be an increasing function $Q(x)$ of the treatment level. Durham and Flournoy (1995) defined response-adaptive biased coin designs *BCD - I* and *BCD - II*, aimed at identifying a given quantile of the response function $Q(x)$. The main area of application is for escalating the dose of a potentially toxic drug in Phase I clinical trials. These experimental protocols have been extended in several ways and are now more commonly known as *Up-and-Down (U&D)* designs. In Giovagnoli and Pintacuda (1998) a general definition is given:

Choose a starting point $x_0 \in \mathcal{D}$. If $n \geq 1$ and $X_n = d_k$, then

Conditionally on $Y_n = 0$:	Conditionally on $Y_n = 1$:
$X_{n+1} = d_{k+1}$ with probability $\alpha_k > 0$	$X_{n+1} = d_{k-1}$ with probability $\beta_k > 0$
$X_{n+1} = d_{k-1}$ with probability $\beta'_k \geq 0$	$X_{n+1} = d_{k+1}$ with probability $\alpha'_k \geq 0$
$X_{n+1} = d_k$ with prob. $1 - \alpha_k - \beta'_k \geq 0$	$X_{n+1} = d_k$ with prob. $1 - \beta_k - \alpha'_k \geq 0$

By the Fundamental Lemma, $\{X_n\}$ is a Markov chain (in fact, a random walk) over $\mathcal{D}$ with transition probabilities $p_k = \alpha_k(1 - F(d_k)) + \alpha'_k F(d_k)$ and $q_k = \beta_k F(d_k) + \beta'_k(1 - F(d_k))$. The experiment converges and $\alpha_k, \alpha'_k, \beta_k, \beta'_k$ can be chosen so as to modify the asymptotic design (Giovagnoli and Pintacuda, 1998). Furthermore when $\alpha'_k = \beta'_k = 0 \forall k$, given condition $\frac{\alpha_k}{\alpha_k + \beta_k} = p$, ($0 < p < 1$), which targets the process to the quantile $x_{(p)}$ of interest, choosing α_k and β_k as large as possible, namely $\alpha_k = \min\{p/(1-p), 1\}$ and $\beta_k = \alpha_k(1-p)/p$, as suggested by Durham and Flournoy's *BCD - II* (1995), minimizes the asymptotic variance of $\bar{N}_n(d_k) \forall k$ (Bortot and Giovagnoli, 2004). This result can be extended to the case of an unlimited set of design points $\mathcal{D} = \{d_k, k = \dots - 1, 0, 1, \dots, M, \dots\}$ as follows: under the very mild assumption that $\inf_k F(d_k) = 0$ and $\sup_k F(d_k) = 1$ the asymptotic design exists by Proposition 1; such a design is the same for every α_k and β_k because of $\lambda_k = \frac{\alpha_k(1 - F(d_k))}{\beta_k F(d_k)}$ in (5). We can then apply Tierney's (1998) extension of Peskun's theorem.

4. Conclusions

Various types of sequential experiments for comparing treatments, which in the design literature are investigated in different contexts, satisfy a Markov condition. This paper shows that a unified study of their asymptotic behaviour may lead to fruitful results.

References

- Baldi Antognini A. (2004) Extensions of Ehrenfest urn designs for comparing two treatments, In A. Di Bucchianico, H. Lauter, and H.P. Wynn eds, *mODa 7-Advances in Model-Oriented Design and Analysis, Contribution to Statistics*, 69-75. Physica Verlag.
- Baldi Antognini A., Giovagnoli A. (2004) A new "biased coin" design for the sequential allocation of two treatments, *J. Roy. Statist. Soc. C*, 53(3), 1-14. Short version in *Proceedings of the XLI Scientific Meeting of the Italian Statist. Soc.*, 5-7 June 2002, 699-702.
- Bortot P., Giovagnoli A. (2004) Up-and-Down experiments of first and second order, *J. Statist. Planning and Inference*, (to appear).
- Efron B. (1971) Forcing sequential experiments to be balanced, *Biometrika*, 58, 403-417.
- Durham S.D. and N. Flournoy (1995) Up-and-down designs I: stationary treatment distributions, in *Adaptive Designs*, N. Flournoy and W.F. Rosenberger eds., IMS Lecture Notes, 25, Hayward CA, 139-157.
- Giovagnoli A., Pintacuda N. (1998) Properties of the frequency distribution induced by a general Up-and-Down method for estimating quantiles, *J. Statist. Planning and Inference*, 74, 51-63.
- Tierney, L. (1998) A note on Metropolis-Hastings kernels for general state spaces, *Annals of Applied Probability*, 8, 1-9.
- Wei, L.J. (1978) The adaptive biased coin design for sequential experiments, *Annals of Statistics*, 6(1), 92-100.