

A NOTE ON GITTINS INDICES FOR PHARMACEUTICAL RESEARCH

YOU-GAN WANG AND
JOHN GITTINS,* *University of Oxford*

Abstract

The Bernoulli/exponential target process is considered. Such processes have been found useful in modelling the search for active compounds in pharmaceutical research. An inequality is presented which improves a result of Gittins (1989), thus providing a better approximation to the Gittins indices which define the optimal search policy.

BANDIT PROCESS; DYNAMIC PROGRAMMING; GITTINS INDEX; STOPPING TIME; TARGET PROCESSES

1. Introduction

Suppose the random variable X has a density with respect to the measure μ on the real line, where μ is the Lebesgue measure with an additional atom at 0, so that $\mu(0) = 1$. There are two parameters involved in the density function, p and $\theta(0 \leq p \leq 1, \theta \geq 0)$. The density function has the following form:

$$f(x | p, \theta) = \begin{cases} 0 & \text{if } x < 0 \\ 1 - p & \text{if } x = 0, \\ p\theta e^{-\theta x} & \text{if } x > 0 \end{cases}$$

and will be described as a *Bernoulli/exponential* density. It represents approximately the distribution of potentially therapeutic activity within a family of chemical compounds if activity is measured on a suitable scale and high activity is more rare than low activity. The atom of probability at $x = 0$ may either simply represent compounds with no detectable activity, or correspond to a truncation of the activity scale at some threshold level up to which any observed activity is deemed to take the threshold value, which becomes zero simply by subtracting an appropriate constant from the scale. Truncating the scale in this way avoids some of the risks inherent in extrapolating a distribution which may not be of the precise form which has been assumed.

A process of sequential sampling from a real-valued random variable forms a *target* process if there is a target value which it is hoped will be exceeded by a sampled value as soon as possible. Given a collection of target processes to be sampled sequentially one at a time and all with the same target, the expected total number of values which need to be sampled before the target is reached is minimised by sampling always from the process for which the current value of the *Gittins index* is largest (see Gittins (1989), Theorem 3.2 and §6.2). The Gittins index $v(S)$ for a target process in the state (of knowledge) S is defined to be

$$(1) \quad \sup_{\tau > 0} \{P_{\tau}(S)/E(\tau)\},$$

where τ is a stopping time for the process which is measurable with respect to the sequence of

Received 15 April 1991.

* Postal address for both authors: Department of Statistics, University of Oxford, 1 South Parks Road, Oxford OX1 3TG, UK.

possible states of knowledge starting from state S , and $P_\tau(S)$ is the probability that at time τ the target has been reached (to attain the supremum τ clearly must not exceed the time at which the target is first reached).

In general the state of a target process is given by the current distribution expressing our beliefs about the probabilities of different possible values of the parameters (p and θ in the Bernoulli/exponential case). The description of the set of possible states is greatly simplified when the prior distribution for the parameters belongs to a parametric family of distributions which is closed under sampling. This happens when the prior distribution is *conjugate* to the sampling distribution (see Raiffa and Schlaifer (1961)).

Conjugate prior densities for the parameters (p, θ) of a Bernoulli/exponential distribution are those of the form

$$\pi(p, \theta \mid m, n_1, \Sigma, n_2) = \frac{\Gamma(m + n_1)}{\Gamma(n_1)\Gamma(m)} p^{n_1-1}(1 - p)^{m-1} \frac{\Sigma^{n_2}}{\Gamma(n_2)} \theta^{n_2-1} e^{-\Sigma\theta}.$$

If a value of x of X is sampled a prior density of this form becomes transformed under the operation of the generalised Bayes theorem to a posterior density of the same form. The parameters (m, n_1, Σ, n_2) become ($m + 1, n_1, \Sigma, n_2$) if $x = 0$, and ($m, n_1 + 1, \Sigma + x, n_2 + 1$) if $x > 0$, and represent the current state of knowledge about p and θ .

An *exponential* target process may be regarded as a Bernoulli/exponential target process for which p is known to take the value 1. There is just one unknown parameter θ , and conjugate prior densities are those of the form

$$\pi(\theta \mid \Sigma, n) = \frac{\Sigma^n}{\Gamma(n)} \theta^{n-1} e^{-\Sigma\theta},$$

so that state of the process is given by the parameters (Σ, n).

The properties of Bernoulli/exponential and exponential target processes and their use in pharmaceutical research are described by Gittins (1989), Bergman and Gittins (1985) and Jones (1975). The purpose of the present note is to establish an inequality which allows the much more easily calculated index for the exponential case to be used to approximate the Bernoulli/exponential index. Our inequality improves the one given by Gittins (1989), p. 182.

2. The improved inequality

The inequality relates index values for a Bernoulli/exponential target process T_1 in the state (m, n_1, Σ, n_2), and for an exponential target process T^2 in the state (Σ, n_2).

Theorem

$$v(m, n_1, \Sigma, n_2) \geq \frac{n_1}{m + n_1} v(\Sigma, n_2).$$

Proof. Any realisation sequence X_1, X_2, \dots , of T_1 defines a corresponding sequence of T_2 , obtained by deleting all the zeros from the T_1 sequence. Thus given a stopping time σ for T_2 there is a corresponding stopping time $\tau(\sigma)$ for T_1 , equal to the number of T_1 values which have been sampled when for the first time σ of these values are non-zero. Let

$$\tau^*(\sigma) = \begin{cases} 1, & \text{if } X_1 = 0 \\ \tau(\sigma), & \text{if } X_1 > 0 \end{cases}.$$

Thus $P_{\tau(\sigma)}(m, n_1, \Sigma, n_2) = P_\sigma(\Sigma, n_2)$, and since the event that T_1 reaches the target by $\tau(\sigma)$ is independent of the event $X_1 > 0$,

$$\begin{aligned} P_{\tau^*(\sigma)}(m, n_1, \Sigma, n_2) &= P(X_1 > 0)P_{\tau(\sigma)}(m, n_1, \Sigma, n_2) \\ (2) \qquad \qquad \qquad &= \frac{n_1}{m + n_1} P_\sigma(\Sigma, n_2). \end{aligned}$$

Also

$$\begin{aligned} E\{\tau^*(\sigma)\} &= E[E\{\tau^*(\sigma) \mid p\}] \\ &= E[P(X_1 = 0 \mid p)E\{\tau^*(\sigma) \mid p, X_1 = 0\} + P(X_1 > 0 \mid p)E\{\tau^*(\sigma) \mid p, X_1 > 0\}] \\ &= E[(1 - p)1 + pE\{\tau^*(\sigma) \mid p, X_1 > 0\}]. \end{aligned}$$

Now $E\{\tau^*(\sigma) \mid p, X_1 > 0\} = 1 + p^{-1}E(\sigma - 1)$, since $E\{\min_{j>0}(j : X_{i+j} > 0)\} = p^{-1}$, for any $i \geq 1$, and σ is independent of p . Thus

$$(3) \quad E\{\tau^*(\sigma)\} = E\{(1 - p) + p + E(\sigma - 1)\} = E(\sigma).$$

Now putting together (1), (2) and (3) we have

$$\begin{aligned} v(m, n_1, \Sigma, n_2) &= \sup_{\tau>0} \{P_\tau(m, n_1, \Sigma, n_2)/E(\tau)\} \\ &\geq \sup_{\sigma>0} [P_{\tau^*(\sigma)}(m, n_1, \Sigma, n_2)/E\{\tau^*(\sigma)\}] \\ &= \frac{n_1}{m + n_1} \sup_{\sigma>0} \{P_\sigma(\Sigma, n_2)/E(\sigma)\} \\ &= \frac{n_1}{m + n_1} v(\Sigma, n_2). \end{aligned}$$

This improves the corresponding inequality of Gittins (1989), which was similar but with the factor $(n_1 - 1)/(m + n_1 - 1)$ in place of $n_1/(m + n_1)$.

References

BERGMAN, S. W. AND GITTINS, J. C. (1985) *Statistical Methods for Planning Pharmaceutical Research*. Marcel Dekker, New York.
 GITTINS, J. C. (1989) *Multi-armed Bandit Allocation Indices*. Wiley, Chichester.
 JONES, D. M. (1975) Search Procedures for Industrial Chemical Research. Ph.D. thesis, Cambridge University Engineering Department.
 RAIFFA, H. AND SCHLAIFER, R. (1961) *Applied Statistical Decision Theory*. Harvard University Press, Cambridge, MA.