Available online at www.sciencedirect.com
ScienceDirect
European Journal of Operational Research 177 (2007) 1105–1112
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

Stochastics and Statistics

Optimal decision indices for R&D project evaluation in the pharmaceutical industry: Pearson index versus Gittins index

Michael A. Talias

School of Humanities and Social Sciences, International University Bremen, P.O. Box 750561, 28725 Bremen, Germany

Received 11 February 2004; accepted 8 January 2006
Available online 10 March 2006

Abstract

This paper examines issues related to various decision-based analytic approaches to sequential choice of projects, with special motivation from and application in the pharmaceutical industry. In particular, the Pearson index and Gittins index are considered as key strategic decision-making tools for the selection of R&D projects. It presents a proof of optimality of the Pearson index based on the Neyman–Pearson lemma. Emphasis is also given to how a project manager may differentiate between the two indices based on concepts from statistical decision theory. This work demonstrates and justifies the correct use of the Pearson index.

Keywords: Gittins index; Pearson index; Project prioritization; Research and development in pharmaceutical industry; Statistical decision theory

1. Introduction

This paper addresses the problem of evaluating research and development (R&D) projects in the pharmaceutical industry by using appropriate indices for determining priorities between potential R&D drugs within the Bayesian decision framework.

An R&D process is divided into a number of distinct phases that must be implemented sequentially (in a fixed order) and all succeed before a drug is marketed to yield any financial benefits. The phases and sequence of a research project's stages are explained by Gittins (1996, 1997). Typically, a project is characterized by the exploratory research stage in pharmaceutical laboratory, which is followed by the development and the marketing stages. There are several important features of the R&D process: It is considered a very uncertain process, it needs at least 15–30 years from the time the research starts until a drug is marketed, and the initial investment is large. Towards the end of each stage the results from the current stage will become known and a decision should be made regarding the fate of the project.

E-mail address: m.talias@iu-bremen.de

0377-2217/$ - see front matter
doi:10.1016/j.ejor.2006.01.011

A project manager must continually make decisions concerning the treatment of projects. This means that the existing resources should be allocated on the basis of the available information concerning each project. For example, the project manager has to decide which new projects should be funded and which existing projects should be continued or abandoned. However, information is gained at considerable cost and decisions must be made under uncertainty. Thus, a continued corrective action based on new information is at the heart of the successful R&D management. For all these reasons, a statistician may recommend a Bayesian decision approach.

R&D project selection models have a history of more than 50 years. The definition of the problem is not unique and is largely dependent on methods that were fashionable at various times in the past. These range from classical methods, such as profiles, checklists, scoring models, and economic indices, to decision theory or analysis models. Models based on mathematical techniques, such as linear programming, non-linear programming and dynamic programming, were developed in the past. A more recent technique is the application of the real option pricing analysis. A summary of R&D selection methods as applied in the pharmaceutical industry may be found in chapter 4 in Bergman and Gittins (1985).

In this paper, we consider profitability indices which can be viewed as a solution to decision problems related to the valuation of research and development projects. A profitability index is given by Pearson (1972) who suggested that the sequential characteristics of decision analysis must be incorporated into a profitability index as follows: Consider a project with k stages with fixed order. Each stage of the project is denoted by the index j for $j = 1,\ldots,k$, with associated cost c_j, conditional probability of success given success at the previous stages p_j, and final reward R to be gained when the last stage is completed successfully. Assuming that the reward R and costs c_j for all j are discounted figures at some interest rate, then the Pearson index is defined as

$$\frac{\text{Expected net value}}{\text{Expected cost}} = \frac{R\prod_{j=1}^{k} p_j - \sum_{j=1}^{k} c_j \prod_{r=0}^{j-1} p_r}{\sum_{j=1}^{k} c_j \prod_{r=0}^{j-1} p_r}, \tag{1}$$

where $p_0 = 1$.

To set up a portfolio of projects using a profitability index, the index value of a project is calculated on the basis of the available information concerning the project. Projects are then ranked in decreasing order of their index values. The projects to be included in the portfolio must have a greater index value than those projects that may possibly be rejected: Such a criterion may help us to achieve the highest possible income for a given expenditure expressed in terms of the total expected cost.

Pearson defined his index as the ratio of expected net present value of the reward of an R&D project to its expected development cost without making any reference either to the pharmaceutical industry or to the R&D manager's objectives. When the manager's goal is to maximize the net present value of the expected profit stream from any project undertaken, it is not clear why one should try to maximize the ratio of expected reward to the expected cost, and not, for example, to seek maximization of their long-run difference.

The history of this index is not long. Gittins (1996, 1997) reconsidered this index when he developed a stochastic model to investigate how many scientist should be allocated at each successive stage to maximize profitability. Senn (1996, 1997, 1998), too, discussed its merits, the relation of the index with games, and whether some option value is captured by the index. Regan and Senn (1997) accounted for the use of the Pearson index as the main approach in project evaluation in drug development from a pharmaceutical company perspective. Zipfel (2003) discussed how to model the probability-cost-profitability architecture of portfolio management in the pharmaceutical industry using the Pearson index.

The objectives of this paper are twofold: first, to develop a way of looking at the Pearson index as an optimal decision rule in the project selection process, and second, to compare the Pearson index with Gittins' index (Gittins, 1979) as a means of evaluating R&D projects to indicate how the two selection methods differ.

The paper is organized as follows: In Section 2 notation is given and the problems of parallel and sequential selection are defined. Parallel selection is approached in Section 3. Optimality of the Pearson index is justified using the Neyman–Pearson lemma in Section 4. Section 5 deals with the applications of sequential analysis in project prioritization. The multi-armed bandit problem is introduced and the Gittins index is studied. Concluding remarks are discussed in Section 6.

2. The R&D selection problem

2.1. An R&D project and some notation

An R&D project can be modelled as a two-action Markovian decision process with a finite number of stages. Let $x_i(t)$ be the state of the ith project at time t. At each stage of the process, the decision maker can either select to stop, in which case the process is terminated, or to pay some cost $C = C(x_i)$ and continue. When the ith project is selected, the expected reward $E[R(x_i)|x_i]$ is to be received at the completion of the current stage of the project. The expected development cost depends on the current stage of the project and is denoted by $E[C(x_i)|x_i]$.

2.2. The problem of parallel and sequential selection

Suppose that we have two varieties of pharmaceutical projects to be tested. It is expected that some of them will perform better than others, but we do not know which projects are the most promising. Information for each candidate project concerns its expected development cost $E[C(x_i)|x_i]$ and expected reward $E[R(x_i)|x_i]$. The R&D selection problem is concerned with how to evaluate and identify the best subset among several proposed projects under some constraints. A decision as to whether a given project is to be engaged is made on the basis of the available information concerning this given project. Let S be the set of projects under consideration and suppose that a decision rule is applied to partition S into, at most, three exclusive and exhaustive regions S_1, S_2, and $S\backslash(S_1 \cup S_2)$, separating projects into those which will be developed, not developed and the projects which are left "unclassified" and for which further information is sought before a final allocation is made.

There are two important special cases of this problem:

In parallel selection methods, the partition $S\backslash(S_1 \cup S_2)$ is an empty set. One could assume that the project scientist will not envisage the possibility of the need for more information and therefore projects clearly are separated into two groups only.

However, in an alternative approach, once a project has been chosen for development, the development of the remaining projects is considered to be ceased temporarily, and at any subsequent decision time the most recently chosen project may be replaced by any temporarily ceased project. This is a case of sequential decision processes where all projects are left "unclassified". An example of a sequential decision process is the multi-armed bandit problem which is formulated in Section 5.2.

3. Parallel selection

The parallel selection method is addressed by the following question. Let G be the set of the n candidate projects for development (n is assumed to be constant). It is required that each project is allocated to one of two partitions of G: these are G_D, the projects which will be developed, and G_U, the set of projects will remain undeveloped. Parallel selection means that no experimentation is currently required to gain additional information for each project and thus current knowledge of each project is considered satisfactory and is utilized for decision making. Decision rules are established on the basis that a large number of projects will need to be classified in the future, and hence the decision variable must be chosen in such a way as to minimize the expected consequences of mistakes made in the allocation procedure. Given the current information, how should one separate these projects into two sets in order to decide which projects will or will not be developed?

The above problem can be approached as a classification problem on the basis that each individual project must be assigned to either G_D or alternatively to G_U, even at the risk of error. The decision maker may commit two kinds of mistakes, according to the population to which the decision maker erroneously allocates a given project. Suppose that each project is characterized by a multivariate random variable X from space W. Having observed $X = x$ for a given project, then $f_1(x)$ and $f_0(x)$ represent the frequency function that the project belongs to G_D and G_U, respectively. We wish to set up a boundary so that as many successful projects as possible lie on one side and as many as possible unsuccessful on the other side. Let A be a region such that the two errors are of the same value as follows:

$$\int_A f_1(x)\mathrm{d}x = \int_{W-A} f_0(x)\mathrm{d}x = 1 - \int_A f_0(x)\mathrm{d}x. \tag{2}$$

We wish to minimize the total error, or equivalently, minimize one of the types of error, say $\int_A f_1(x)\mathrm{d}x$. The problem is then to find an unconditional minimum of

$$\int_A [f_1(x) - \lambda\{f_0(x) + f_1(x)\}]\mathrm{d}x,$$

or, equivalently, of

$$\int_A \{\mu f_1(x) - f_0(x)\}\mathrm{d}x,$$

where the constants λ or μ are determined from the constraint (2). The minimization is achieved by taking into A only those points for which $\mu f_1(x) - f_0(x) < 0$. Therefore, the boundary is given by

$$f_0(x)/f_1(x) = \mu, \tag{3}$$

which is the likelihood ratio. This is a reasonable criterion and remind us of the Neyman–Pearson lemma. An example of such a criterion is the Pearson index.

4. The Pearson index and the Neyman–Pearson lemma

4.1. Interpretation of the Pearson index

The meaning of the Pearson index is the expected net reward per unit expected cost which can be seen as an expected cost-effectiveness criterion. More precisely, given $n(\in \mathbb{N})$ of different potential new drugs at different stages in the testing process, projects should be chosen in such a way as to maximize the total expected profit without the total expected cost of the selected projects exceeding the total budget B, that is

$$\text{Maximize} \quad \sum_{i=1}^{n} y_i E[R(x_i)|x_i] \tag{4}$$

subject to

$$\sum_{i=1}^{n} y_i E[C(x_i)|x_i] \leqslant B,$$
$$y_i \in \{0,1\}, \quad i = 1,\ldots,n, \tag{5}$$

where y_i is a decision variable that maximizes (4) when and only when

$$y_i = \begin{cases} 1 & \text{if } E[R(x_i)|x_i] \geqslant \lambda E[C(x_i)|x_i], \\ 0 & \text{if } E[R(x_i)|x_i] < \lambda E[C(x_i)|x_i], \end{cases} \tag{6}$$

for some λ, $0 < \lambda < \infty$.

Selecting projects having an index value equal to or greater than an arbitrary positive value λ results in maximizing the total expected net benefit subject to a budgetary constraint on the total expected cost of the selected projects. However, the project manager may need to apply a randomized solution to consume the total expected cost B. This implies that the last project may be selected partially, as will be discussed in the next section. The constant λ plays the role of a dual variable in a linear programming problem or as a Largrange multiplier in control theory. In the Pearson index, time consideration is not taken into account in the sense that the selected projects are expected to yield rewards at the same time.

It is perhaps more interesting to a decision maker in pharmaceutical industry to know that his choices are expected to yield the fewest possible mistakes. This seems to be a property of the index and it is the main conclusion if one recalls the Neyman–Pearson lemma.

4.2. An application of the Neyman–Pearson lemma

A similar constrained maximization problem occurs in the Neyman–Pearson theory of hypothesis testing (Neyman and Pearson, 1933). The simple null hypotheses $H_0 : X \sim f_0$ is tested against the simple alternative $H_1 : X \sim f_1$, based on the observed values of some random variable X which takes values in a space S. Define any critical region $W \subset S$ such that (i) if $x \in W$, reject H_0 with error probability $\alpha \equiv P(W|H_0)$ and (ii) if $x \in S - W$, fail to reject H_0 with error probability $\beta \equiv P(S - W|H_1)$. The performance of a test is characterized in the two error probabilities of a test. A test is ideal if the sizes of both errors are as small as possible. The best size critical region can be determined using the Neyman–Pearson lemma.

Lemma 1 (Neyman–Pearson lemma). *For any test of size α, power $1 - \beta$, and a decision function $\Phi(x) \in \{0, 1\}$, there exist a constant k, $0 \leqslant k < \infty$, such that the maximum of*

$$1 - \beta = 1 - P(S - W|H_1) = \int_X \Phi(x) f_1(x) \mathrm{d}x, \tag{7}$$

subject to

$$P(W|H_0) = \int_X \Phi(x) f_0(x) \mathrm{d}x \leqslant \alpha \tag{8}$$

is attained through the critical region W^* and it occurs when

$$\Phi(x) = \begin{cases} 1 & \text{if } f_1(x) \geqslant k f_0(x), \\ 0 & \text{if } f_1(x) < k f_0(x), \end{cases} \tag{9}$$

and $P(W^*|H_0) = \alpha$.

The following remarks are the key results of this section.

Remark 1. The Pearson index is an optimal criterion for selecting projects in a "predata" spirit. On the basis of the Neyman–Pearson lemma, it is reasonable to suggest choosing those projects for development whose ratio of expected reward to expected cost ratio is "big" enough (greater than k).

Remark 2. In case the total expected cost of the selected projects in (9) is less than B, a randomized decision rule may be applied. The sample space S is partitioned into three non-overlapping regions, each region characterising the type of decision-based on the observed value of x. The three possible decisions are

$$\Phi(x) = \begin{cases} 1 & \text{if } f_1(x) > \lambda f_0(x), \\ \gamma(x) & \text{if } f_1(x) = \lambda f_0(x), \\ 0 & \text{if } f_1(x) < \lambda f_0(x), \end{cases} \tag{10}$$

for some λ, $0 < \lambda < \infty$, and $0 < \gamma(x) < 1$. $\Phi(x) = \gamma(x)$ means that the project is chosen with probability $\gamma(x)$. If $\lambda f_0(x) = f_1(x)$ has positive probability, this is not unique. The portfolio will possibly include enough projects with ratio equal to λ to spend exactly B. Comparing solutions (9) and (10), the latter would just include an additional project of which only a fraction $\gamma(x)$ of its total cost is being paid.

Remark 3. In Bayesian methods, the cut-off point depends only on the utility ratio and the prior probabilities. This contrasts with the classical recommendations of choosing the cut-off point to achieve some desired value of size α. Let $P(H_0) = \pi_0$ and $P(H_1) = \pi_1$ be the prior probabilities which express the scientist's uncertainty based on his belief, experience and past data, that a project is to be classified as a successful one or not. Let $U = u_{ji}$ be the utility of action j when $\theta = \theta_i$ for $i, j = 0$ or 1. The overall expected utility of using decision rule d is given by $\pi_0 U(d, \theta_0) + \pi_1 U(d, \theta_1) = (\pi_0 u_{00} + \pi_1 u_{11}) - [\pi_0(u_{00} - u_{10})\alpha + \pi_1(u_{11} - u_{01})\beta]$. The optimal policy is one which minimizes $\pi_0(u_{00} - u_{10})\alpha + \pi_1(u_{11} - u_{01})\beta$ which is equivalent to say that we minimize an arbitrary linear combination $k_1\alpha + k_2\beta$, for $k_1, k_2 > 0$. In the above case the rule (9) becomes

$$\Phi(x) = \begin{cases} 1 & \text{if } f_1(x)/f_0(x) \geqslant \pi_0(u_{00} - u_{10})/\pi_1(u_{11} - u_{01}), \\ 0 & \text{if } f_1(x)/f_0(x) < \pi_0(u_{00} - u_{10})/\pi_1(u_{11} - u_{01}). \end{cases} \tag{11}$$

By using a Bayesian argument, we have justified theoretically a rule of the form presented in (9) and established the appropriate value of k of this rule or for the bound μ in (3). This is equivalent to the reasonable procedure of maximizing the posterior probability that a project with observed vector x satisfies a certain hypothesis.

One may argue that a hypothesis statement is either true or false. Can we really associate a hypothesis with an R&D project? We do this in the sense that we maximize the probability of making a correct decision in an uncertain environment.

Viewing the project selection problem as a hypothesis testing problem, the decision making as seen in Eq. (10) may be thought of as a special case of the following scenario: Reject the null hypothesis H_0 for a satisfactorily large likelihood ratio; accept H_0 for satisfactorily small likelihood ratio (smaller cut-off point); and for intermediate values of the likelihood ratio, the decision maker collects more data. This kind of test is known as the sequential probability ratio test.

5. Sequential analysis and project selection

5.1. The sequential probability ratio test (SPRT)

The term sequential analysis refers to methods for testing hypothesis or estimating parameters when the sample size is not fixed in advance but is determined during the statistical experiment (procedure) by criteria which depend on the observations as they occur. In other words, the time which observations are terminated is random, namely, stopping time.

Under the scenario of testing a simple hypothesis against another simple hypothesis optimal test are known. In Section 4.2 we presented the Neyman–Pearson lemma where the error probability β is minimized for a given α and sample size n, which remind us of the Pearson index.

In Wald's Sequential probability test (Wald, 1945) the size α and the power β of a test are fixed in advanced and the number of observations is random. The investigator chooses constants $0 < A < 1 < B < \infty$ and defines

$$\lambda_n = \prod_{i=1}^{n} \frac{f_1(x_i)}{f_0(x_i)}.$$

Let $T = \min\{n \in \mathbf{N} : \lambda_n \notin (A, B)\}$ be the random time at which the sequence λ_n leaves the open interval (A, B) for the first time. Sampling is stopped at time T and reject H_0 if $\lambda_T \geqslant B$, and accept H_0 if $\lambda_T \leqslant A$. The error probabilities of the test are defined by $\alpha = P(\lambda_T \geqslant B|H_0)$ and $\beta = P(\lambda_T \leqslant A|H_1)$. Suppose that α^* and β^* are the two error probabilities for any other test, sequential or not, and let T^* be the corresponding sample size. Among all tests for which $\alpha^* \leqslant \alpha$ and $\beta^* \leqslant \beta$, the SPRT is optimal in the sense that $E\{T^*|H_0\} \geqslant E\{T|H_0\}$ and $E\{T^*|H_1\} \geqslant E\{T|H_1\}$. The optimal character of the sequential probability ratio test is due to Wald and Wolfowitz (1948) who proved that the SPRT simultaneously minimizes the expected sample sizes under both hypotheses for a given α and β.

A truncated variation of this problem is the bandit problem. A priority index which may be used to choose projects sequentially is the Gittins index, which solves the multi-armed bandit problem.

5.2. The Gittins index

The multi-armed bandit problem is concerned with the question of how to dynamically allocate a single resource among several alternatives as follows. Suppose that there are N independent projects, each divisible into stages, and only one project can be worked on at a time. Project i has state $x_i(t)$ at time $t = 1, 2 \ldots$, for $i = 1, 2, \ldots, N$. At each time t, one must operate exactly one project. If project i is selected, it gives immediate reward $R_i(x_i(t))$, a function of the current state $x_i(t)$ of the chosen project i, and its state $x_i(t)$ then changes according to a stationary Markov transition rule. The states of the other projects remain frozen. The states

of all projects are observed, and the problem is to schedule the order in which the projects are operated so as to maximize the expected present values of the sequence of immediate rewards, that is,

$$E_\pi\left\{\sum_{t=1}^{\infty}\sum_{i=1}^{N}\alpha^t R_i(x_i(t))\Delta x_i(t)\right\}, \tag{12}$$

where $0<\alpha<1$ is a fixed discounting factor, $\Delta x_i(t)\in\{0,1\}$ which takes value one if the project i, whose state is $x_i(t)$, is chosen at time t and zero otherwise. The strategy π is the optimal policy for choosing between projects and the maximization is over the vector $(\Delta x_1(t),\Delta x_2(t),\ldots,\Delta x_N(t))$ subject to $\sum_{i=1}^{N}\Delta x_i(t)=1$, for every t.

What makes the problem attractive is its solution as suggested by Gittins and Jones (1974) and analyzed further by Gittins (1979), that the problem is solved by a priority index, that is, to associate to each project i an index $v_i(x_i(t))$, which is a function only of its state, and at each time operate the project with the largest current index. The Gittins index has the following form:

$$v_i(x_i)=\sup_{\tau>1}\frac{E\left\{\sum_{t=1}^{\tau-1}\alpha^t R_i(x_i(t))\middle|x_i(1)=x_i\right\}}{E\left\{\sum_{t=1}^{\tau-1}\alpha^t\middle|x_i(1)=x_i\right\}}, \tag{13}$$

where the maximization is over the set of all stopping times $\tau>1$. Its numerator is the expected discounted reward for a certain project up to the chosen stopping time τ, and its denominator is the expected discounted time up to the stopping time τ. This index is interpreted as the maximum expected discounted reward per unit of expected discounted time. Gittins called his index a *Dynamic Allocation Index*. This index is sequential in the sense that it needs to be recalculated at each decision point in order to guide the reallocation process.

The theory of Gittins' index, and more generally of any priority index, is based on the idea that the index for each project depends on the past history of the given (or chosen) project only and not on the history of other projects. Effort is allocated to the project with the highest current index value. Bergman and Gittins (1985) referred to two different versions of dynamic allocation index related to the Pearson index. These are given by (14) and (15):

$$\frac{Rp_1p_2p_3-c_1-c_2p_1-c_3p_1p_2}{1-(1-p_1)D_1-p_1(1-p_2)D_2-p_1p_2D_3}. \tag{14}$$

The index given by (14) is for a three-stage project, with final reward R, costs c_1,c_2,c_3 for stage one, two, and three, respectively, and probability of success of each stage p_1, p_2 and p_3. D_i denotes the present discounted value of one monetary unit at time t_i for $i=1,2,3$. The numerator is the expected net profit and the denominator expresses the discounted probability that all stages of the project will be implemented. The above can be approximated by the following index as the discounting rate approaches zero:

$$\frac{Rp_1p_2p_3-c_1-c_2p_1-c_3p_1p_2}{t_1+t_2p_1+t_3p_1p_2}. \tag{15}$$

One should observe that the Pearson index (1) and the index given by (15) differ only in their denominators. However, there are differences between the Pearson and Gittins indices and a comparison is made below.

6. Concluding remarks

The primary objective was to provide theoretical justification for the use of the Pearson index as an optimal decision tool for the parallel classification of R&D projects.

The basic model viewed the project selection as a statistical hypothesis problem in a Bayesian framework. A hypothesis statement is either true or false. In an uncertain environment this is not known and one has to make decisions-based on the available information as to whether the alternative hypothesis is valid or not. In project prioritization, the uncertainty refers to the probability of success of a given project. A project will either succeed or fail, a fact that will be ascertained after the event. The key issue is how the decision maker chooses projects at an early stage (uncertain environment) in order to maximize the probability of stating correctly which projects will eventually be successful given a budget constraint.

The Pearson index can be viewed as the right decision rule used when a set of projects needs to be divided into two groups: the projects that are going to be implemented and the projects that will never be considered again. This is equivalent to maximizing the posterior probability that a project is classified correctly when a budget constraint is imposed.

Differences of the two indices are highlighted by the fact that the optimality of the Pearson index is a result of statistical application of the Neyman–Pearson lemma, whereas the Gittins index may be equivalently regarded as a sequential testing procedure with properties of the SPRT.

The Gittins index solves a sequential selection problem since the optimization is with respect to time and the decision maker is looking for the stopping time that maximizes the expected reward per unit expected discounted time for an arbitrary time period. The allocation is varied in time to meet changing conditions.

The decision maker who uses the Pearson index is not concerned about this type of maximization and selects projects without maximizing the objective function with respect to time. Therefore, the Pearson index maximizes the expected reward instantly, as opposed to the Gittins index, which maximizes the expected reward sequentially.

The difference is due to the principle of forward induction, which in its simplest form is known as a one-step look ahead policy. The decision maker applies the one-step look ahead policy when he compares stopping immediately with stopping after a period. In sequential selection, the decision maker has to predict how long he has to use a certain project in order to achieve the maximum attainable reward within that period.

A shared plus point is that both indices can be supported using concepts from statistical decisions theory; consequently, this is one reason why they should be used in the pharmaceutical industry. Managers need to recognize that the option to pause a project does exist and is valuable. When they decide to look beyond the dilemma to develop a project continuously or abandon it, they are bound to appreciate the option to pause a project as well. A manager can combine the two methods by selecting a subset of projects using the Pearson index and then prioritizing them sequentially by using the Gittins index. In addition, it might be useful for the analyst to follow Gittins (1997) advice to allocate more scientists to a project even beyond the point where output per man is greatest.

Acknowledgements

We thank Professor A. Philip Dawid and Professor Stephen Senn for their insightful comments. This work is based on the Ph.D. thesis of the author at the University College London.

References

Bergman, S.W., Gittins, J.C., 1985. Statistical Methods for Pharmaceutical Research Planning. Marcel Dekker.

Gittins, J.C., 1979. Bandit processes and dynamic allocation indices. Journal of the Royal Statistical Society Series B (14), 148–177.

Gittins, J.C., 1996. Quantitative methods in the planning of pharmaceutical research. Drug Information Journal 30, 479–487.

Gittins, J.C., 1997. Why crash pharmaceutical research? R&D Management 27 (1), 79–85.

Gittins, J.C., Jones, D.M., 1974. A dynamic allocation index for sequential design of experiments. In: Gani, J.M., Sarkadi, K., Vincze, I. (Eds.), Progress in Statistics, pp. 241–266.

Neyman, J., Pearson, E.S., 1933. On the probability of the most efficient tests of statistical hypotheses. Philosophical Transactions of the Royal Society of London Series A 231, 289–337.

Pearson, A., 1972. The use of ranking formulae in R&D projects. R&D Management 2, 69–73.

Regan, P.J., Senn, S., 1997. Project evaluation in drug development: A case study in collaboration. In: French, S., Smith, J.Q. (Eds.), The Practice of Bayesian Analysis. Arnold, pp. 259–277.

Senn, S.J., 1996. Some statistical issues in project prioritization in the pharmaceutical industry. Statistics in Medicine 15 (1), 2689–2702.

Senn, S.J., 1997. Statistical issues in drug development. Statistics in Practice. John Wiley & Sons.

Senn, S.J., 1998. Further statistical issues in project prioritization in the pharmaceutical industry. Drug Information Journal 32, 253–259.

Wald, A., 1945. Sequential tests of statistical hypothesis. Annals of Mathematical Statistics 16, 117–186.

Wald, A., Wolfowitz, J., 1948. Optimum character of the sequential probability ratio test. Annals of Mathematical Statistics 19, 326–339.

Zipfel, A., 2003. Modeling the probability-cost-profitability architecture of portfolio management in the pharmaceutical industry. Drug Information Journal 37, 185–205.