THE UNIVERSITY of York

CENTRE FOR HEALTH ECONOMICS

A Dynamic Programming Approach to the Efficient Design of Clinical Trials

K. Claxton
K.M. Thompson

CHE Technical Paper Series 13

A DYNAMIC PROGRAMMING APPROACH TO THE EFFICIENT DESIGN OF CLINICAL TRIALS

K. Claxton[a] and K.M. Thompson[b]

[a] Department of Economics, University of York, UK
[b] Harvard School of Public Health, Boston, MA.

Abstract

If the prospective evaluation of all feasible strategies of patient management is not possible or efficient then this poses a number of questions: i) which clinical decision problems will be worth evaluating through prospective clinical research; ii) if a clinical decision problem is worth evaluating which of the many competing alternatives should be considered "relevant" and be compared in the evaluation; iii) what is the optimal (technically efficient) scale of this prospective research; iv) what is an optimal allocation of trial entrants between the competing alternatives; and v) what is the value of this proposed research? The purpose of this paper is to present a Bayesian decision theoretic approach to the value of information which can provide answers to each of these questions. An analysis of the value of sample information was combined with dynamic programming and applied to numerical examples of sequential decision problems. The analysis demonstrates that this approach can be used to establish: optimal sample size; optimal sample allocation; and the societal payoff to proposed research. This approach provides a consistent way to identify which of the competing alternatives can be regarded as "relevant" and should be included in any evaluative study design. Bayesian decision theory can provide a general methodological framework which can ensure consistency in decision making between service provision, research and development and the design, conduct and interpretation of clinical research.

Key words: Bayesian decision theory, value of information, clinical trials, stochastic CEA

Correspondence: Karl Claxton, Department of Economics, University of York, Heslington, York. YO1 5DD. United Kingdom. Tel: (44) 01904 433798. Fax: (44) 01904 433759. E-mail: kpc1@york.ac.uk; claxton_mayer@msn.com.

1 Introduction

In many circumstances valid inferences cannot made about the expected costs and benefits of alternative strategies of patient management simply by observing current clinical practice for two reasons. First, many key parameters, particularly measures of efficacy, are vunerable to selection bias. Second, some feasible strategies of patient management are not part of current practice and have simply never been observed. There are potentially a large number of possible strategies in most sequential clinical decision problems. However the prospective evaluation of all these strategies is unlikely to be regarded as ethical and certainly not efficient (or even possible) given limited resources for research and development, and a recognition of the opportunity costs (health benefits forgone) for those enrolled in the trial and for the population of patients awaiting the results of the research.

In practice many feasible strategies are ruled out as irrelevant during the design of prospective research. However if the identification of "relevant alternatives" to be included in prospective research is either arbitrary, or uses an implicit decision rule inconsistent with those that will be applied when the study is complete, then there is a danger that research will revolve around an arbitrarily selected sub set of strategies. The optimal strategy may have already been ruled out of the analysis as an "irrelevant alternative". The evaluation of a particular clinical decision problems should, at least initially, consider all feasible alternatives rather than focus only on those currently used or those identified as of interest in some arbitrary way. The "relevant alternatives" which should be compared in prospective research should be identified explicitly and consistently.

If alternatives cannot be ruled out a priori but the prospective evaluation of all feasible strategies is not possible or efficient then this poses a number of questions: i) which clinical decision problems will be worth evaluating through prospective clinical research; ii) if a clinical decision problem is worth evaluating which of the many competing alternatives should be considered "relevant" and be compared in the evaluation; iii) what is the optimal (technically efficient) scale of this prospective research; iv) what is an optimal allocation of trial entrants between the competing alternatives; and v) what is the value of this proposed research? These are questions of how to establish technical efficiency in research design (including the selection of relevant alternatives as well as optimal scale) and how to achieve allocative efficiency in research and development.

The purpose of this paper is to present a Bayesian decision theoretic approach to the value of information which can provide answers to each of these questions. Our aim is to provide a general methodological framework which can ensure consistency in decision making between service provision, research and development and the design, conduct and interpretation clinical research.

2 Methodological background

A Bayesian decision theoretic approach to the value of sample information has been available for some time but has only recently been applied to simple decision problems (two alternatives) in the evaluation of health care technologies (Claxton and Posnett (1996); Hornberger and Eghtesady (1998); Claxton 1999). This approach is generalised to the analysis of sequential clinical decision problems which offer a choice between many strategies. Estimates of the expected net benefit of sample information are used in a dynamic program to establish optimal (efficient) allocation of trial entrants, optimal (efficient) sample size, technically efficient research design and the expected net benefits of the proposed research.

There is a body of literature which considers the optimal allocation of trial entrants in sequential clinical trials where the results of the trial accumulate over time and can be used to assign entrants to the different arms (e.g., Armitage (1985)). An example of this type of approach is Bather's "play the winner rule" where patients are assigned to the arm of the trial which appears to be most effective given the accumulated trial results (e.g., Bather (1981); Bather (1985)). This approach and others addressing the same problem (e.g., Zelen (1969); Hoel et al (1975); and Igiewicz (1983)) do not consider the marginal cost of sampling and tend to focus on minimising the potential health cost to individuals enrolled in the trial. They are primarily concerned with individual medical ethics rather than the collective ethical concerns for the costs of acquiring sample information and the future patients who will benefit from the information generated by the research (e.g., Upton and Lee (1976)). These approaches are also concerned with sequential clinical trials where the accumulated results from earlier participants in the trial are available and are used to allocate those entering the trial. Here we addresses a more fundamental problem of optimal allocation in a fixed sample design where sample information is only available at the end of the trial. The value of sample information, optimal sample size, and the allocation of patients at each stage (and the identification of relevant alternatives), must be established before any sample information is available.

The premise of our analysis is that clinical decision rules and research methods should be consistent with the explicit objective of the health care system. In this case we take an appropriate objective to be: maximise health gain subject to the budget constraint on service provision. This analysis is also is based on an explicit recognition that information is costly to acquire both in terms of resources and health outcomes and that the information generated by clinical research is valuable insofar as it reduces the cost of uncertainty surrounding clinical decisions that can not be deferred. This analysis provides a technically efficient approach to research design which is based on an explicit consideration of the expected benefits and costs of acquiring sample information.

The optimal allocation of trial entrants is based on the marginal benefits and marginal costs of allocating a trial entrant to a particular arm of the trial. This optimal sample allocation also enables "relevant alternatives" which should be compared in prospective research to be identified explicitly and consistently. Relevant alternatives can be defined as those where it will be efficient to allocate some of the sample in a subsequent evaluative trial. When it is not efficient to allocate patients to an arm of a trial then that alternative can be safely ruled out as irrelevant. It is only by explicitly considering the marginal costs and benefits of sample information that a consistent and rational definition of what constitute relevant and irrelevant alternatives in any evaluation is available.

The stylised numerical examples of single, two and four stage decision problems presented below (and detailed in appendix A) demonstrate that optimal patient allocation increases the value of sample information and can be used to identify and rule-out irrelevant alternatives which should not be included in an efficient trial design. Indeed if arbitrary fixed allocation rules are used then the research design will be technically inefficient (the expected net benefits of sampling will be underestimated and optimal sample size will be biased), there will be a danger that cost-effective research proposals may be rejected, and any subsequent evaluation may only consider an arbitrarily selected sub set of programs.

3 A single stage decision problem

A simple single stage decision problem which involves a choice between only two alternative strategies of patient management is illustrated in figure 1a where treatment t_0 can be regarded as current practice (in this case no treatment) and t_1 is a new treatment. Some (prior) information is available about the path probabilities and the associated costs and utilities of both treatments, where t_1 and t_0 have a prior expected utilities of $E(U|t_1)$ and $E(U|t_0)$, and prior expected costs of $E(C|t_1)$ and $E(C|t_0)$ respectively (this prior information is detailed in appendix A). A decision rule which is consistent with the objective of maximising health gain subject to the budget constraint on service provision would be to maximise net health benefits, or select t_1 if the prior incremental net benefit (δ_0) is positive:

$$\delta_{0(2)} = (E(U|t_1)\text{-}g.E(C|t_1)) - (E(U|t_0)\text{-}g.E(C|t_0)) > 0 \qquad (1)$$

where g is the minimum acceptable effectiveness cost ratio for the new treatment and 1/g can be regarded as the acceptable price per effectiveness unit (e.g., Weinstein and Zeckhauser (1972); Phelps and Mushlin (1991); Claxton and Posnett (1996); Stinnett and Mullahy (1998)). Clearly whenever a decision is made, or CEA is used to make a statement about what is cost-effective, a monetary valuation of health outcome is applied.

Figure 1a

The treatment decision based only on existing information (the prior mean incremental net benefit in (1) is uncertain because the prior incremental net benefit has a prior variance of ($\sigma_{(t)}^2$). The expected cost of this uncertainty will be determined by the probability that a decision based on prior mean net benefit is wrong and the size of the opportunity loss if the wrong decision is made. Additional information will be valuable insofar as it reduces this expected opportunity loss.

To decide whether acquiring additional information through prospective research will be efficient it is necessary to establish the marginal benefits and the marginal cost of sampling, given an optimal (random) allocation of trial entrants between the two arms of the proposed trial. It is convenient to separate the sample allocation decisions at stage 2 from the selection of optimal sample size at stage 1 in figure 1a, because it helps the explanation of the dynamic programming approach which is used to solve the two and four-stage decision problems in section 4.

The value of sample information

The benefit of sample information is the reduction in the cost of uncertainty surrounding the treatment decision. For an individual patient the expected benefit of a sample of $n_{(2)}$ entering stage 2 in figure 1a with n_{t1} allocated to t_1 and $n_{(2)} - n_{t1}$ allocated to t_0 can be expressed as the Expected Value of Sample Information at stage 2 ($EVSI_{(2)}|n_{(2)},n_{t1}$) (see Raiffa and Schlaifer (1959); Schlaifer (1961); Raiffa (1968); Claxton and Posnett (1996); Claxton (1999))

$$EVSI_{(2)}|n_{(2)},n_{t1} = \frac{1}{g}.\sqrt{V\,|\,n_{(2)},n_t 1}\,.\sigma_{0(2)}.L(D_{(2)}|n_{(2)},n_{t1}) \qquad (2a)$$

where:

$L(D_{(2)}\|n_{(2)},n_{t1})$	= unit normal loss integral for standardised distance $D_{(2)}\|n_{(2)},n_{t1}$
$D_{(2)}\|n_{(2)},n_{t1}$	$= (\delta_{0(2)}-\delta_b)/\sqrt{V\,\|\,n_{(2)},n_t 1}\,.\sigma_{0(2)}$
$V\,\|\,n_{(2)},n_t 1$	$= \sigma_{0(2)}^2/(\sigma_{0(2)}^2 + \sigma_{n(2)}^2)$
$\delta_{0(2)}$	$= (E(U\|t_1)-g.E(C\|t_1)) - (E(U\|t_0)-g.E(C\|t_0))$
δ_b	= point of indifference between t_1 and t_0 ($\delta_b = 0$)
$\sigma_{0(2)}^2$	= prior variance of $\delta_{0(2)}$
$\sigma_{n(2)}^2$	$= (\sigma_{t1}^2/n_{t1}) + (\sigma_{t0}^2/(n_{(2)}-n_{t1}))$
σ_{t1}^2 and σ_{t0}^2	= variance of the net benefit of t_1 and t_0 respectively
$\sigma_{t1}^2 \neq \sigma_{t0}^2$	

The benefit of sample information is the reduction in sample variance ($\sigma_{n(2)}^2$ in $V|n_{(2)},n_{t1}$ in 2a) which reduces the posterior variance of the difference in net benefit. The EVSI in 2a represents the benefit of sample information to an individual patient but this information is non rival and can be used to inform the treatment decisions for the population of current and future patients:

$$EVSI\,|\,n_{(2)},n_{t1}.\sum_{t=1}^{T}\frac{(I_t - n_{(2)})}{(1+r)^t} \qquad (2b)$$

where the incidence of new patients entering the decision problem (I) at time t is discounted at rate r over the effective lifetime (T) of the new technology. The population of current and future patients is also the population from which the sample will be drawn. In this sense patients entering the trial ($n_{(2)}$) are "used up" and will not benefit from the information generated by the research [1].

The marginal cost of additional sample information will include the additional treatment cost compared to current practice plus any marginal reporting costs (C_r). The cost of a sample (Cs) of $n_{(2)}$ at stage 2 with n_{t1} allocated to t_1 and $n_{(2)}-n_{t1}$ allocated to t_0 is given by:

$$Cs_{(2)}|n_{(2)},n_{t1} = (E(C|t_1) - E(C|t_0)).n_{t1} + C_r.n_{(2)} \qquad (2c)$$

where $(E(C|t_1) - E(C|t_0))$ is the additional cost of the new treatment. The additional treatment cost for those allocated to t_0 are zero since trial entrants would have incurred the costs of current practice anyway [2]. The difference between the expected benefit of sample information (2b) and the cost of the sample (2c) is the Expected Net Benefit of Sample Information and represents the societal payoff from proposed research:

$$ENBS_{(2)}|n_{(2)},n_{t1} = EVSI_{(2)}|n_{(2)},n_{t1} - Cs_{(2)}|n_{(2)},n_{t1} \qquad (2d)$$

Optimal allocation of trial entrants

The payoff at stage 2 of a sample of $n_{(2)}$ with n_{t1} allocated to t_1 ($\Pi_{(2)}|n_{(2)},n_{t1}$) is the expected net benefit of sample information:

$$\Pi_{(2)}|n_{(2)},n_{t1} = ENBS_{(2)}|n_{(2)},\,n_{t1} \qquad (2e)$$

Table 1 Optimal Allocation of the Sample Entering Stage 2 (1/g=£4,000)

n_{t1} / $n_{(2)}$		Payoff at stage 2 $\Pi_{(2)}\|n_{(2)},n_{t1} = ENBS_{(2)}\|n_{(2)},n_{t1}$ Allocation to t_1 (n_{t1}) (with n_{t1} allocated to t_1, and $n_{(2)}-n_{t1}$ allocated to t_0)														Maximum payoff and optimal allocation (n_{t1}^*)	
		0	1	2	3	4	5	6	7	8	9	10	11	12	..	$\Pi_{(2)}\|n_{(2)},n_{t1}^*$	n_{t1}^*
Sample Entering Stage 2 $n_{(2)}= n_{(1)}$	0	0	-	-	-	-	-	-	-	-	-	-	-	-	-	0	0
	1	0	-£10,400	-	-	-	-	-	-	-	-	-	-	-	-	0	0
	2	0	-£9,952	-£20,800	-	-	-	-	-	-	-	-	-	-	-	0	0
	3	0	-£8,395	-£18,405	-£31,200	-	-	-	-	-	-	-	-	-	-	0	0
	4	0	-£7,031	-£8,728	-£26,905	-£41,600	-	-	-	-	-	-	-	-	-	0	0
	5	0	-£6.015	£712	-£8,166	-£35,812	-£52,000	-	-	-	-	-	-	-	-	£712	2
	6	0	-£5,257	£8,216	£11,289	-£9,355	-£45,065	-£62,400	-	-	-	-	-	-	-	£11,289	3
	7	0	-£4,677	£14,053	£27,397	£19,266	-£12,344	-£54,572	-£72,800	-	-	-	-	-	-	£27,397	3
	8	0	-£4,224	£18,645	£40,301	£43,646	£24,092	-£16,786	-£64,259	-£83,200	-	-	-	-	-	£43,646	4
	9	0	-£3,860	£22,314	£50,664	£63,540	£55,766	£26,223	-£22,300	-74,084	-£93,600	-	-	-	-	£63,540	4
	10	0	-£3,562	£25,304	£59,072	£79,752	£81,979	£64,182	£26,199	-£28,711	-£84,010	-£104,000	-	-	-	£81,979	5
	11	0	-£3,314..	£27,795..	£66,028..	£93,093..	£103,658..	£96,005..	£69,558..	£24,479..	-35,753£	-£94,013..	-£114,400..	-	-	£103,658	5
	12	0	-£3,108	£29,906	£71,840	£104,237	£121,602	£122,536	£106,255	£72,503	£21,399	-£43,304	-£104,078	-£124,800	-	£122,536	6
	..	0	..	..	..	..	..	..	..	..	..	..	..	..	..	..	..

In order to establish the optimal allocation of trial entrants an estimate of the payoff at stage 2 for every feasible allocation of $n_{(2)}$ to t_1 and t_0 is required. The optimal allocation of each sample size considered ($n_{(2)}$ = 1, ...,500) to t_1 (n_{t1}^*) will be where the marginal net benefit of allocating an entrant to t_0 is equal to t_1, or where $ENBS_{(2)}|n_{(2)}$, n_{t1} reaches a maximum for each $n_{(2)}$. This optimal contingent allocation is illustrated in table 1 where the sample sizes considered are represented by each row, the feasible allocation to t_1 are represented by the columns and the payoffs ($\Pi_{(2)}|n_{(2)},n_{t1}$) are in the body of the table. The optimal allocation of each sample is then the row maximum. The optimal contingent allocations (n_{t1}^*) and associated payoffs ($\Pi_{(2)}|n_{(2)},n_{t1}^*$) are reported in the far right columns.

The optimal allocation of trial entrants is also illustrated in figure 1b. In this example a greater proportion of the trial entrants are allocated to t_0 than with an equal allocation rule (represented by the rising diagonal) because the marginal costs of assigning an entrant to t_1 is higher than t_0 where the additional treatment costs are zero. This difference in marginal sampling cost offsets the higher marginal benefits of assigning entrants to t_1 ($\sigma_{t1}^2 > \sigma_{t0}^2$ when 1/g<£20,000) 3. The proportion of the sample allocated to t_0 also falls as total sample size increases because the marginal benefits of sampling fall with increased sample size but the differences in the marginal cost of sampling are constant.

Optimal allocation is also dependent on the monetary valuation of health outcome for two reasons: first because this determines the "weight" placed on the differences in the benefits and costs of assigning a trial entrant to each alternative; and second the value of g partly determines the variance of the expected costs and therefore the marginal benefits of assignment to t_1 and t_0. In this example when the value of 1/g is increased the difference between σ_{t1}^2 and σ_{t0}^2 falls (because the variance of the expected cost is a larger component of σ_{t1}^2), consequently the effect on the marginal benefits and on the weight attached to differences in marginal costs both work in the same direction, and a greater proportion of the sample is assigned to t_1.

Figure 1b

Once the payoff for each sample size (given that it will be allocated optimally) is established at stage 2 the optimal sample size at stage 1 ($n_{(1)}$*) will be where the payoff ($\Pi_{(1)}|n_{(1)}$) reaches a maximum:

$$\Pi_{(1)}|n_{(1)} = \Pi_{(2)}|n_{(2)},n_{t1}^* = ENBS_{(2)}|n_{(2)},\ n_{t1}^*,\ n_{(1)} = n_{(2)} \qquad (2f)$$

The optimal sample size selected at stage 1 will be allocated according to the contingent allocation decisions that have already been made at stage 2 and illustrated in figure 1b. If the maximum payoff ($\Pi_{(1)}|n_{(1)}$*) is positive and greater than the fixed costs of the research (C_f) then it will be efficient to acquire sample information to inform this clinical decision problem at this technically efficient scale of research.

Figure 1c

Optimal sample size and ENBS for this example are illustrated in figure 1c. The ENBS reaches an maximum of £725,000 at a sample size of 116 when optimal allocation is used (n_{t1}*=29). However if an equal allocation rule was used the ENBS would be lower reaching a maximum of £627,000 at a sample size of 76 (n_{t1}=38). Clearly the optimal allocation of trial entrants increases the value of proposed research. If arbitrary allocation rules are used then sample size (and sample allocation) will be inefficient, the value of proposed research will be underestimated and there will be a danger that cost-effective research will be rejected.

This analysis has provided answers to each of the questions posed at the start of the paper. It should be recognised that in this simple choice between only two strategies we can say that t_1 is a relevant alternative if it is efficient to allocate trial entrants to that arm of the trial (n_{t1}*>0) and if $\Pi_{(1)}|n_{(1)}$*>C_f. If either of these conditions do not hold then the proposed research is not efficient and t_1 can be regarded as an irrelevant alternative. In these circumstances no additional research is necessary and the choice between treatment strategies can be safely based on existing information (prior incremental net benefit).

4 Generalizing to sequential clinical decision problems

The above analysis has assumed that only two strategies of patient management are possible. However in almost all clinical decision problems there are potentially very large numbers of strategies. In practice many feasible alternatives are ruled out a priori as "irrelevant" and the alternatives regarded as "relevant" are selected from those currently used or identified as of interest in some implicit and arbitrary way. These alternatives represent a very small and arbitrarily selected sub set of strategies. The optimal strategy may have long since been ruled out as irrelevant or simply never seriously considered. Without a rational and consistent way to identify relevant alternatives which should be included in any evaluation there is a danger that the research agenda will simply revolve around arbitrarily selected alternatives.

The evaluation of a clinical decision problem should, at least initially, consider all feasible alternatives rather than focus only on those currently used or those identified as of interest. The analysis of the simple choice between two alternatives in

section 3 can be generalised to a S stage sequential clinical decision problem where a clinician faces a choice between two alternatives at each stage. This provides S+1 possible strategies and those that are relevant can be defined as those where it will be efficient to allocate some of the sample in a proposed trial. Where it is not efficient to allocate patients to an arm of a trial then that alternative can be safely ruled out as irrelevant. The analysis of a two stage and four stage decision problem in section 4.1 and 4.2 demonstrates that it is only by explicitly considering the marginal costs and benefits of sample information that a consistent and rational definition of what constitute relevant and irrelevant alternatives is available.

4.1 A two stage decision problem

A two stage clinical decision problem is illustrated in figure 2a. This decision problem is identical to the simple single stage choice between two treatments in figure 1a except that a diagnostic test has been made available (see table 1, appendix A). This provides a choice between three possible strategies of patient management: t_0; t_1; or test (t_e) and treat with t_1 following a positive results (t_e^+) and t_0 following a negative result (t_e^-).

Figure 2a

This two stage decision problem can be solved as a three stage dynamic programme when optimal sample allocation is required. At stage 3 an optimal allocation of trial entrants between t_1 and t_0 contingent on each sample entering stage 3 is made based on the expected net benefit of every feasible allocation of each sample. At stage 2 an optimal allocation between the test and no test arm of the proposed trial is made contingent on each sample entering stage 2. This allocation is based on the expected net benefit of sampling at stage 2 and the payoff at stage 3 from the sample assigned to the no test arm. Finally the optimal sample size which will maximise expected net benefits can be selected at stage 1. Solving contingent allocation decisions before the optimal sample size is selected at stage 1 is consistent with the principle of backward induction and can be characterised as a simple three-stage dynamic programme, where the payoff for each sample size considered at stage 1 is the expected net benefit given that an optimal patient allocation policy will be followed from stage 1 to the end.

Optimal allocation at stage 3

The choice between t_0 and t_1 at stage 3 is identical to the single stage choice considered in section 2. Both the payoffs and the optimal allocation of a sample entering stage 3 ($n_{(3)}$) is identical to the single stage decision problem considered in section 3 (see figure 1b). The payoff of a sample of $n_{(3)}$ allocated optimally, with and $n_{(3)} - n_{t1}^*$ allocated to t_0, and n_{t1}^* allocated to t_0, will be:

$$\Pi_{(3)}|n_{(3)},n_{t1}^* = ENBS_{(3)}|n_{(3)},n_{t1}^*. \qquad 3a$$

Optimal allocation at stage 2

A sample entering stage 2 ($n_{(2)}$) which is allocated between test (n_{te}) and no test alternatives ($n_{(2)}$-n_{te}) will generate benefits and costs of sampling at stage 2 ($ENBS_{(2)}|n_{(2)},n_{te}$), but the sample allocated to the no test arm of the trial it will enter stage 3 ($n_{(3)}=n_{(2)}-n_{te}$) and will be allocated optimally, generating payoffs of $\Pi_{(3)}|n_{(3)},n_{t1}^*$ in 3a. This is a recursive relationship (see Bellman (1957); Bellman and Dreyfus (1962)) where the payoff at stage 2 is partly determined by the payoffs associated with the optimal allocation policy which will be followed at a later stage.

$$\Pi_{(2)}|n_{(2)},n_{te} = ENBS_{(2)}|n_{(2)},n_{te} + \Pi_{(3)}|n_{(3)},n_{t1}^*, \qquad n_{(3)} = n_{(2)}-n_{te} \qquad 3b$$

The payoffs at stage 2 require estimates of the $ENBS_{(2)}|n_{(2)},n_{te}$ for every feasible allocation of $n_{(2)}$ to the test and no test arms of the trial. The expected benefits of a sample of $n_{(2)}$ with n_{te} allocated to the testing and $n_{(2)}-n_{te}$ allocated to the no test arm can be expressed as the expected value of sample information:

$$EVSI_{(2)}|n_{(2)},n_{te} = \frac{1}{g}.\sqrt{V\,|\,n_{(2)},n_{te}}\,.\sigma_{0(2)}\,.L(D_{(2)}|n_{(2)},n_{te}) \qquad 3c$$

where:

$$D_{(2)}|n_{(2)},n_{te} = (\delta_{0(2)}-\delta_b)/\sqrt{Vn_{(2)},n_{te}}$$

$$\delta_{0(2)} = (E(U|t_e)-g.E(C|t_e)) - (E(U|nt_e)-g.E(C|nt_e))$$

$$\sqrt{Vn_{(2)},n_{te}} = \sigma_{0(2)}^2/(\sigma_{0(2)}^2 + \sigma_{n(2)}^2)$$

$$\sigma_{n(2)}^2 = (\sigma_{te}^2/n_{te}) + (\sigma_{nte}^2/(n_{(2)}-n_{te}))$$

Although 3c appears to be similar to the expression for EVSI in 2a, the prior incremental net benefit ($\delta_{0(2)}$), the prior variance of $\delta_{0(2)}$ ($\sigma_{0(2)}^2$), and the variance of the net benefits of not testing (σ_{nte}^2) will depend on the sample allocated to the no test arm which will be allocated between t_1 and t_0 at stage 3.

The prior net benefits of not testing ($E(U|nt_e)-g.E(C|nt_e)$) and therefore $\delta_{0(2)}$ is dependent on the sample allocated to the no test alternative because there is a probability ($p(\delta_{x(3)}>\delta_{x(3)}^*)$) that this sample of $n_{(2)}-n_{te} = n_{(3)}$ will generate a posterior mean

net benefit at stage 3 ($\delta_{1(3)}$) which will lead to t_1 being selected with net benefits of $E(U|t_1)$-g.$E(C|t_1)$. There is also a probability ($1-p(\delta_{x(3)}>\delta_{x(3)}^*)$) that the same sample will generate a posterior mean net benefit which will lead to t_0 being selected with net benefits of $E(U|t_0)$-g.$E(C|t_0)$.

$$E(U|nt_e)\text{-g.}E(C|nt_e) = p(\delta_{x(3)}>\delta_{x(3)}^*).(E(U|t_1)\text{-g.}E(C|t_1)) + 1-p(\delta_{x(3)}>\delta_{x(3)}^*).(E(U|t_0)\text{-g.}E(C|t_0)) \quad 3d$$

where

$\delta_{1(3)} = (I_{0(3)}.\delta_{0(3)} + I_{x(3)}.\delta_{x3s})/(I_{0(3)} + I_{x(3)})$

$\delta_{x(3)}^* = ((I_{0(3)} + I_{x(3)})\delta_b - I_{0(3)}.\delta_{0(3)})/I_{x(3)}$ = critical value of sample mean [4]

$I_{0(3)} = 1/\sigma_{0(3)}^2$,

$I_{x(3)} = 1/\sigma_{n(3)}^2$

Since the choice of sample size must be made before any sample information is available we assume that the sample mean is normally distributed, centred on the prior mean, with sample variance of $\sigma_{n(2)}^2$ [5].

$$p(\delta_{x(2)}>\delta_{x(2)}^*) = p(Z>((\delta_{0(2)}>\delta_{x(2)}^*)/\sigma_{n(2)})) \quad 3e$$

The prior variance of $(E(U|nt_e)\text{-g.}E(C|nt_e))$ and therefore $\sigma_{0(2)}^2$ are also dependent on the sample allocated to the no test arm. This is for two reasons: first, it determines the value of $p(\delta_{x(3)}>\delta_{x(3)}^*)$ and second, this sample will reduce the uncertainty surrounding the net benefit of t_1 and t_0 at stage 3. Therefore the prior variance of $(E(U|nt_e)\text{-g.}E(C|nt_e))$ at stage 2 is a combination of the <u>posterior</u> variance [6] of $E(U|t_1)$-g.$E(C|t_1)$ and $E(U|t_0)$-g.$E(C|t_0)$ and is calculated for each sample assigned to the no test arm assuming that an optimal allocation policy will be followed at stage 3. The population variance of the net benefit of not testing (σ_{nte}^2) is also partly determined by the sample entering and allocated at stage 3, because this determines the value of $p(\delta_{x(3)}>\delta_{x(3)}^*)$.

The $EVSI_{(2)}|n_{(2)},n_{te}$. can now be established taking into account the relationship between the sample considered at stage 2 and sample information it will generate at stage 3. It measures the <u>additional</u> value of sample information given that a sample of $n_{(2)}$ will generate a sample of $n_{(2)}$-n_{te} at stage 3, which will reduce the uncertainty surrounding the no test alternative and will also change the expected prior net benefits of choosing not to test. So unlike the single-stage decision problem, the prior mean and prior variance for the initial decision is not independent of the sample size considered.

The cost of sampling at stage 2 ($Cs_{(2)}|n_{(2)},n_{te}$) includes the additional treatment cost, compared to current practice, and any marginal reporting costs. The marginal cost of assigning an entrant to the no test arm of the trial will be zero because the additional treatment and reporting cost have already been taken into account in the estimates of the ENBS at stage 3. The marginal cost of assigning an entrant to the testing arm of the trial will include the additional cost of the testing strategy and marginal reporting costs.

$$Cs_{(2)}|n_{(2)},n_{te} = ((E(C|t_e)\text{-}E(C|t_0))+C_r).n_{te} \quad 3f$$

The $ENBS_{(2)}|n_{(2)},n_{te}$ is the difference between the expected benefit and the expected cost of a sample of $n_{(2)}$ with n_{te} allocated to the test arm of the trial and $n_{(2)}$-n_{te} allocated to the no test arm.

$$ENBS_{(2)}|n_{(2)},n_{te} = EVSI_{(2)}|n_{(2)}, n_{te} - Cs_{(s)}|n_{(2)}, n_{te} \quad 3g$$

The payoff (see 3b) for every feasible allocation to the test and no test arms of each sample considered at stage 2 can be established assuming that an optimal allocation policy will be followed at stage 3 (see 3a). The optimal allocation will be where the marginal payoff of assigning a trial entrant to test arm is equal to the marginal payoff of assigning the entrant to the no test arm, or where the $\Pi_{(2)}|n_{(2)},n_{te}$ reaches a maximum for each $n_{(2)}$. The optimal sample allocation contingent on $n_{(2)}$ for this numerical example is illustrated in figure 2b for three values of 1/g [7].

Figure 2b

Optimal allocation depends on both the sample size entering stage 2 and the monetary valuation of health outcome. In this example the variance of the test arm is greater than the no test arm so the marginal benefits of allocating an entrant to t_e will be greater than nt_e. However the marginal costs of assigning an entrant to t_e is greater than nt_e. When the value of 1/g=£4,000 n_{te}=0 and it is not efficient to allocate any sample to the testing arm of the proposed trial. In these circumstances the testing strategy can be safely ruled out as an irrelevant alternative. The proposed trial should simply compare t_1 and t_0 with the optimal allocation and payoffs given in section 3. However if the value of 1/g is increased to £10,000 then less weight is placed on the additional costs of assigning entrants to the testing arm of the trial and more weight is placed on the additional benefits. In these circumstances it is efficient to assign entrants to t_e and the test strategy should be regarded as a relevant alternative.

Optimal sample size and ENBS at stage 1

The first stage is simply to select the optimal sample size given that it will be allocated optimally between the test and no test arms, and between t_1 and t_0 at stage 3. The payoff from a sample of $n_{(1)}$ at stage 1 ($\Pi_{(1)}|n_{(1)}$) is simply the payoff given that an optimal sample allocation policy is followed at each subsequent stage.

$$\Pi_{(1)}|n_{(1)} = \Pi_{(2)}|n_{(2)},n_{te}^{*} = ENBS_{(2)}|n_{(2)},n_{te}^{*} + ENBS_{(3)}|n_{(3)},n_{t1}^{*} \qquad 3h$$

$n_{(1)} = n_{(2)}$, $n_{(3)} = n_{(2)}-n_{te}^{*}$

The optimal sample size at stage 1 ($n_{(1)}^{*}$) will be where $\Pi_{(1)}|n_{(1)}$ reaches a maximum. The optimal sample sizes and ENBS using optimal and fixed sample allocation rules are illustrated in figure 2c

Figure 2c

The payoff reaches a maximum of £724,970 at a sample size of 116 when optimal allocation is used (n_{te}=0, n_{t1}=29). In this example all the sample is allocated to the no test arm, the testing strategy is ruled out as irrelevant, so the problem is simplified to the single stage decision problem considered in section 3. If an equal allocation rule is used the maximum payoff would be £362,366 at a sample size of 80. The equal allocation rule forces half the sample to be allocated to testing where sampling costs exceed expected benefits. This demonstrates that optimal allocation increases the value of proposed research and allows irrelevant alternatives to be ruled-out explicitly and consistently because it is possible to allocate none of the sample to a particular arm of the trial (in this case testing). If arbitrary allocation rules are used then: irrelevant alternative may be included in a proposed trial design; sample size and allocation between alternatives will be inefficient; the value of proposed research will be underestimated; and there will be a danger that efficient research proposals will be rejected.

4.2 A four stage decision problem

The dynamic programming approach to optimal patient allocation can also be applied to more complex sequential clinical decision problems. In this section the approach is applied to a four stage decision problem which is illustrated in figure 3a. This decision problem is identical to the two stage test/treatment decision presented in section 4.1 except that another treatment (t_2) which is less costly but less effective than t_1 is available (see appendix A). The clinician now faces a choice between 5 strategies of patient management: t_0; treat (t_r) with either t_1 or t_2; test and treat with t_1; or test and treat with t_2.

Figure 3a

This four stage decision problem can be characterised as a five stage dynamic programme where optimal contingent sample allocations between t_1 and t_2 are made at stage 5 and stage 4 based on the expected net benefits of sampling for every feasible allocation of each sample entering the stage. At stage 3 contingent allocation decisions must be made between treatment (with either t_1 or t_2) and no treatment. The optimal allocation is based on the expected net benefits of sampling at stage 3 and the payoff from stage 4 given the optimal allocation of the sample allocated to the treatment arm of the trial. At stage 2 an optimal allocation between the test and no test arms of the trial is made for each sample entering stage 2. This contingent allocation is based on the expected net benefits of sampling at stage 2, the payoff from the optimal allocation of the sample which will enter stage 3, and the payoff from the optimal allocation of the sample which will enter stage 5. Finally the optimal sample size can be selected at stage 1 given that an optimal allocation policy will be followed at each subsequent stage.

This dynamic programming approach exploits the recursive relationship between payoffs at: stage 3 and 4; stage 2, 3 and 5; and stage 1 and 2. The ENBS at earlier stages takes into account the relationship between the sample considered and the information it will generate when it is allocated at a subsequent stage. This ensures that the estimates of the benefits of sample information at earlier stages measure only the additional value of information given that the sample will be allocated at later stages: reducing the uncertainty surrounding those contingent decisions and changing the prior net benefit of those alternatives. So, in the same way as allocation of the sample at stage 2 in section 4.1, the prior mean, prior variance and population variance of the alternatives at earlier stages are not independent of the sample size considered and the way it is allocated. The details of measures of payoff, optimal contingent allocations at each stage, and optimal sample size are given in Appendix B.

Optimal allocation

The ENBS at each stage, given the optimal contingent allocation of the sample entering stage 1, is illustrated in figure 3b. The $ENBS_{(s)}$ represents the payoff at each stage and the total payoff for a sample selected at stage 1 ($ENBS_{(1)}$) is the vertical summation of ENBS for each sample considered. Figure 3b illustrates that the sample allocation at an earlier stages is partly determined by the payoffs generated by the allocated sample entering subsequent stages. For example, the payoff at stage 2 is negative ($ENBS_{(2)} <0$) but it is still worth allocating patients at stage 2 because this will generate positive payoffs at stage 3 and 4 ($ENBS_{(3)} >0$, $ENBS_{(4)} >0$)which more than offset the negative $ENBS_{(2)}$.

Optimal sample size and ENBS
Optimal sample size and ENBS using optimal allocation can be compared to the results of a fixed and equal allocation rule in figure 3c. ENBS reaches an maximum of £3,866,000 at a sample size of 165 when optimal allocation is used (n_{te}*=67, n_{tr}*=27, n_{t1}*=12 at stage 4, and n_{t1}*=19 at stage 5). However ENBS and optimal sample size is lower when an equal allocation rule is used (£3,530,000, $n_{(1)}$*=148).

Figure 3b

In this example at $n_{(1)}$* it is efficient to assign some of the trail entrants to all the potential arms of the proposed trial including testing (all are relevant alternatives see figure 3b). In section 4.1 testing was not a relevant alternative but once the alternative treatment t_2 becomes available it does become relevant because treatment with t_2 following a positive test result makes the testing strategy more cost-effective. This demonstrates that ruling out alternatives a priori will be inefficient and will seriously bias the design of the proposed research. In this case if t_2 is ruled out (possibly because it is less effective than t_1) then the testing strategy would also not be included in the design. The ENBS would be seriously underestimated at £725,000 rather than £3,865,000 and there would be a danger that cost-effective research will be rejected.

Figure 3c

The dynamic programming approach
This four-stage decision problem is solved using a simple five-stage dynamic programme where contingent allocation decisions at each stage are solved before the optimal sample at stage 1 can be selected. The payoff for each sample considered at stage 1 is established conditional on an optimal allocation policy at each subsequent stage. This approach reduces the computation required to solve this problem even more dramatically than in the two-stage decision problem considered in section 4.1 The solution to this five stage problem, considering a maximum sample size of N, requires $(N+1)^2+(N+1))/2$ estimates of $\Pi_{(5)}|n_{(5)},n_{t1}$ at stage 5, 4, 3, and 2.

This reduces the number of estimates of the ENBS from $A^S = 125{,}751^4$ if full enumeration is required to A*S = 125,751*4 when a maximum sample of 500 is considered at each stage (A are the feasible allocations of each sample at each stage and S are the number of stages in the decision problem). The solution for this numerical example requires 503,004 estimates of the ENBS, which takes approximately 14 hours of computing time (10 estimates per second). This compares very favourably to the full enumeration of all possible alternatives which would require over 7.9 billion years of computing time, a task so enormous it can currently be regarded as impractical.

Although a dynamic programing approach was adopted in the two and four stage decision problems, this was simply because the dimensionality of the problem made the full enumeration, of all feasible allocations at every stage for every sample size, impossible. In many ways a the full enumeration of the problem would be more transparent (although less efficient) but this was only practicable for the single stage decision problem in section 3. However if developments in computing speed do make the scale of full enumeration possible the solution will still require the general framework of Bayesian decision theory to provide estimates of the marginal benefits and costs of sampling so that optimal sample size, optimal allocation of trail entrants, and the expected value of proposed research can be established.

5 Discussion

The maximum payoffs (ENBS at optimal sample size) for a range of possible monetary valuations of health outcome are illustrated in figure 4 for the single, two and four stage decision problems. These results demonstrate that the value of proposed research (ENBS) is determined by the monetary valuation of health outcome. In figure 4 there is clearly a strong relationship between ENBS and 1/g although this is not necessarily positive because 1/g determines both the prior mean (probability of error) and the value placed on consequences of error (slope of the loss function). For example, the ENBS for the single stage decision problem falls with 1/g (when 1/g>£6,500 and δ_0>0) because reduction in the probability of error offsets the increase in the value of errors if they occur.

Figure 4

What is clear from these example is that the value of proposed research cannot be established without making some assessment of societies monetary valuation of health outcome. This valuation cannot be avoided and can either be made explicitly, as in this case, or implicitly whenever decisions about research and development are actually made. These results also demonstrate that research design including: optimal sample size, the allocation of the sample and which alternatives should be regarded as relevant are also determined by the monetary valuation of health outcome. The optimal allocation of patients in a clinical trail is a fundamental design issue that must be addressed even in the fixed sample designs discussed above. This can only be done by considering the marginal benefits and marginal costs of alternative allocations. Therefore research cannot be efficiently designed independently of economic decision rules and the budget constraint on service provision.

A comparison of ENBS for the three decision problems demonstrates that ruling out possible strategies of patient management a priori (by excluding t_2 and simplifying the four stage to a two stage problem or ruling out testing and simplifying the two to a single stage problem) can seriously underestimate the value of proposed research. There will be a danger that cost-effective research will be rejected and research will focus on an arbitrary selected subset of strategies. This approach to the value of information does provide a way to safely rule out alternatives as irrelevant when it is not efficient to allocate any sample to that arm of the trial. This occurs in the two stage decision problem where it will not be efficient to allocate patients to the testing arm (when $1/g<£4,000$) and it can be ruled out as irrelevant. In other circumstances none of the alternatives are relevant and it is not efficient to conduct any prospective research. In these circumstances it will be efficient to base decisions on existing (prior) information. In these examples prospective research will not be efficient when $1/g\leq£3,000$ if t_2 is excluded and when $1/g\leq£2,000$ if t_2 is included.

It is also worth noting that the notion that dominated programmes (whether strongly or extendedly dominated) can be excluded from an analysis and from subsequent evaluations is fundamentally flawed. It is possible that a dominated alternative in a deterministic cost-effectiveness analysis will be relevant in subsequent evaluation while non dominated but not optimal alterative may not. It is only by explicitly considering the marginal costs and benefits of sample information that a consistent and rational definition of what constitute relevant and irrelevant alternatives is available. In these circumstances the concept of dominance is not only unnecessary (when using net benefit) it is also positively dangerous

6 Conclusions

Bayesian decision theory does provide a general methodological framework which can ensure consistency in decision making between health care provision, research and development priorities and the design, conduct and interpretation clinical research. The examples presented above demonstrate that an analysis of the value of sample information can be applied to more complex sequential clinical decision problems and provide answers to the questions posed in the introduction: i) which clinical decision problems will be worth evaluating through prospective clinical research; ii) if a clinical decision problem is worth evaluating which of the many competing alternatives should be considered "relevant" and be compared in the evaluation; iii) what is the optimal scale of this prospective research; iv) what is an optimal allocation of trial entrants between the competing alternatives; and v) what is the value of this proposed research to society?

These are questions of technical efficiency in research design, and allocative efficiency in research and development activities. Technically efficient research design requires the analyst to: explicitly consider the marginal costs and benefits of acquiring sample information; abandon the traditional and arbitrary power calculation and allocation rules; and establish the technically efficient scale of research by selecting the design that maximises the societal payoff to proposed research. Importantly it provides, possibly for the first time, a consistent, and rational definition of what constitute relevant alternatives in the economic and clinical evaluation of health care technologies. Allocative efficiency in research and development is possible once the societal payoff to proposed research has been established. The valuation of proposed research can be used to allocate resources between proposed research in the same clinical area; across broad clinical areas; and between research and development activities in general and the provision of health services. In short it can be used to efficiently allocate resources between technically efficient research proposals, and between a technically and allocatively efficient research and development programme and the provision of health care in a way which is consistent with the explicit objectives of the health care system.

Notes

1 If patients can have the disease and can be treated more than once then population of interest is not the population of patients at risk in t but expected episodes of illness in t.

2 Clearly it is possible that a new treatment has a lower cost then current practice if this difference exceeds marginal reporting costs then Cs<0 for the experimental arm of the trial. However this will not lead to infinite sample sizes because there is another substantial opportunity cost of sampling ie patients are "used up" in the trial, see (2b) Also we assume that the follow-up period of the trial corresponds to the time horizon of cost and health effects. With minor changes it is possible use this framework to consider different lengths of follow up, selecting the design with the highest societal payoff.

3 The marginal benefit of allocating a sample to the experimental arm will be determined by its contribution to the reduction in sample variance of incremental net benefit. Clearly the reduction in sample variance will be determined by the population variance of the net benefits for that arm and the size of the sample already allocated (see Vln(2),nt1 in 2a). So any differences in the marginal benefit of allocating to control or experimental are determined by differences in population variance not prior variance. Differences in the prior variance of the net benefits between control and experimental are only important insofar as they contribute to the prior variance of the incremental net benefit.

4 The critical value of the sample mean ($\delta_{x(3)}^{*}$) is the sample mean which generates a posterior mean that changes the prior decision. If the sample mean is greater than this critical value ($\delta_{x(3)}>\delta_{x(3)}^{*}$) the posterior mean will be greater than δ_b and the clinician should select t_1 at stage 3, but when $\delta_{x(3)}<\delta_{x(3)}^{*}$ the posterior mean is less than δ_b and the clinical should select t_0.

5 For the purposes of demonstrating this type of approach we have taken a parametric approach and assumed normality. This provides conjugate priors and eases exposition. However this assumption can be relaxed and non parametric numerical methods are available (Markov chain Monte Carlo) with sofware for their implementation (eg., Win BUGS)

6 Posterior variance is a combination of prior and sample variance and will be less than either the prior variance or the sample variance. In general the posterior variance ($\sigma_{1(s)}^2$) with a sample of $n_{(s)}$ is a combination of prior ($\sigma_{0(s)}^2$) and sample variance ($\sigma_{n(s)}^2$):
$\sigma_{1(s)}^2 = (\sigma_{0(s)}^2/(\sigma_{0(s)}^2 + (\sigma_{n(s)}^2))).(\sigma_{n(s)}^2)$

7 In the single stage decision problem there are no discontinuities in optimal sample allocation in figure 1b except for the steps which simply reflect the fact that trial entrants are discrete. However this allocation results in a smooth relationship between sample size and ENBS in figure 1c as expected. The marked discontinuities in the optimal allocation at stage 2 which are illustrated in figure 2b occur for three reasons: I) the discrete nature of trial entrants is more important at earlier stages (ie to allocate one more patient at a later stage may require a number of entrants at an earlier stage); ii) optimal allocation at earlier stages is not only determined by the value of sampling at that stage (which would generate a smooth allocation) but also by the payoff gained a subsequent stages; iii) as sample size changes some strategies become relevant. For example, at very small sample sizes it is not efficient to allocate sample to t_1 in figure 1b and 2c. This also occurs in the four stage problem where it is not efficient to allocate sample to the testing strategy when $n_{(1)}<48$ in figure 3b. These three issues interact with the discrete nature of trial entrants to produce the jaggedness in optimal allocation. Finally the discontinuities in allocation do generate a smooth relationship between ENBS at stage 1 and sample size in figure 2c and figure 3b.

Bibliography

Armitage, P. The search for optimality in clinical trials. *International Statistical Review* 1985; **53**:15-24.
Bather, JA. Randomised allocation of treatments in sequential experiments. *Journal of the Royal Statistical Society* B 1981; 43: 265-292.
Bather, JA. On the allocation of treatment in sequential medical trials. *International Statistical Review* 1985; 53: 1-13.
Berry DA, Stangl DK. *Bayesian Biostatistics*. New York: Marcel Dekker, 1996.
Bellman RE. *Dynamic Programming*. Princeton University Press, 1957.
Bellman, RE. and Dreyfus, SE. *Applied Dynamic Programming*. Princeton University Press, 1962.
Claxton K. The irrelevance of inference: a decision making approach to the stochastic evaluation of health care technologies. *Journal of Health Economics*1999; 8: 34-364.
Claxton K. Bayesian approaches to the value of information: implications for the regulation of new health care technologies. *Health Economics*1999; 8: 268-274.
Claxton K. and Posnett J. An economic approach to clinical trial design and research priority setting. *Health Economics* 1996; 5: 513-524.
Felli JC. and Hazen GB. Sensitivity analysis and the expected value of perfect information. *Medical Decision Making* 1998; 18: 95-109
Hoel, DG. Sodel, M. and Weiss, GH. A survey of adaptive sampling for clinical trials. In: Elashoff, R. editor. *Perspectives in Biometry*. New York: Academic Press, 1975.
Hornberger J, Eghtesady P. The cost-benefit of a randomized trial to a health care organization. *Controlled Clin Trials* 1998;19:198-211.
Igiewicz, B. Alternative designs: sequential, multi-stage, decision theory and adaptive designs. In: Buyse, ME. Staquet, J. and Sylvester, RJ. editors. *Cancer clinical trials: methods and practice*. Oxford University Press, 1983.
Phelps, CE. and Mushlin, AI. On the near equivalence of cost-effectiveness and cost-benefit analysis. *International Journal of Technology Assessment in Health Care* 1991; **7**: 12-21.
Raiffa H. and Schlaifer, R. *Probability and Statistics for Business Decisions*. McGraw-Hill, 1959.
Raiffa, H. *Decision analysis: introductory lectures on choices under uncertainty*. Addison-Wesley, 1968.
Schlaifer, R. *Introduction to Statistics for Business Decisions*. McGraw-Hill, 1961.
Spiegelhalter DJ., Freedman LS. and Parmar MKB. Bayesian approaches to randomised trials. *Journal of the Royal Statistical Society A* 1994; 157: 357-416.
Stinnett A. and Mullahy J. (1998). Net health benefits: a new framework for the analysis of uncertianty in cost-effectiveness analysis Medical Decision Making.Vol 18, pp. S68-S80.
Thompson KM. and Graham JD. Going beyond the single number: using probabilistic risk assessment to improve risk management. *Human and Ecological Risk Assessment* 1996; 2: 1008-1034.
Thompson KM. and Evans JS. The value of improved national exposure information for perchloroethylene (perc): a case study for dry cleaners. *Risk Analysis* 1997; 17:253-271.
Upton, GJG. and Lee, RD. The importance of the patient horizon in the sequential analysis of binomial clinical trials. *Biometrika* 1976; 63: 335-342.
Weinstein, MC. and Zeckhauser, RJ. Critical ratios and efficient allocation. Journal of Public Economics 1972; **2**: 716-721.
Zelen, M. Play the winner rule and the controlled clinical trial. Journal of the American Statistical Association 1969; 64: 131-146.

Appendix A

Table 1 Numerical Examples

	Prior Mean	Prior SD	Population SD	Quasi prior sample size (n_0)
U_{11}	6	0.5164	1.2649	12
U_{01}	8	0.5164	1.2649	12
U_{10}	2	0.2582	0.8942	6
U_{00}	10	0.2582	0.8942	6
U_{12}	4	0.5164	1.2649	12
U_{02}	4	0.5164	1.2649	12
C_{11}	£12,000			
C_{01}	£8,000			
C_{10}	0			
C_{00}	0			
C_{12}	£2,400			
C_{02}	£2,400			
C_{te}	£8,000			
p(D)	0.6	0.1	0.4899	24
$p(t_e^+|D)$	0.9	0.0866	0.3	12
$p(t_e^-|ND)$	0.8	0.1155	0.4	12

$p(t_e^+|D)$ = probability of a true positive result (sensitivity)
$p(t_e^-|ND)$ = probability of a true negative result (specificity)

Appendix B

Optimal allocation at stage 5

The payoff from a sample entering stage 5 ($n_{(5)}$)with n_{t1} allocated to t_1 and $n_{(5)}$-n_{t1} allocated to t_2 is the expected net benefit of sampling or the difference between the benefit of sample information ($EVSI_{(5)}|n_{(5)}, n_{t1}$) and the cost of sampling($Cs_{(5)}|n_{(5)}, n_{t1}$)

$\Pi_{(5)}|n_{(5)},n_{t1} = ENBS_{(5)}|n_{(5)},n_{t1} = EVSI_{(5)}|n_{(5)}, n_{t1} - Cs_{(5)}|n_{(5)}, n_{t1}$

where: $Cs_{(5)}|n_{(5)},n_{t1} = (E(C|t_1)-E(C|t_0)).n_{t1} + (E(C|t_2)-E(C|t_0)).(n_{(5)}-n_{t1})$

$\delta_{0(5)} = (E(U|t_e^+, t_1)-g.E(C|t_e^+, t_1)) - (E(U|t_e^+, t_2)-g.E(C|t_e^+, t_2))$

The optimal allocation of the sample entering stage 5 (n_{t1}*) will be where $\Pi_{(5)}|n_{(5)},n_{t1}$ reaches a maximum for each $n_{(5)}$ considered.

Optimal allocation at stage 4 ($\Pi_{(4)}|n_{(4)},n_{t1}^*$)

The payoff from a sample entering stage 4 ($n_{(4)}$)with n_{t1} allocated to t_1 and $n_{(4)}$-n_{t1} allocated to t_2 is the expected net benefit of sampling or the difference between the benefit of sample information ($EVSI_{(4)}|n_{(4)}, n_{t1}$) and the cost of sampling($Cs_{(4)}|n_{(4)}, n_{t1}$)

$\Pi_{(4)}|n_{(4)},n_{t1} = ENBS_{(4)}|n_{(4)},n_{t1} = EVSI_{(4)}|n_{(4)}, n_{t1} - Cs_{(4)}|n_{(4)}, n_{t1}$

where: $Cs_{(4)}|n_{(4)},n_{t1} = (E(C|t_1)-E(C|t_0)).n_{t1} + (E(C|t_2)-E(C|t_0)).(n_{(4)}-n_{t1})$

$\delta_{0(4)} = (E(U|t_1)-g.E(C|t_1)) - (E(U|t_2)-g.E(C|t_2))$

Optimal allocation at stage 3

A sample entering stage 3 ($n_{(3)}$) which is allocated between the treatment and no treatment arms will generate an expected net benefit of sampling at stage 3, but the sample allocated to the treatment arm will enter stage 4 and will be allocated optimally between t_1 and t_2 generating a payoff of $\Pi_{(4)}|n_{(4)},n_{t1}^*$. So there is a recursive relationship between the payoff at stage 3 and stage 4, where the payoff and the optimal allocation at stage 3 is determined by the ENBS at stage 3 and the payoff given an optimal allocation policy at stage 4.

$\Pi_{(3)}|n_{(3)},n_{tr} = ENBS_{(3)}|n_{(3)},n_{tr} + \Pi_{(4)}|n_{(4)},n_{t1}^*, \qquad n_{(4)} = n_{tr}$

where: $ENBS_{(3)}|n_{(3)},n_{tr} = EVSI_{(3)}|n_{(3)},n_{tr} - Cs_{(3)}|n_{(3)},n_{tr}$

$Cs_{(3)}|n_{(3)},n_{tr} = 0$

$\delta_{0(3)} = (E(U|t_r)\text{-}g.E(C|t_r)) - (E(U|t_0)\text{-}g.E(C|t_0))$

$(E(U|t_r)\text{-}g.E(C|t_r)) = p(\delta_{x(4)}>\delta_{x(4)}^*).(E(U|t_1)\text{-}g.E(C|t_1)) +$
$1\text{-}p(\delta_{x(4)}>\delta_{x(4)}^*).(E(U|t_2)\text{-}g.E(C|t_2))$

$\delta_{x(4)}^* = ((I_{0(4)} + I_{x(4)})\delta_b - I_{0(4)}.\delta_{0(4)})/I_{x(4)},$

$I_{0(4)} = 1/\sigma_{0(4)}^2, \; I_{x(4)} = 1/\sigma_{n(4)}^2$

$\sigma_{0(3)}^2$ and $\sigma_{(3)}^2 = f(n_{(4)},n_{t1}),$

Optimal allocation at stage 2

A sample entering stage 2 ($n_{(2)}$) which is allocated to the test (n_{te}) and no test alternatives ($n_{(2)}\text{-}n_{te}$) will generate an expected benefit and cost of sampling at stage 2. Also $n_{(2)}\text{-}n_{te}=n_{(3)}$ will be allocated optimally at stage 3 generating a payoff of $\Pi_{(3)}|n_{(3)},n_{tr}^*$, and $p(t_e^+).n_{te}=n_{(5)}$ will be allocated optimally at stage 5 generating a payoff of $\Pi_{(5)}|n_{(5)},n_{t1}^*$. This is a recursive relationship where the payoff and the optimal contingent allocation at stage 2 is determined by the expected net benefits of sampling at stage 2 but also by the payoffs given an optimal allocation policy at subsequent stages.

$\Pi_{(2)}|n_{(2)},n_{te} = ENBS_{(2)}|n_{(2)},n_{te} + \Pi_{(3)}|n_{(3)},n_{tr}^* + \Pi_{(5)}|n_{(5)},n_{t1}^*, \; n_{(3)} = n_{(2)}\text{-}n_{te}, \; n_{(5)} = p(t_e^+).n_{te}$

where: $ENBS_{(2)}|n_{(2)},n_{te} = EVSI_{(2)}|n_{(2)},n_{te} - Cs_{(2)}|n_{(2)},n_{te}$

$Cs_{(2)}|n_{(2)},n_{te} = C_{te}.n_{te}$

$\delta_{0(2)} = (E(U|t_e)\text{-}g.E(C|t_e) - (E(U|nt_e)\text{-}g.E(C|nt_e))$

$E(U|nt_e)\text{-}g.E(C|nt_e) = p(\delta_{x(3)}>\delta_{x(3)}^*).(E(U|t_r)\text{-}g.E(C|t_r)) +$
$1\text{-}p(\delta_{x(2)}>\delta_{x(2)}^*).(E(U|t_0)\text{-}g.E(C|t_0))$

$E(U|t_e)\text{-}g.E(C|t_e) = p(t_e^-).(E(U|t_e^-, t_0)\text{-}g.E(C|t_e^-, t_0)) +$
$p(t_e^+).(p(\delta_{x(5)}>\delta_{x}^{*}{}_{(5)}).(E(U|t_e^+, t_1)\text{-}g.E(C|t_e^+, t_1))$
$+ \quad 1\text{-}p(\delta_{x(5)}>\delta_{x}^{*}{}_{(5)}).(E(U|t_e^+, t_2)\text{-}g.E(U|t_e^+, t_2)))$

$\delta_{x(3)}^*, \; \delta_{x(3)}^*, \; \sigma_{0(3)}^2$ and $\sigma_{(2)}^2 = f(n_{(3)}, n_{tr}; n_{(5)}, n_{t1})$

Optimal sample size at stage 1

The first stage is simply to select the optimal sample size given that an optimal allocation policy will be followed at each subsequent stage. The payoff given a sample of $n_{(1)}$ selected at stage 1 ($\Pi_{(1)}|n_{(1)}$) is simply the payoff from stage 2 given an optimal allocation of $n_{(2)}$ between the test arm (n_{te}^*) and the no test arm $n_{(2)}\text{-}n_{te}^*$.

$\Pi_{(1)}|n_{(1)} = ENBS_{(1)}|n_{(1)} = \Pi_{(2)}|n_{(2)},n_{te}^*$

$ENBS_{(1)}|n_{(1)} = ENBS_{(2)}|n_{(2)},n_{te}^* + ENBS_{(3)}|n_{(3)},n_{tr}^* + ENBS_{(4)}|n_{(4)},n_{t1}^* + ENBS_{(5)}|n_{(5)},n_{t1}^*$

$n_{(2)} = n_{(1)}, \; n_{(3)} = n_{(2)}\text{-}n_{te}^*, \; n_{(4)} = n_{tr}^*, \; n_{(5)} = p(t_e^+).n_{te}^*$

The optimal sample size at stage 1 ($n_{(1)}^*$) can be selected and is where $\Pi_{(1)}|n_{(1)}$ or the $ENBS_{(1)}|n_{(1)}$ reaches a maximum.

Figure 1a
A Single-Stage Decision Problem (a two stage dynamic programme)

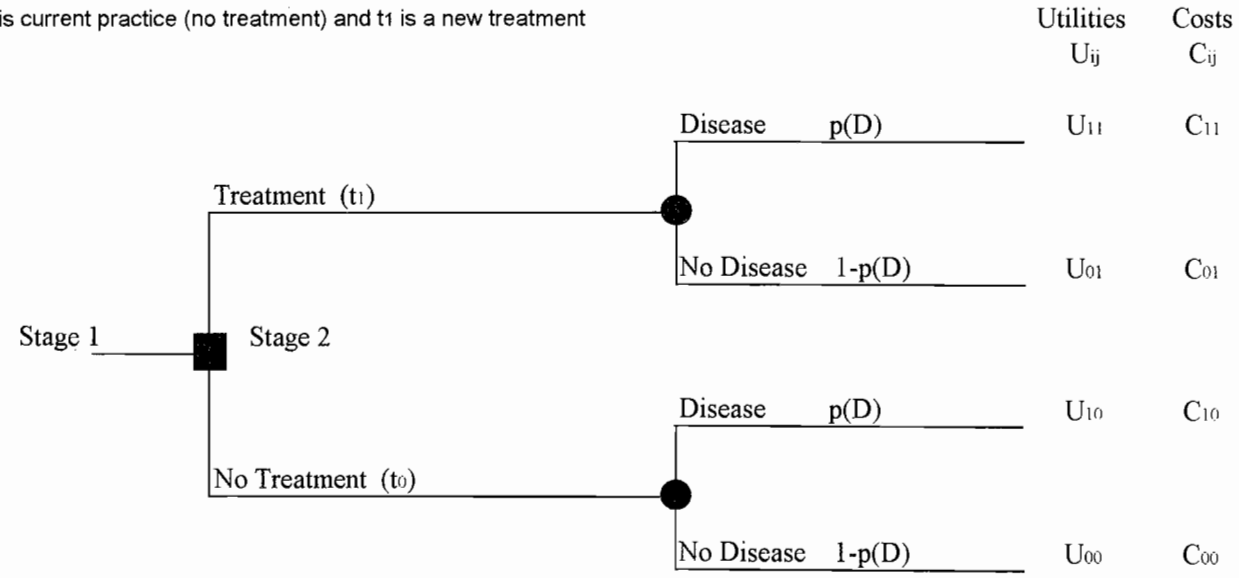

Where:

t0 is current practice (no treatment) and t1 is a new treatment

p(D) is the prior probability of disease

Figure 1b
Single Stage Problem - Optimal Allocation at Stage 2 (1/g=£4,000)

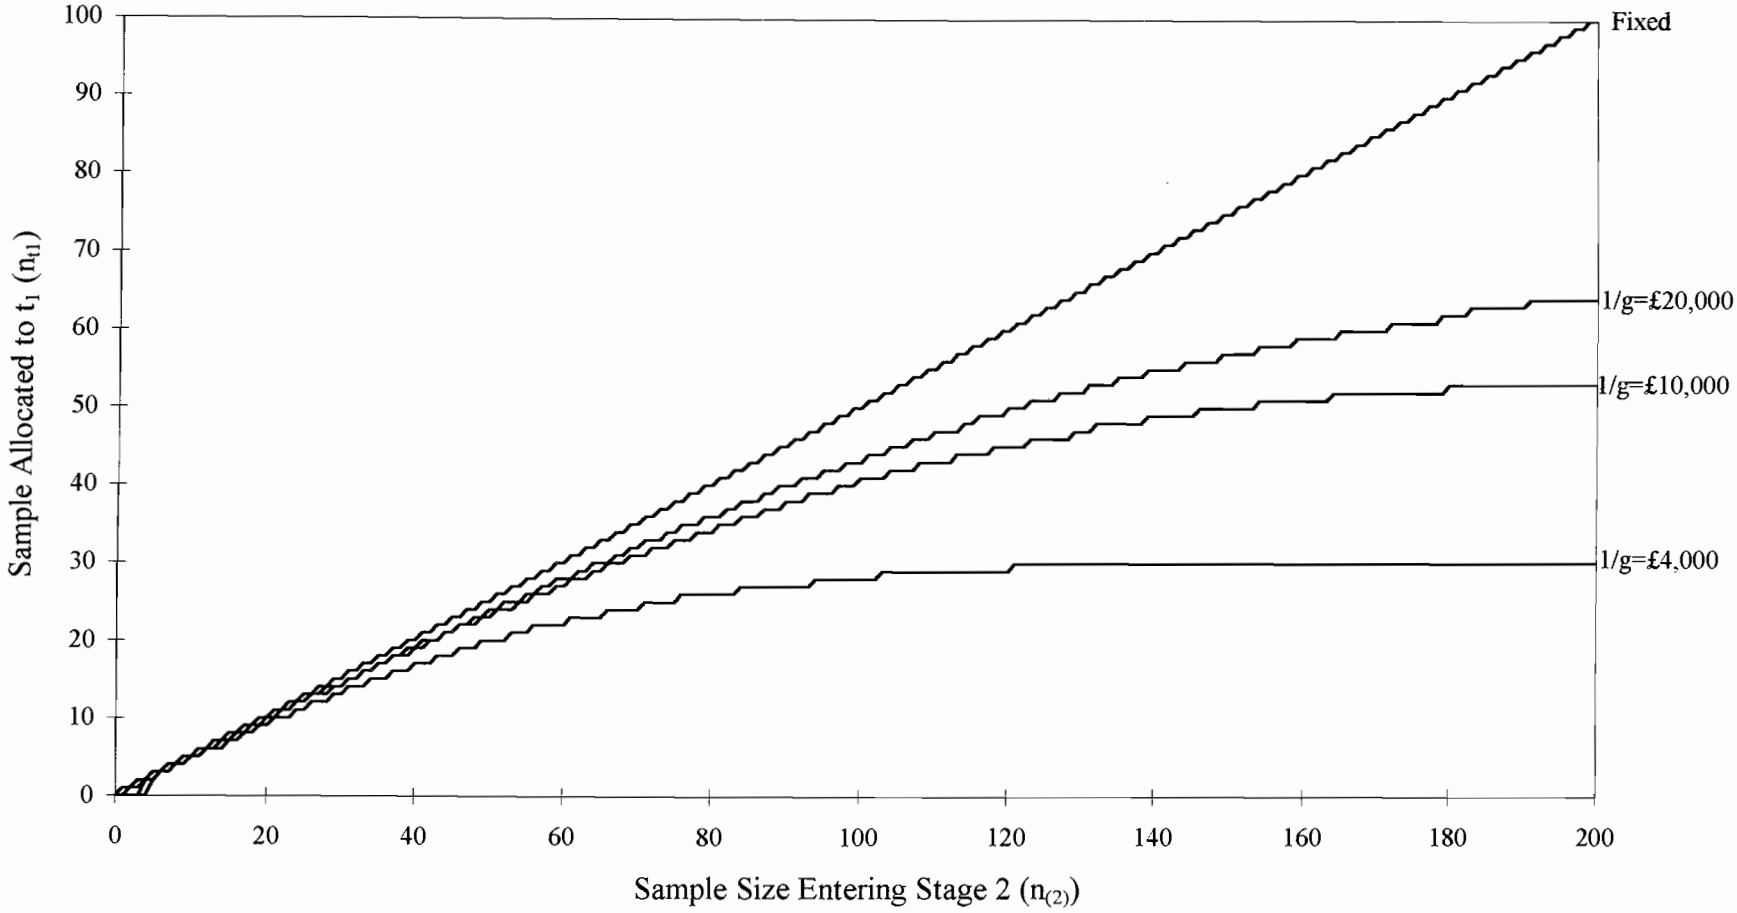

Figure 1c
Single Stage Problem - Difference in ENBS (1/g=£4,000)

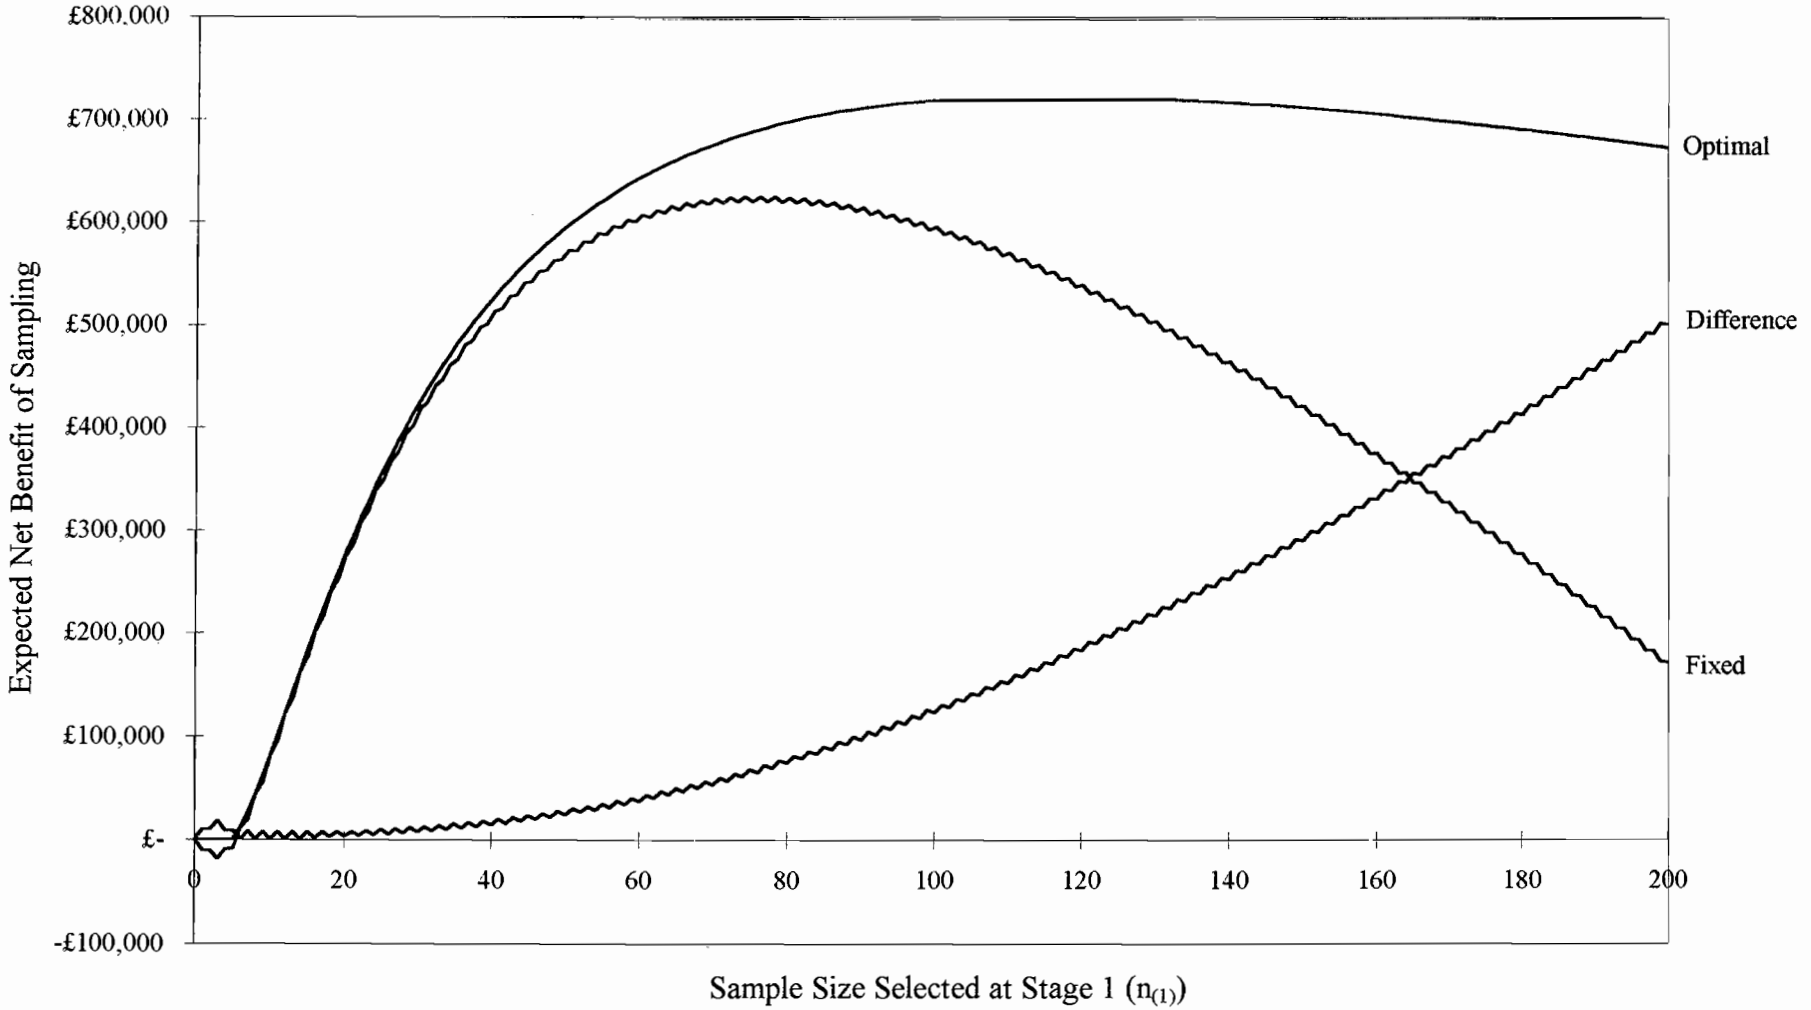

Figure 2a
A Two Stage Decision Problem (a three stage dynamic programme)

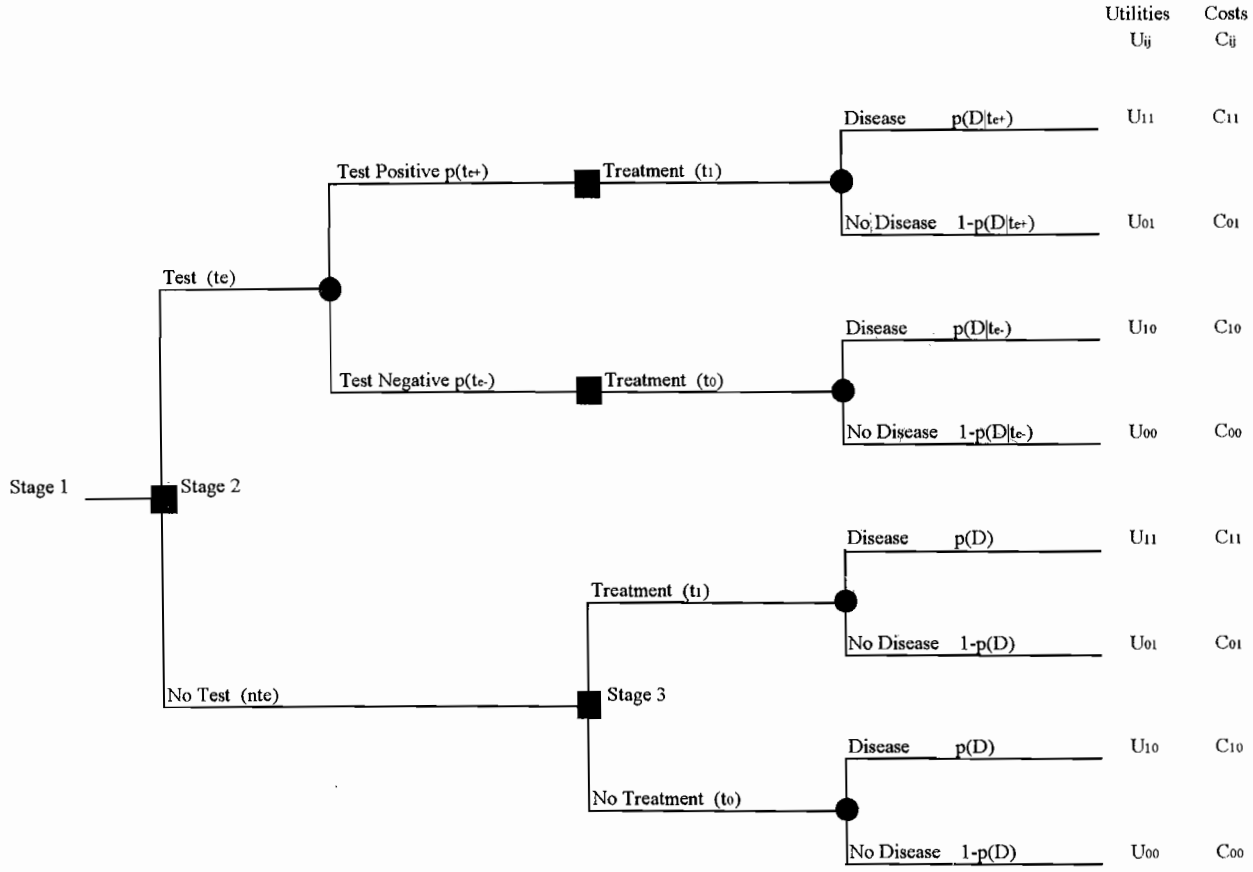

Where:
$p(te+)$ is the probability of a test positive result
$p(D|te+)$ is the probability of disease conditional on a positive test (predictive value positive)

Figure 2b
Two Stage Problem - Optimal Allocation of Sample at Stage 2 to te (n_{te})

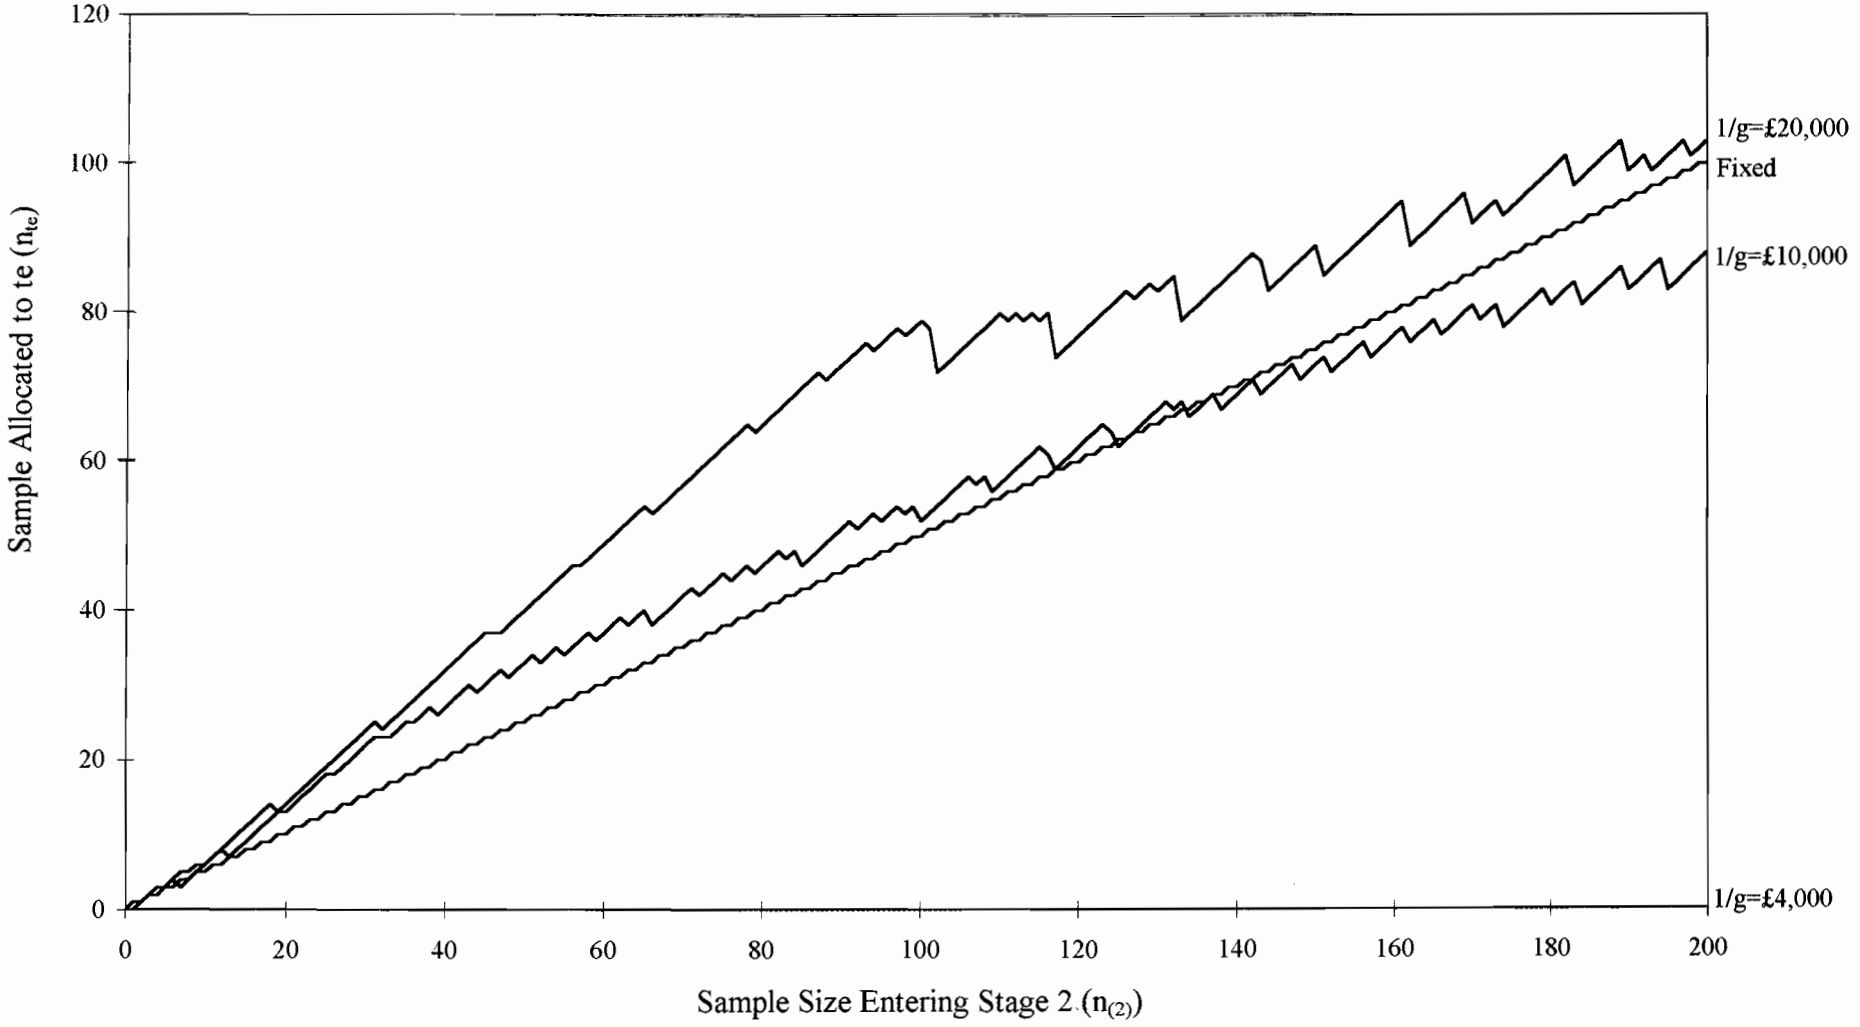

Figure 2c
Two Stage Problem - Difference in ENBS (1/g=£4,000)

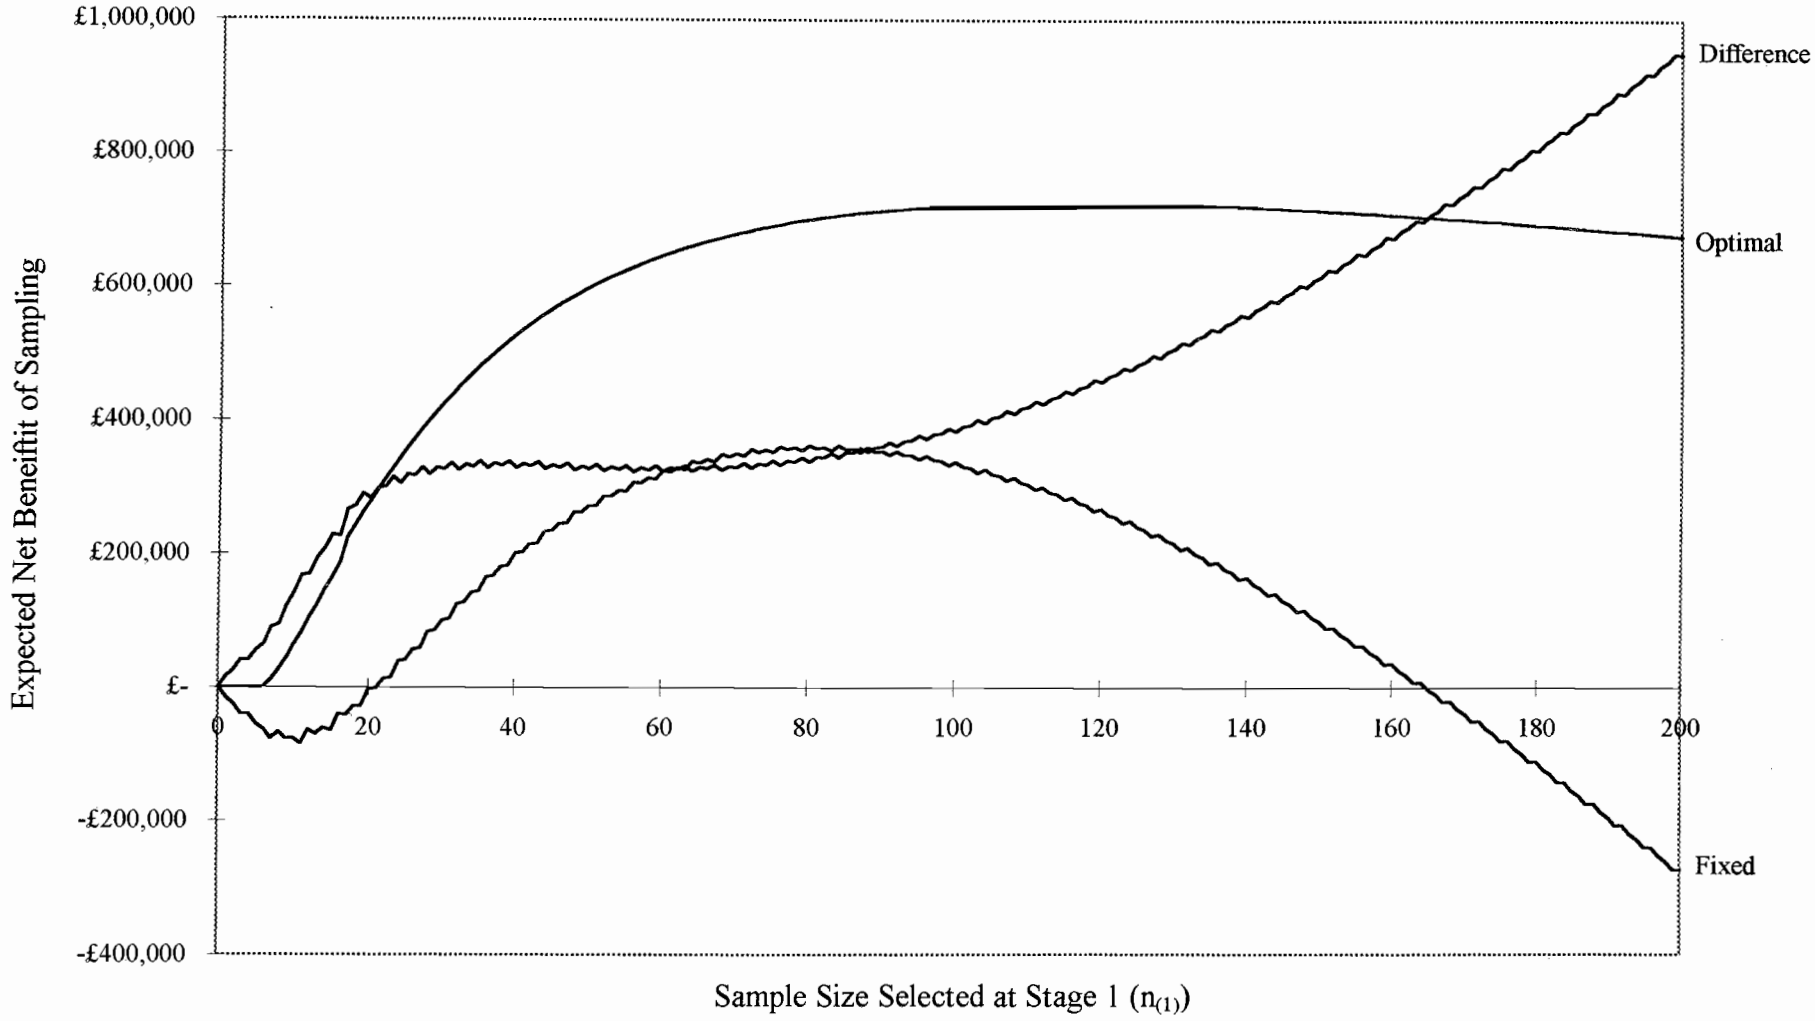

Figure 3a
A Four Stage Decision Problem

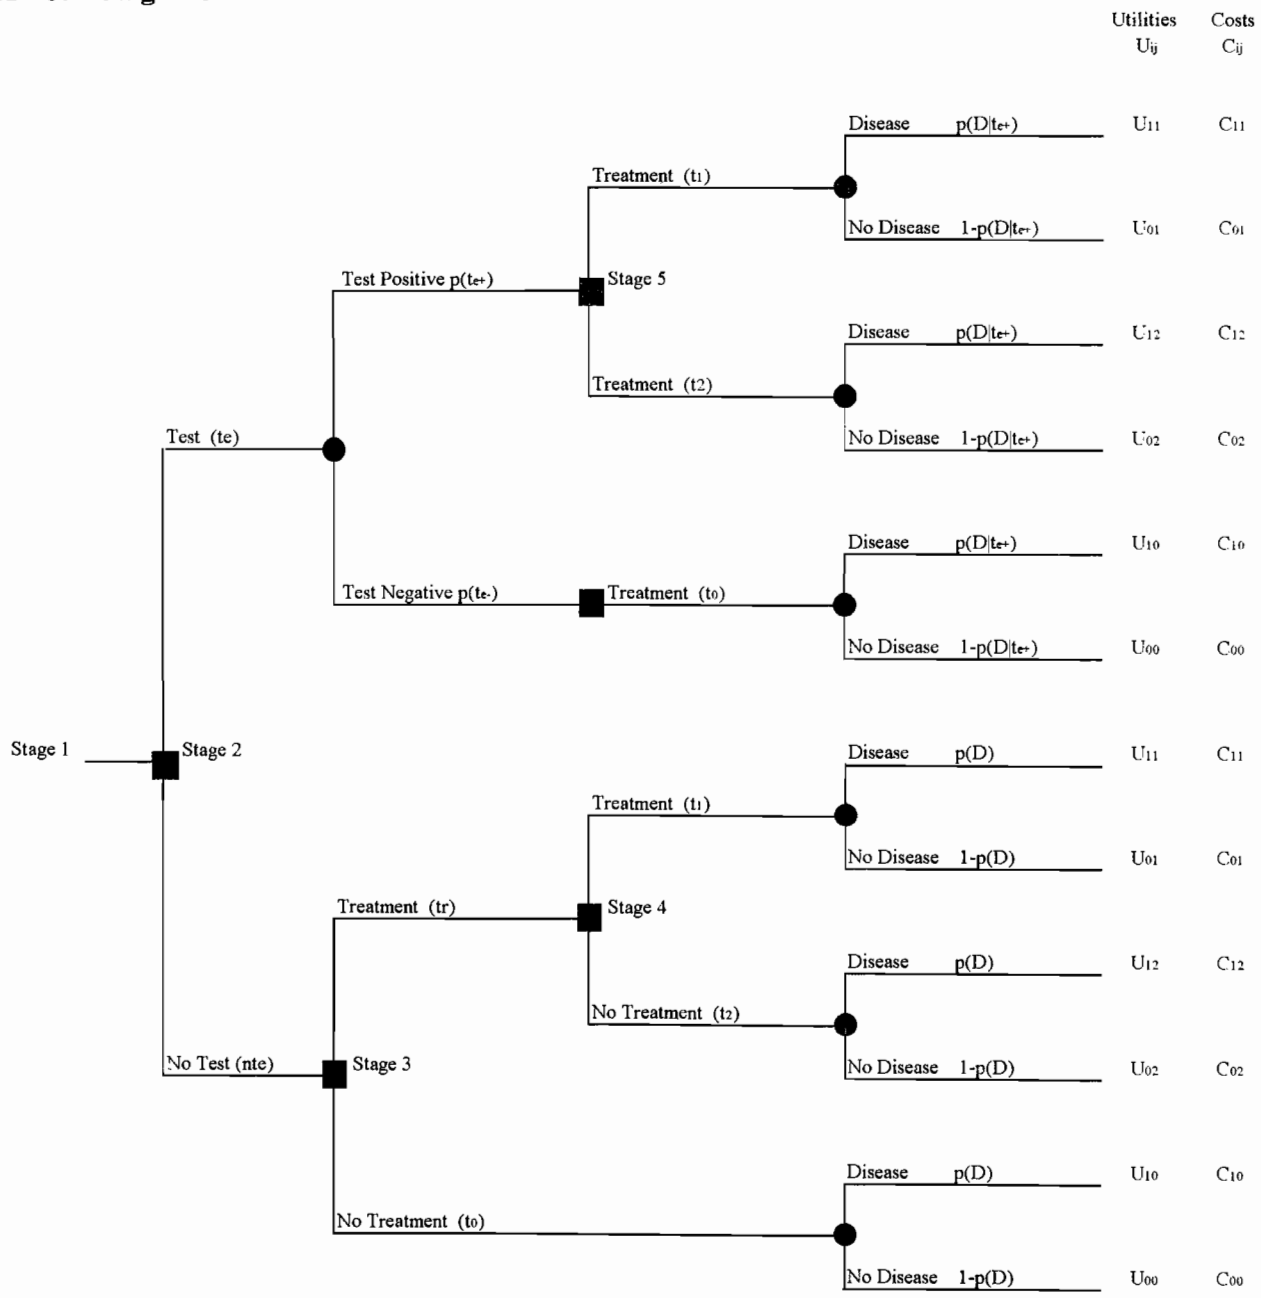

Where:
$p(te+)$ is the probability of a test positive result
$p(D|te+)$ is the probability of disease conditional on a positive test (predictive value positive)

Figure 3b

Four Stage Problem - ENBS with Optimal Allocation (1/g=£4,000)

($ENBS_{(1)} = ENBS_{(2)}+ENBS_{(3)}+ENBS_{(4)}+ENBS_{(5)}$)

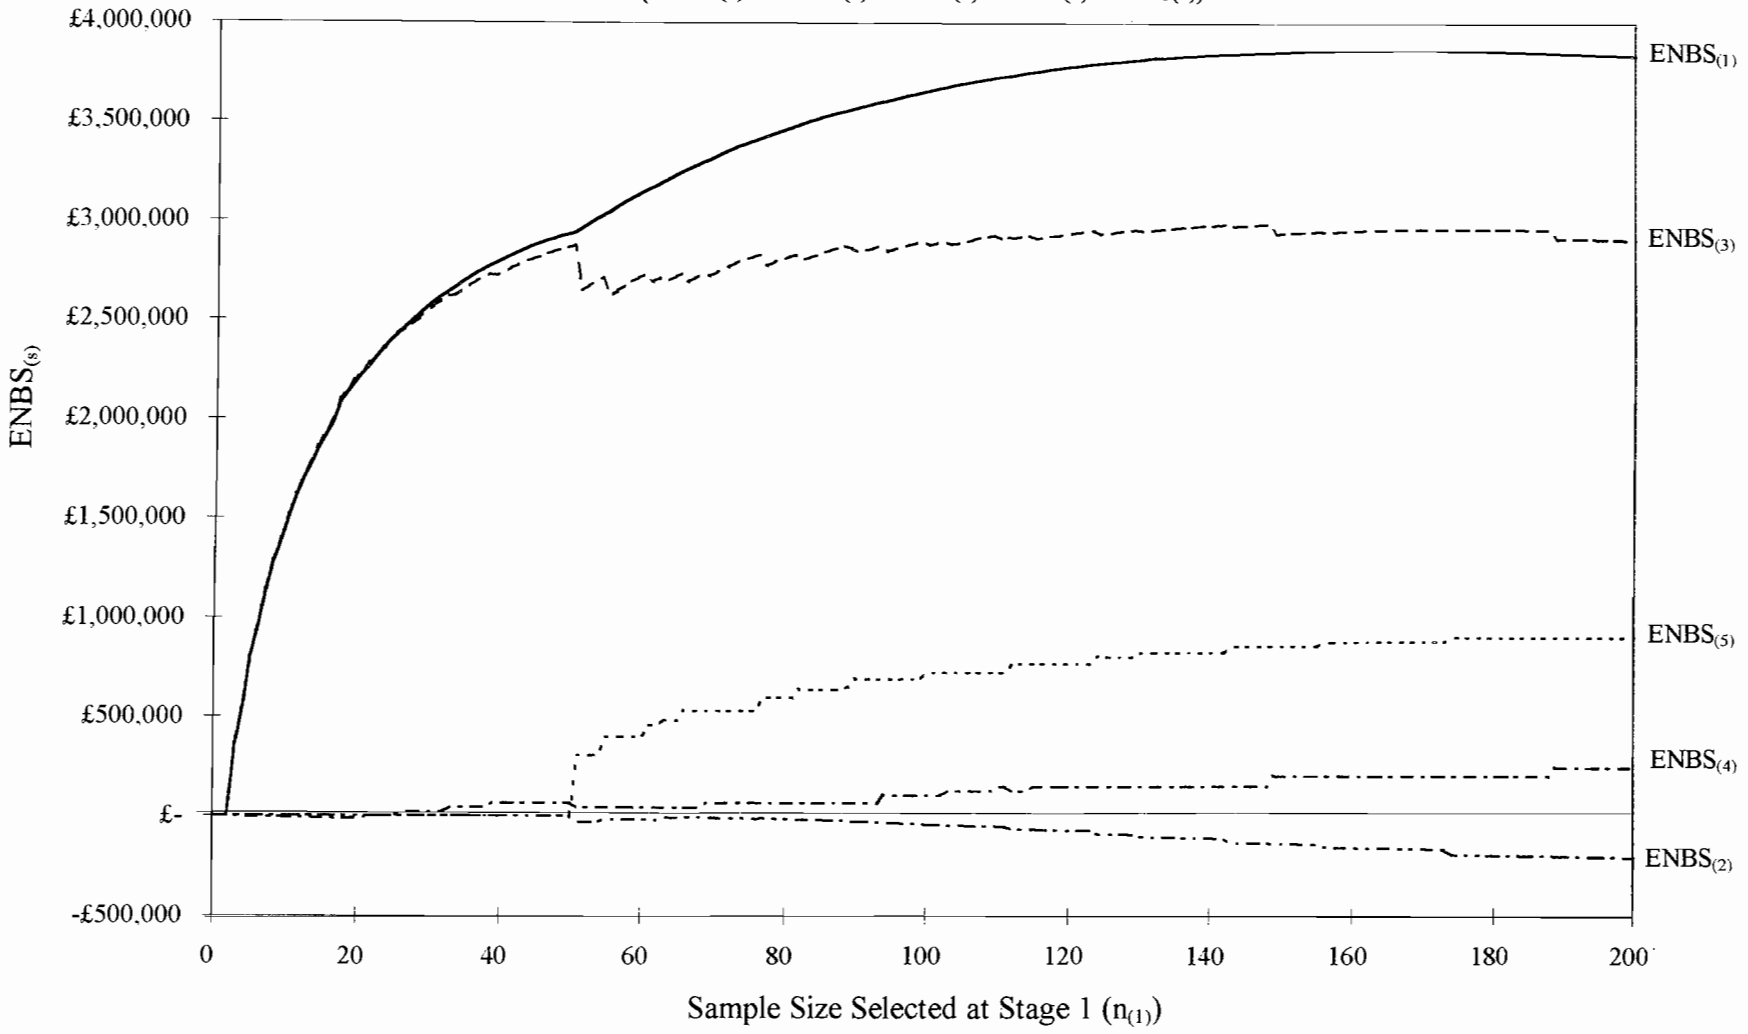

Figure 3c
Four Stage Problem - Difference in ENBS (1/g=£4,000)

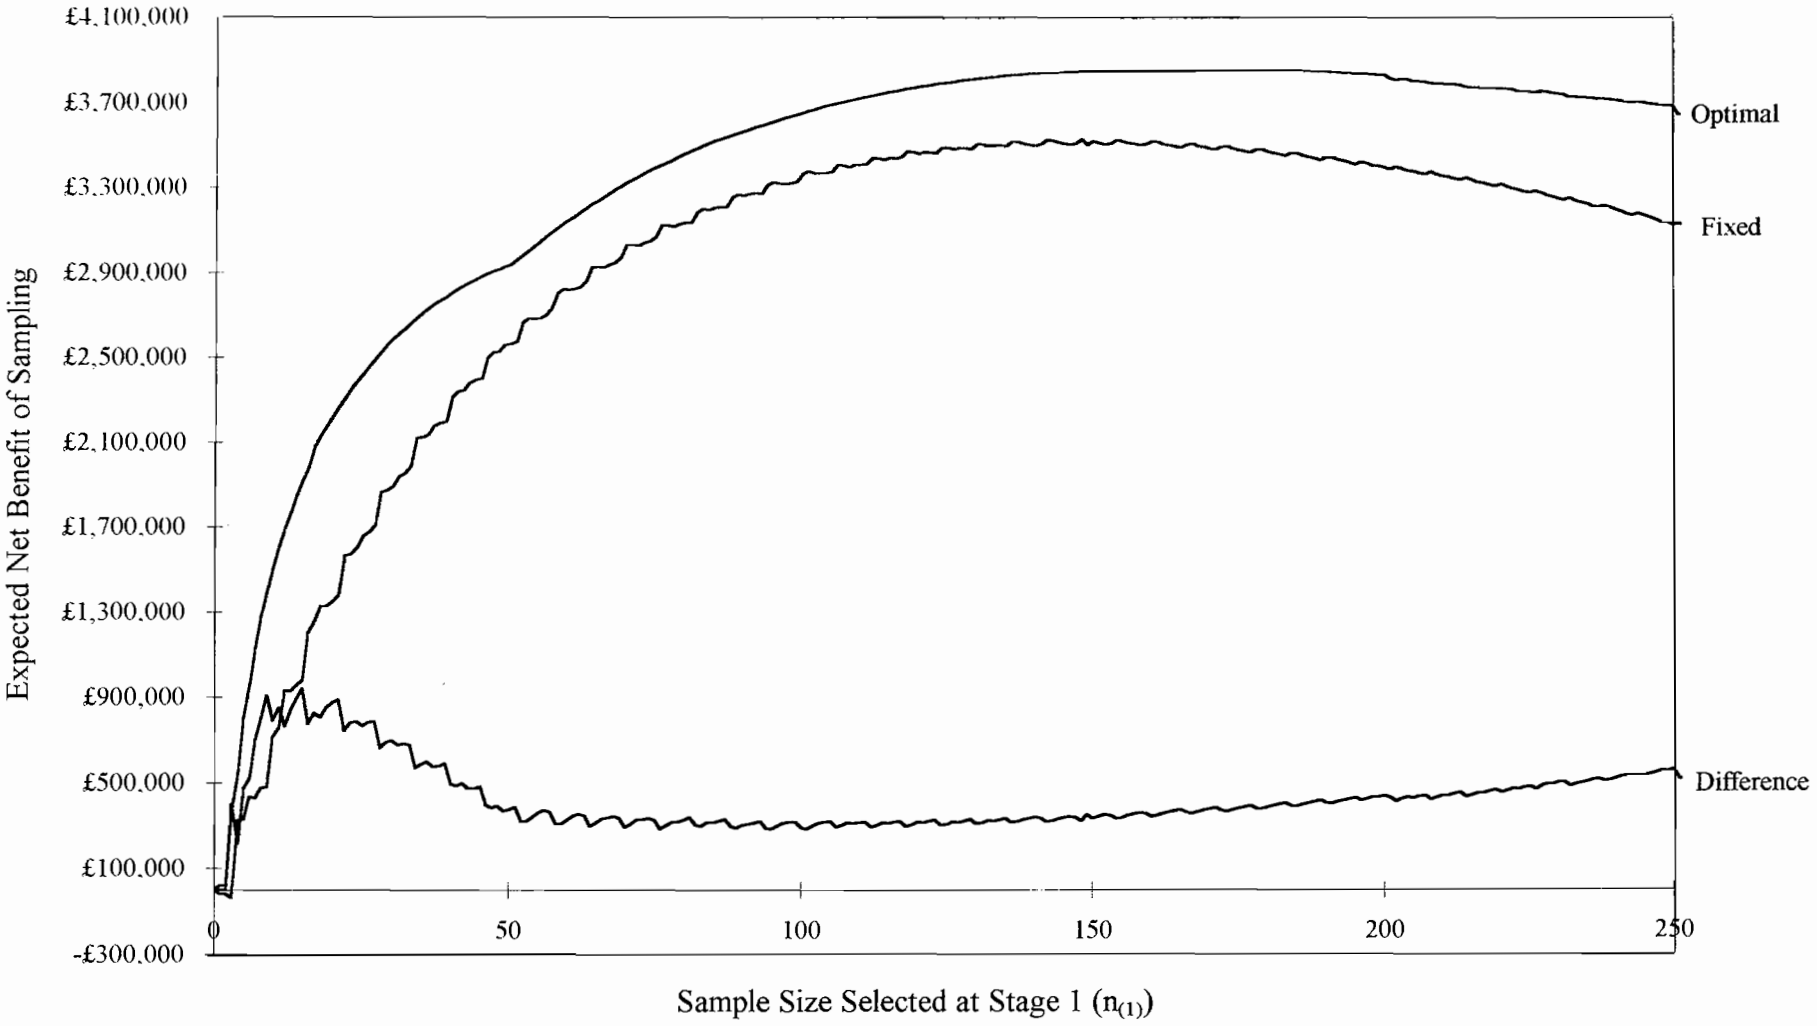

Figure 4
ENBS for the Single Two and Four Stage Problem