



THE UNIVERSITY *of York*

*Discussion Papers in Economics*

No. 2010/24

**Optimal Sequential Sampling Rules for the Economic Evaluation of Health Technologies**

By

**Paolo Pertile, Martin Forster, Davide La Torre**

Department of Economics and Related Studies  
University of York  
Heslington  
York, YO10 5DD

# Optimal sequential sampling rules for the economic evaluation of health technologies

Paolo Pertile\*, Martin Forster,<sup>†</sup> Davide La Torre<sup>‡</sup>

December 1, 2010

## Abstract

Referring to the literature on optimal stopping under sequential sampling developed by Chernoff and collaborators, we solve a dynamic model of the economic evaluation of a new health technology, deriving optimal rules for technology adoption, research abandonment and continuation as functions of sample size. The model extends the existing literature to the case where an adoption decision can be deferred and involves a degree of irreversibility. We explore the model's applicability in a case study of the economic evaluation of Drug Eluting Stents (DES), deriving dynamic adoption and abandonment thresholds which are a function of the model's economic parameters. A key result is that referring to a single cost-effectiveness threshold may be sub-optimal.

**JEL codes:** I10 ; D92 ; C61

**Keywords:** Cost-effectiveness analysis; Sequential sampling; Dynamic programming

---

\*Department of Economics, University of Verona, 37129 Verona, Italy. e-mail: paolo.pertile@univr.it, tel.: +39 045 8028438.

<sup>†</sup>Corresponding author. Department of Economics and Related Studies, University of York, York YO10 5DD, U.K.. e-mail: mf8@york.ac.uk, tel.: +44 (0)1904 433797.

<sup>‡</sup>Department of Economics, Business and Statistics, University of Milan, 20122 Milan, Italy. e-mail: davide.latorre@unimi.it, tel.: +39 02 50321462.

# 1 Introduction

The decision about whether or not to carry out more research to reduce the uncertainty surrounding the estimate of cost-effectiveness associated with a new health care technology has generated a voluminous literature (examples include Claxton (1999), Chalkidou et al. (2007) and Griffin et al. (2010)). Key to such decisions are the estimates of the value of waiting for new information, which must balance the expected benefit associated with the reduction in uncertainty that new information provides with the costs of obtaining that information: sampling, simulation or research costs themselves, together with any expected benefits foregone if action is delayed. Central in this literature is the idea that, if the decision to adopt or not adopt the new technology cannot be deferred, it should be separated from the decision about whether to continue to research (Claxton, 1999, pages 347 - 350). However, assuming that the decision cannot be deferred implies a ‘now-or-never’ view of the adoption decision (Dixit and Pindyck, 1995); to what extent do decision rules change if the decision-maker has the option to delay the decision and carry out further research? More than ten years ago, Palmer and Smith (2000) recognised that such a scenario, characterised as it is by uncertainty over future states of the world and flexibility over the timing of a decision with potential irreversibilities, fits well within a real option framework. To date, however, there has been limited application of the real option approach to health technology assessment (exceptions being Pertile (2009a,b)) and little application elsewhere in the health economics literature (examples include Driffield and Smith (2007) and Levaggi and Moretto (2008)).

Such limited uptake could be due, in part at least, to criticism of the suitability of stochastic processes used by some real option models (such as geometric Brownian motion) for use in economic evaluation, and the fundamental differences between decisions which are based upon the ‘passive’ accumulation of information (such as the evolution of stock market prices) and those in which decision-makers must decide whether or not more research, such as additional clinical trials, should be commissioned (Eckermann and Willan, 2008). In this paper, we argue that these problems can, to some extent, be overcome by adopting the framework of a Bayesian decision-maker developing optimal policy rules under sequential sampling that was developed in the mathematical statistics literature by Chernoff and collaborators during the 1960s and 1970s (Chernoff and Ray, 1965; Chernoff, 1961, 1967, 1972) and which has recently been extended and linked to the real option literature and the literature on the value of information (Lai and Lim, 2005;

Chick and Gans, 2009).<sup>1</sup>

Chernoff showed how a sequential decision problem to test for the mean of a normal random variable in the presence of sampling costs could be recast in continuous time and modelled as a free boundary problem using a Brownian motion with drift as the stochastic process. Approximate solutions to the problem could then be obtained in (sample size  $\times$  posterior estimate of mean) space as thresholds which separate the region in which it is optimal to continue sampling from the regions in which the null hypothesis should be rejected in favour of the appropriate alternative. Chernoff's approach involves solving a partial differential equation derived using Bellman's principle of optimality: '[A]n optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision' (Bellman, 1957). Put another way, a decision-maker who, today, faces the choice about whether to adopt a health care technology based upon existing information or to commission more research, must make their decision accounting for the fact that they should act optimally in all possible future states of the world; the value of one's 'options' today must incorporate valuation of one's optimal behaviour in the future.

The aim of the present paper is to bridge the gap that exists between the absence of key economic variables from standard statistical approaches to HTA and the absence of important statistical dimensions from the real option approach to HTA, in situations when an adoption decision can be deferred. We interpret the health technology assessment as a project with uncertain returns, which depend on clinical (effectiveness) and economic (cost) variables, including the costs of treatment, carrying out research, sunk investments and delaying adoption. We apply the model in a case study of the economic evaluation of Drug Eluting Stents and compare our results with two alternative approaches: a traditional, inferential, statistical approach and the 'irrelevance of inference' approach (Claxton, 1999). Our results highlight the importance of considering the size of the population to be treated should adoption take place and the degree of (im)patience of the decision maker, which are typically ignored in the statistical approach. They also define a dynamically optimal policy for the abandonment of research, which is typically overlooked in the literature.

Section 2 outlines the theory behind sequential sampling for an unknown population mean, section 3 presents our case study and section 4 discusses

---

<sup>1</sup>Although Chick and Gans (2009) focus on simulation selection problems, their methods are applicable also to problems involving sampling, based as they are on the work of Chernoff.

the results. Section 5 concludes.

## 2 Optimal rules under sequential sampling

Our model is based on that proposed by Chick and Gans (2009) for comparing a single simulation system to a known alternative. A risk-neutral decision-maker (DM) is carrying out a project to investigate whether to adopt a new health care technology which, in the case of adoption, will provide benefit to  $P$  patients such that, if the decision is made to adopt the technology,  $P$  is the number of patients that the DM will commit to treating in the interval  $(0, \bar{s})$ , where  $\bar{s}$  is the time at which the DM is able to revise his decision.<sup>2</sup> We define the project as including both the process of research and the eventual decision (whether to adopt the technology or to abandon research). The DM faces a sunk investment cost  $I \geq 0$  if the technology is adopted.

Let the impact of the technology at the individual level be expressed in terms of net incremental monetary benefit (NIMB)  $W$ :

$$W = \lambda(E_1 - E_0) - (C_1 - C_0), \quad (1)$$

where  $\{E_1, C_1\}$  and  $\{E_0, C_0\}$  are pairs of random variables denoting the effectiveness and costs of the new and existing (base) technology, respectively, and  $\lambda$  is the acceptance threshold for the Incremental Cost-Effectiveness Ratio (ICER), such that  $W > 0$  corresponds to  $ICER < \lambda$ .<sup>3</sup> The DM assumes that  $W \sim N(\mu_W, \sigma_W^2)$ , where  $\mu_W$  is unknown and  $\sigma_W^2$  is known. Let  $X = PW$  be the random variable representing the total NIMB associated with adopting the technology.<sup>4</sup> Given the assumptions on  $W$ , the distribution of  $X$  is Gaussian with unknown mean  $\mu_X = P\mu_W$  and known variance  $\sigma_X^2 = P^2\sigma_W^2$ .

---

<sup>2</sup>The size of  $P$  depends on a number of factors, including  $\bar{s}$ , the prevalence and incidence of the disease requiring treatment, the nature of the technology (the decision to reverse the use of a technology can be made more easily for some technologies than for others) and the characteristics of the decision-maker. For example, the DM could be the authority responsible for the decision to admit the use of the new technology in a particular jurisdiction, where the jurisdiction could be defined at a regional level or a national one.

<sup>3</sup>Consistent with our dynamic approach, unlike in standard cost-effectiveness analyses, the per patient component of  $I$  does not enter the measure of NIMB because, as a sunk cost, it is incurred the time of adoption.

<sup>4</sup>We assume an exogenous ‘arrival process’ of patients presenting for treatment, such that the rate of arrival equals the rate at which patients are treated (with the base or the new technology), implying that  $P$ , the instantaneous stock of patients to be treated, is constant over time. Strictly speaking, benefits upon adoption should be discounted according to  $\int_0^{\bar{s}} pWe^{-\rho t} ds$ , where  $\bar{s}$  is the time at which the adoption decision can be revised and  $\int_0^{\bar{s}} pds = P$ . For the sake of simplicity, we ignore this discounting, which is

The DM has prior beliefs about  $\mu_X$ , represented by a Gaussian random variable  $\mathcal{X}$  with expected value  $\mu_{\mathcal{X},0}$  and variance  $\sigma_{\mathcal{X},0}^2$ . Assume the DM can observe sequentially drawn, independent and identically distributed realisations of  $W$ , such that  $X_i = PW_i, i = 1, \dots, N$ , where  $N$  could feasibly equal infinity, by paying a constant, per-realisation, research cost  $c$ . We interpret this process as the DM carrying out, or monitoring the results of, research (such as a trial). After observing each realisation, the DM updates his beliefs according to Bayes' theorem. The posterior distribution, after observing  $m$  realisations, will therefore also be Gaussian with  $E[\mathcal{X}_m] = (\mu_{\mathcal{X},0}t_0 + \sum_{i=1}^m X_i)/(t_0 + m)$  and  $\text{var}(\mathcal{X}_m) = \sigma_X^2/(t_0 + m)$ , where  $t_0 = \sigma_X^2/\sigma_0^2$ . We define each realisation of  $W$  as a stage of the project, immediately after which point the DM must decide to take one of the following actions: to adopt the technology immediately, to pay for another observation (to continue with research) or to abandon research. Paying for another observation reduces the DM's uncertainty concerning the value of  $\mu_X$  by reducing  $\text{var}(\mathcal{X}_i)$ , and so reduces the chances that the DM makes an incorrect decision, but it costs  $c$  and delays the payoff (of adoption or abandonment). In choosing the optimal policy rule, the DM's objective is to maximise total net incremental monetary benefit minus project costs, which comprise the research costs  $c$ , the sunk cost  $I$  that is incurred if adoption takes place, and discounting costs.

Define  $Y_m = \mu_{\mathcal{X},0}t_0 + \sum_{i=1}^m X_i$  and  $t = t_0 + m$ . The DM's updating process for  $E[\mathcal{X}_m]$  may be recast in continuous time ( $m$  becomes continuous and therefore so does  $t$ ) in which  $\text{var}(\mathcal{X}_t) = \sigma_X^2/t$  follows a deterministic path and  $Y_t$  follows a random walk with drift  $\mu_X$  and variance  $\sigma_X^2$ :

$$dY_t = \mu_X dt + \sigma_X dV_t, \quad (2)$$

where  $dV_t$  is the increment of a Wiener process. The expected value of the DM's posterior beliefs concerning  $\mu_X$  at any  $t$  is given by  $E[\mathcal{X}_t] = y_t/t$  and  $\text{var}(\mathcal{X}_t) = \sigma_X^2/t$ , where  $y_t$  is the realisation of  $Y_t$ .

Let the continuous time expected discounted reward of the project be defined by the value function  $B(y, t)$ . The DM must obtain the optimal policy rule  $\pi^*$  which, in  $(t, y)$  space, defines the DM's optimal actions from the set  $u_t \in \{1, 2, 3\}$ , where 1 means 'adopt immediately', 2 'pay for another observation' (continue to research) and 3 'abandon research'. The Bellman equation is:

$$B(y_t, t) = \max \left\{ (y_t/t) - I, \quad -cdt + \frac{1}{1 + \rho dt} E[B(y_t + dy, t + dt) | y_t, t], \quad 0 \right\}, \quad (3)$$

---

consistent with the common practice in HTA as long as  $\bar{s}$  is less than one year. Extending the model to account for discounting would be straightforward.

for all  $t$  and  $y$ , where  $\rho$  is the discount factor. The first term in brackets on the right hand side of (3) is the expected value of the DM's posterior beliefs about  $\mu_X$ , the expected total NIMB if the DM adopts the technology, minus the sunk cost  $I$ , the second is the continuation value - the value to the DM of continuing to carry out research - and the third is the value associated with abandoning research and not adopting the technology. The problem is an optimal stopping problem with a free boundary whose solutions are two thresholds for the posterior mean defining three regions such that, above the upper threshold, immediate adoption of the technology is optimal, below the lower threshold, stopping research is optimal and, in the 'continuation region'  $\mathcal{C}$  lying between the two thresholds, continuing to research - that is, paying for another draw from  $W$  - is optimal.

Given the Bellman equation, in  $\mathcal{C}$ , the following equation must hold:

$$B(y_t, t) = -c dt + \frac{1}{1 + \rho dt} E[B(y_t + dy, t + dt) | y_t, t], \quad (4)$$

where:

$$E[B(y_t + dy, t + dt) | y_t, t] = B(y_t, t) + E[dB], \quad (5)$$

and, using Ito's Lemma:

$$dB = B_t dt + B_y dy_t + (1/2) B_{yy} (dy_t)^2, \quad (6)$$

where  $dy_t$  is defined by Eq. (2) ( $y_t$  is the realisation of  $Y_t$ ). Through substitutions we can obtain the partial differential equation satisfied by  $B$  in  $\mathcal{C}$ :<sup>5</sup>

$$0 = -c - \rho B + B_t + B_y \frac{y}{t} + \frac{\sigma_X^2}{2} B_{yy}. \quad (7)$$

The value-matching and smooth pasting conditions which define the upper and lower thresholds are:

$$B(y, t) = D(y, t) = \max\{0, (y/t) - I\} \text{ on } \partial\mathcal{C}; \quad (8)$$

$$B_y(y, t) = D_y(y, t) \text{ on } \partial\mathcal{C}, \quad (9)$$

where  $\partial\mathcal{C}$  is the free boundary of  $\mathcal{C}$  and  $D$  is the value of  $B$  in the adoption and abandonment regions. The objective of the problem is to solve for  $B$ , the project's maximum expected discounted reward, and the free boundaries, which together are defined by (7), (8) and (9).

---

<sup>5</sup>We substitute (6) into (5) and (5) into (4). Applying the expectation operator, noting that  $dV_t^2 = dt$  and sending  $dt$  to 0, gives (7).

To date there exists no closed form solution to this problem. For our case study we use finite difference methods (Dixit and Pindyck (1994, pages 319-339)) to obtain a numerical solution to the PDE and to derive the boundary conditions.

### 3 Case study

In this section, we apply the model of section 2 to the economic evaluation of drug eluting stents (DES) versus bare metal stents (BMS) for an unselected population, using data from a published study (the ‘SIRIUS trial’, reported in Cohen et al. (2004)). We do so for the purposes of illustration only; the case study is not intended to be a contribution to the literature on whether or not DES should be adopted by health care systems.

Stents may be employed in percutaneous coronary intervention (PCI), with BMS the base technology and DES the new technology. DES was launched some years ago with the promise of a potential reduction in the rates of angiographic and clinical restenosis when compared with BMS, albeit at a higher cost, and hence its value in the current health care environment has been questioned. The cost-utility analysis asks whether the benefits of the technology to the patient justify this higher cost.<sup>6</sup> Between February and August 2001, 1058 patients undergoing PCI were randomly assigned to receive either a DES<sup>7</sup> or a BMS. The trial reported the average incremental cost associated with using DES to be \$309 per patient and calculated the incremental cost-effectiveness ratio to be \$27540 per QALY gained. It concluded that sirolimus-eluting stents ‘may be viewed as reasonably attractive’ from the perspective of the U.S. health care system (Cohen et al., 2004).

To apply the model of section 2 to the case study, we assume a noninformative prior for the DM (that is, we assume that  $\sigma_{\mathcal{X},0}^2 = \infty$  and so  $t$  becomes the sample size). We use the information on the point estimates of incremental cost and the ICER to obtain a point estimate of mean NIMB equal to \$252 and we calculate the standard deviation of individual NIMB to be  $\sigma_W = \$17538$ .<sup>8</sup> Since our prior is noninformative, this point estimate is equivalent to the DM’s posterior estimate of expected NIMB. Since we

---

<sup>6</sup>Other studies have concentrated on the adoption of DES for high risk populations, where the benefits of the technology are likely to be greater. Most of these studies indicate cost-effectiveness of DES for these selected patients.

<sup>7</sup>The type of DES used in the trial is Sirolimus.

<sup>8</sup>We estimate incremental effectiveness using a first order Taylor series approximation for the ICER and we use Cohen et al.’s result that  $\Pr(ICER < \$50000) = 0.632$  to infer that  $\Pr(E[NIMB] > 0) = 0.632$ . Under the DM’s assumption that NIMB has a Gaussian distribution, this yields  $\sigma_W = \$17538$ .

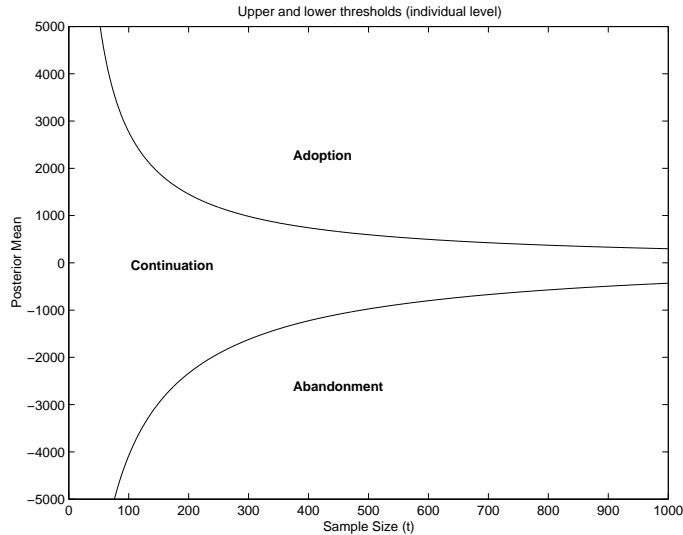


Parameter	Definition	Source	Value
1. $\sigma_{\mathcal{X},0}$	Standard deviation of prior beliefs on $\mu_{\mathcal{X}}$	Assumption	$\infty$
2. $c$	Marginal cost of sampling	Assumption	\$100
3. $C_1 - C_0$	Incremental cost	Cohen et al. (2004)	\$309
4. $E_1 - E_0$	Incremental effectiveness	First order approximation of ICER in Cohen et al. (2004)	0.01122 QALY
5. $\lambda$	Shadow value of effectiveness	Cohen et al. (2004)	\$50000
6. $\sigma_W$	Population standard deviation of individual NIMB	Cohen et al. (2004) result for ICER	\$17538
7. $E[\text{NIMB}]$	Estimate of expected net incremental monetary benefit	Rows 3, 4 and 5 of this table	\$252
8. $\rho$	Discount rate (annual)	Assumption	0.03
9. -	Rate of arrival of observations (annual)	Cohen et al. (2004)	904
10. $P$	Number of patients that will be treated before adoption decision can be revised	Assumption	1000
11. $I$	'Sunk cost' associated with adoption of technology	Assumption	0
12. $T$	Maximum sample size	Assumption	2000

**Table 1:** Parameter values used for the case study (baseline case) of section 3.

are dealing with a medical device, we assume that the irreversible investment cost may be considered negligible and so we set  $I$  equal to zero. We adopt the standpoint of an authority responsible for the decision to introduce DES for the treatment of unselected patients undergoing PCI and, in the absence of information on the value of  $P$ , we run our application assuming that  $P = 1000$ .

In section 2, we implicitly assume that the number of patients allocated to each arm of the trial is the same, allowing us to interpret each pair of observations as equating to one realisation. In the SIRIUS trial the arms are slightly unbalanced, although the split is small: 533 for DES versus 522 for BMS. In order to approximate this situation, we set the number of




**Figure 1:** Optimal thresholds and adoption, continuation and abandonment regions (range for  $t$  restricted to  $(0, 1000)$ ) for the baseline model.

observations at the level such that the variance of the sample mean with 533 individuals in the DES group and 522 individuals in the base group equals the variance of a hypothetical situation in which the numbers in each group are equal and the DM’s prior is noninformative.<sup>9</sup> This implies an equivalent sample size of 527.

Consistent with the model in section 2, a discount rate accounts for the cost of foregone treatments when the decision is postponed. We set  $\rho = 0.03$  (an annual rate). Since the unit of ‘time’ in our problem corresponds to the time needed to take one observation in each arm of the trial, the discount rate must be adjusted for consistency with the sampling process. Assuming a constant pace of accumulation of observations,  $n = 527$  observations taken over the seven months of the trial correspond to 904 observations in one year. Hence the adjusted, per observation, discount rate is  $\rho = 0.03/904$ . Lacking any useful information to estimate the cost of an additional pair of observations, we assume that  $c = \$100$ . Table 1 summarises the parameter values, together with their sources.

---

<sup>9</sup>See Joseph and Belisle (1997).



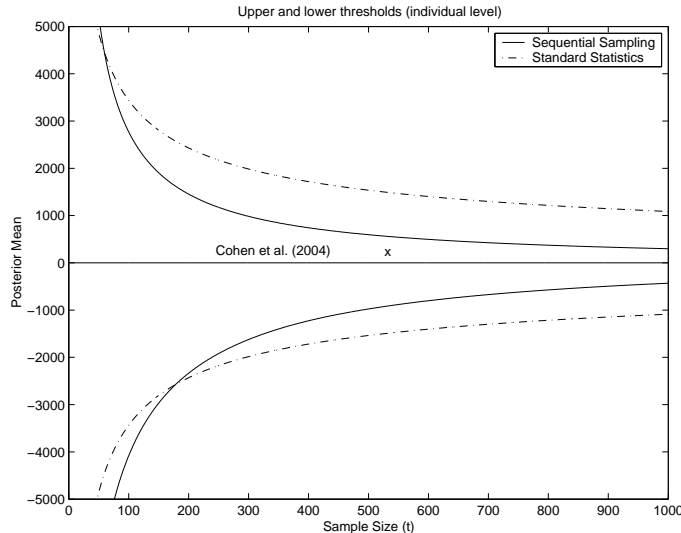
**Figure 2:** Sensitivity of optimal thresholds to changes in: (a) research cost  $c$ ; (b) population standard deviation of individual NIMB ( $\sigma_W$ ).

### 3.1 Optimal thresholds

With the values reported in Table 1, to which we refer as the ‘baseline case’ parameters, we numerically solve Eq. (7), together with the boundary conditions, using the Crank-Nicholson finite-difference scheme and a mesh size of  $2000 \times 2000$ . Figure 1 shows the upper and the lower thresholds and the three regions - adoption, continuation and abandonment - at the individual level, assuming a maximum sample size  $T = 2000$ . In Figure 1, we restrict the range for  $t$  to  $(0, 1000)$  to show clearly the three regions.<sup>10</sup> The area above the upper threshold indicates the set of points  $(t, y)$  for which adoption of DES is optimal. In the area between the thresholds it is optimal to continue with research. Below the lower threshold continuing to research is not optimal, that is, the expected value of additional sample information is outweighed by the costs of obtaining it. The dynamically optimal thresholds are nonlinear and not symmetrical and, when  $I = 0$ , they converge to zero as the sample size increases. It can be optimal for the DM not to adopt, but continue to carry out research, when the posterior estimate of expected NIMB is positive and it can be optimal for the DM to continue with research even though the posterior estimate of expected NIMB is negative.

Figures 2(a) and (b) show the sensitivity of the optimal thresholds to changes in the sampling cost  $c$  and the standard deviation of NIMB,  $\sigma_W$ , where the solid line corresponds to the baseline case of Figure 1. In Figure

<sup>10</sup>Choice of  $T$  affects the positioning of the thresholds, since  $t$  appears in the PDE and associated conditions (Eqs. (7) - (9)) which are solved using the finite difference method. Chick and Gans (2009) use a loss function approach to choose the optimal  $T$ . In this preliminary work we fix  $T$  to illustrate the main characteristics of the model.



**Figure 3:** Comparison of Cohen et al. (2004) point estimate with dynamic thresholds from the baseline model and the statistical thresholds.

2(a), the dotted and the dashed lines are obtained by setting  $c = 300$  and  $c = 10$ , respectively. The figure shows that the continuation region shrinks as the cost of sampling increases, which is an unsurprising result.

As is well known from the real option literature, a greater degree of uncertainty at the population level increases the option value, thus increasing the value of waiting. Figure 2(b) confirms this result, showing how the thresholds increase when the standard deviation is increased/decreased by 30%. The figure shows that a larger population standard deviation implies a larger continuation region: for any  $t$ , the value of the posterior estimate of expected NIMB required to induce optimal adoption is higher. In a similar manner, abandonment of sampling is optimal for lower (higher in absolute terms) values than in the baseline case.

### 3.2 A comparison with alternative decision rules

It is relatively straightforward to compare our thresholds with the adoption/research rules that would be used in a traditional, inferential statistical analysis and in the ‘irrelevance of inference’ (IoI) approach of Claxton (1999). To assist our assessment, we assume a hypothetical situation in which we are running the SIRIUS trial and, having sampled 527 pairs of subjects, we wish to decide whether we should adopt DES, continue to sample, or abandon our project, according to different decision rules. Figure 3 compares the

thresholds from our baseline model with those of a standard statistical test (two tail, 5% significance level).<sup>11</sup> For  $n = 527$ , the figure also plots the posterior estimate of expected NIMB, calculated using the parameter values reported in Table 1. Figure 3 shows that the point estimate of expected NIMB lies between the upper and lower statistical thresholds, indicating that the point estimate is not significant at a 5% significance level, implying that the technology should not be adopted. This conclusion is consistent with our decision rule, since the point lies within the continuation region defined by our thresholds. However, although both approaches advise non-adoption, the statistical approach has nothing to say about whether sampling should be continued or abandoned whereas, with our thresholds, the DM's optimal strategy is to continue to research by paying for another realisation of  $W$ .

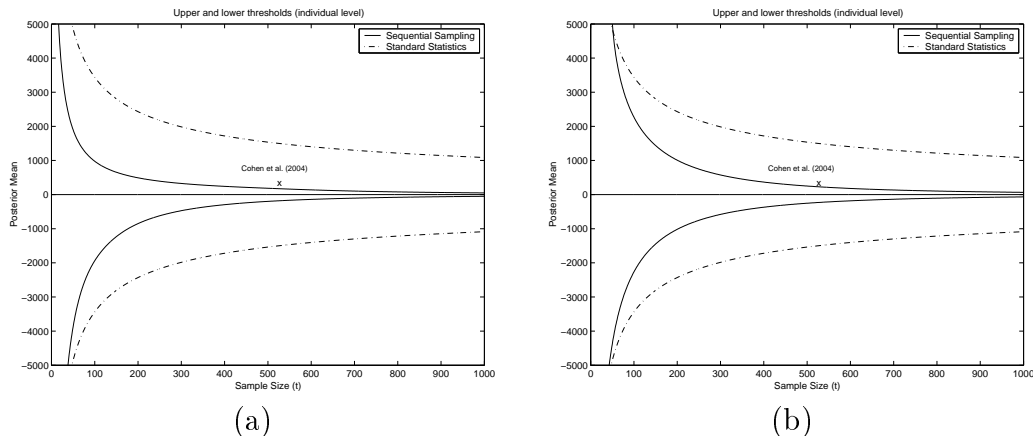
Figure 4(a) shows that, if the sampling cost  $c$  is sufficiently high ( $c = 1300$ ), our thresholds are pushed closer together and it becomes optimal to adopt the technology immediately based on the estimate of NIMB from Cohen et al. (the point estimate lies above our adoption threshold). Figure 4(b) shows a similar result for a sufficiently small value of  $P$ ,  $P = 100$ , which pushes the thresholds closer together because it reduces  $\sigma_X^2$ . Since neither parameter change affects the statistical thresholds, our decision rule is no longer consistent with that of the simple statistical approach.

According to the irrelevance of inference (IoI) approach, the decision concerning adoption should be separated from that concerning whether to continue with research. The former should be based on the point estimate, the latter on the comparison between the expected value of perfect information (EVPI) and the cost of obtaining that information. Since the point estimate of NIMB is positive for the SIRIUS study, IoI advises adoption of the new technology and a comparison of the costs of obtaining more information with the EVPI to indicate whether more research should be carried out.<sup>12</sup> In our model, since the DM has the option to defer the adoption decision in order to carry out further research, these decisions are not separated: above the upper threshold, it is optimal to adopt and not carry out research; in the continuation region, it is optimal to carry out research (and not adopt); below the lower threshold, it is optimal to abandon the project. Our result is therefore not consistent with the IoI result.

---

<sup>11</sup>At the individual level, the general form of the thresholds for a standard statistical test would be  $I/P \pm 1.96(\sigma_W/\sqrt{t})$ . In our case study,  $I = 0$ .

<sup>12</sup>In principle, it will be possible to calculate the expected value of perfect information, but to do so would require information on the population likely to benefit from the new technology.



**Figure 4:** Effect of: (a)  $c = \$1300$  and (b)  $P = 100$  on optimal thresholds, showing adoption to be optimal at Cohen et al. (2004) point estimate.

## 4 Discussion

Our case study shows how a sequential sampling approach to HTA can yield clear, dynamic, policy rules based on maximising the expected net incremental monetary benefit of a health technology assessment project, accounting for the costs associated with conducting research. Whereas the simple statistical approach relies on only the standard error and a chosen significance level to define the individual thresholds of statistical significance, our approach sees the HTA project as a dynamic process with economic dimensions - the marginal cost of sampling  $c$ , the population size  $P$  and the discount rate  $\rho$  - which vary across different technology-adoption decisions and trial contexts, leading to project-specific adoption and abandonment thresholds. The model's dynamic nature permits its use for analysing adoption decisions which can be deferred while further research is carried out, in order to narrow the uncertainty surrounding the DM's beliefs about mean NIMB. It thus extends the irrelevance of inference approach of Claxton (1999), which assumed that treatment decisions could not be deferred. In doing so, it shows that, for positive values of posterior mean NIMB lying in the continuation region, it is optimal for the DM not to adopt the technology but to carry out further research. The model's results are applicable to the population of  $P$  patients who are to be treated should adoption occur. The decision to treat or to conduct research for other groups of patients, such as those with different characteristics, or a wider patient population (including those to be treated after  $\bar{s}$ , the time at which the DM has the opportunity to revise his decision), should be separately addressed.


The sensitivity of the results with respect to changes in  $c$  and  $P$  has

been discussed in section 3.1. Given the comparatively rapid pace of accumulation of observations in the case study (904 per year), the impact of the discount rate  $\rho$  on the optimal thresholds is almost negligible. In general, however, a greater level of impatience tends to restrict the continuation region, by forcing the upper and lower thresholds together. If  $\rho$  equals infinity, that is, the decision to adopt the technology cannot be deferred, Eq. (3) becomes  $\max\{(y_t/t) - I, -c dt, 0\}$ , implying that adoption is optimal if the posterior mean minus  $I$  is positive. This is consistent with the result of the IoI approach.

The sensitivity of our result to changes in the value of  $P$  - as shown in Figure 4(b) - is an important one: while large differences in the cost of sampling  $c$  are more likely to exist for different technologies than for different studies of the same technology, the size of  $P$  may vary substantially in different contexts, even for the same technology. The implication is that, *ceteris paribus*, the expected sample size required to make the adoption/abandonment decision is lower the lower is  $P$ . As a result, it may be optimal to adopt a new technology earlier in contexts where a smaller population is involved, or the decision can be rapidly reversed on the basis of the new evidence. The value of  $P$  may be further influenced by the characteristics of the DM. For example, for adoption decisions being made at a national level, one would expect a larger number of subjects to be sampled before the adoption/abandonment decision is made (because  $P$  is high) whereas, at a local level, one would expect fewer subjects to be sampled (because  $P$  is lower). According to our model, the dynamically optimal decision rules are different for these two contexts, since they are functions of  $P$ : in the jargon of the real options literature,  $P$  is directly related to the degree of irreversibility of the decision.

Another form of irreversibility arises when the DM must pay a sunk cost  $I$  to adopt the technology. Although our case study sets  $I$  equal to zero, it is straightforward to consider the case when  $I > 0$ . Figure 5 compares the individual-level thresholds for the baseline version of the case study (solid lines) with the case when  $I = \$100000$  (dashed lines, the range for  $t$  is restricted to (100, 2000)). The inclusion of the sunk cost means the individual-level thresholds shift upwards, converging to  $I/P$  at the terminal time.

Although the size of the population to be treated in case of adoption, before the decision can be revised, has yet to be formally recognized in the HTA literature, it has been implicitly recognized in some health policy documents. For instance, the recommendations released by the UK National Institute for Health and Clinical Excellence (NICE) typically include an indication of the timing of review of the guidance. In the guidance for the use of Drug Eluting Stents issued in July 2008, it is stated that ‘the guidance on this technology will be considered for review in April 2009’ and that ‘This



**Figure 5:** Comparison of thresholds for the baseline case with the case when  $I = 100000$  (range for  $t$  restricted to (100,2000)).

decision will be taken in the light of information gathered by the Institute, and in consultation with consultees and commentators' (National Institute for Health and Clinical Excellence, 2008, page 40). Such documents can then become the basis for an estimate of  $P$ , given the prevalence and incidence of the disease in the population.

Aside from the absence of a closed form solution to the problem, necessitating reliance on numerical simulations, a possible limitation of our model concerns the arrival rate of new sample information: our assumption that the DM can observe this information in a sequential, deterministic manner, one pair of observations at a time, might not always be appropriate. For example, a DM carrying out a single trial might have more control over, and/or knowledge about, the rate of arrival of new information than one assimilating the results of published trials reporting in different jurisdictions. In the latter case, it could be more appropriate to assume a stochastic arrival process for new information.

## 5 Conclusion

The present paper aims to bridge the gap that exists between the absence of key economic variables from existing statistical approaches to HTA and the absence of important statistical dimensions from the real option approach to



HTA, in situations when an adoption decision can be deferred. We interpret the health technology assessment as a project with uncertain returns, which depends on clinical (effectiveness) and economic (cost) variables. Our results highlight the relevance of variables specific to each adoption decision, such as the size of the population involved, the cost of carrying out research and the degree of (im)patience of the decision maker, which are typically ignored in the statistical approach. We are also able to define a dynamically optimal policy for the abandonment of research, something which is typically overlooked in the health economics and statistical literature.

The costs of moving away from strictly statistical criteria mainly lie in the loss of standardization of the decision rules. However, we show that standardization is, in general, inconsistent with the dynamic consistency of an economic decision which seeks to maximise expected NIMB minus research, sunk and discounting costs.

## Acknowledgements

We thank Marco Realdon, Peter Spencer, Peter Simmons and participants in the Department of Economics and Related Studies seminar series and the 2010 Italian Health Economics Association (AIES) 15th Annual Conference for comments. Analysis is carried out using Matlab R2009b (v. 7.9.0.529). The authors take full responsibility for any errors.

## References

- Bellman, R. E. (1957). *Dynamic Programming*. Princeton University Press, Princeton, New Jersey, First edition.
- Chalkidou, K., Lord, J., Fischer, A., and Littlejohns, P. (2008). Evidence-based decision making: when should we wait for more information? *Health Affairs*, 27(6):1642–1653.
- Chernoff, H. (1961). Sequential tests for the mean of a normal distribution. Proceedings of the fourth Berkeley Symposium on Mathematical Statistics and Probability. pages 79–91.
- Chernoff, H. (1967). Sequential models for clinical trials. Proceedings of the fifth Berkeley Symposium on Mathematical Statistics and Probability. pages 805–812.
- Chernoff, H. (1972). *Sequential analysis and optimal design*, volume 8. Siam, Philadelphia, first edition.
- Chernoff, H. and Ray, S. N. (1965). A Bayes sequential sampling inspection plan. *Ann. Math. Statist.*, 36(5):1387–1407.
- Chick, S. E. and Gans, N. (2009). Economic analysis of simulation selection problems. *Management Science*, 55(3):421–437.

- Claxton, K. (1999). The irrelevance of inference: a decision-making approach to the stochastic evaluation of health care technologies. *Journal of Health Economics*, 18(3):341–364.
- Cohen, D. J., Bakhai, A., Shi, C., Githiora, L., Lavelle, T., Berezin, R., and others (2004). Cost-effectiveness of sirolimus-eluting stents for treatment of complex coronary stenoses. *Circulation*, 110:508–514.
- Dixit, A. and Pindyck, R. (1995). The options approach to capital investment. *Harvard Business Review*, 73(3):105–15.
- Driffield, T. and Smith, P. C. (2007). A real options approach to watchful waiting: theory and an illustration. *Medical Decision Making*, 27:178–188.
- Eckermann, S. and Willan, A. R. (2008). The option value of delay in health technology assessment. *Medical Decision Making*, 28(3):300–305.
- Griffin, S., Claxton, K., Palmer, S., and Sculpher, M. (2010). Dangerous omissions: the consequences of ignoring decision uncertainty. *Health Economics*, doi: 10.1002/hec.1586.
- Joseph, L. and Belisle, P. (1997). Bayesian sample size determination for normal means and differences between normal means. *Journal of the Royal Statistical Society, Series D*, 64(2):209–226.
- Lai, T. and Lim, T. (2005). Optimal stopping for brownian motion with applications to sequential analysis and option pricing. *Journal of Statistical Planning and Inference*, 130:21–47.
- Levaggi, R. and Moretto, M. (2008). Investment in hospital care technology under different purchasing rules: a real option approach. *Bulletin of Economic Research*, 60(2):159–181.
- National Institute for Health and Clinical Excellence (2008). Drug-eluting stents for the treatment of coronary artery disease. Part review of NICE technology appraisal guide 71.
- Palmer, S. and Smith, P. C. (2000). Incorporating option values into the economic evaluation of health care technologies. *Journal of Health Economics*, 19:755–766.
- Pertile, P. (2009a). An extension to the real option approach to the evaluation of health care technologies: the case of positron emission tomography. *International Journal of Health Care Finance and Economics*, 9(3):317–332.
- Pertile, P. (2009b). The timing of adoption of positron emission tomography: a real options approach. *Health Care Management Science*, 12(3):217–227.