UNIVERSITY of York

Discussion Papers in Economics

No. 14/04

Don't Stop 'Til You Get Enough: a quickest detection approach to HTA

Daniele Bregantini

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

Don't Stop 'Til You Get Enough: a quickest detection approach to HTA

Daniele Bregantini*,

March 5, 2014

Abstract. Within the context of the value of information approach we compare static versus quickest detection rules for research design in health care technology assessment (HTA). We show for a research design that the optimal decision rule cannot be correctly predicted at the start of the trial. We make use of the sequential value of information (S-VoI) decision making model for HTA under uncertainty to show that the static value of information approach leads to lower expected benefit and poses costs, both in terms of resources and forgone health gains, on the health care system.

Keywords: Optimal stopping, HTA, Bayes, Value of Information
JEL codes: C00, C11, C12,C15

1. Introduction

Regarding health care technology assessment, health care systems are faced with the following questions: (i) should a health care technology be reimbursed in light of current evidence (ii) is additional evidence required to support the use of the technology, and if this is the case, what type and how much research is required ? In England and Wales, the National Institute for Clinical Excellence (NICE) issues guidance to the National Health Service (NHS) for the need of further research and gives recommendation on adoption. It has made adoption conditional on further research and the production of further evidence (Conti and Claxton, 2009).

When uncertainty about the net benefits is present there is a positive probability of making an incorrect decision. The expected value of information developed by Raiffa and Schlaifer (See Pratt et al. (1995)) and later applied to the case of health technology assessment (HTA) and clinical research design by Claxton and Posnett (1996) and Claxton (1999) can be used to quantify the expected opportunity loss associated with this uncertainty. When the expected opportunity loss is less than the cost of a new study the information is deemed to be sufficient and a decision can be made. When this static decision making approach is implemented to clinical research design it suggests to compute an ex-ante optimal (fixed) sample size deemed to be sufficient for the purposes of decision making. Claxton (1999) put forward the idea that inference is irrelevant to decision making and suggested that

*Department of Economics and related studies, University of York, YO10 5DD, York. e-mail: db566@york.ac.uk

the question of whether more evidence is needed should be determined by the value of information framework developed by Raiffa and Schaifer (See Pratt et al. (1995)).

Recently, William and M.Pinto (2005) suggested a method for computing the ex-ante sample size n^* for a clinical trial that maximises the difference between the cost of a trial and the expected value of the results using the incremental net benefit as the main outcome for the trial. Willan and Kowgier (2008) developed the model of William and M.Pinto (2005) to a multistage adaptive-design involving an early termination rule based on the expected net gain from the trial computed for each stage j. If the EVSI in the next stage $j+1$ is less than the total cost at $j+1$ then the trial terminates at the end of the jth stage and the decision rule can be applied. Although it is theoretically possible to construct a purely multistage model that jointly determines the value of n_j^* for all j maximizing the expected net gain, due to its complexity Willan and Kowgier (2008) suggest to proceed in two-stages steps where at each stage j the (ex-ante) two stage calculation is repeated and the maximisation process is repeated at each j. Another early termination approach is found in Berry and Ho (2003) who take the point of view of a pharmaceutical company that wishes to maximise profits and uses a one-sided decision-theoretic approach in order to determine if experimentation of a newly developed drug should be stopped early in case of negative evidence.

In recent years the literature has seen the application of the real option approach to investment decisions in health technology assessment (Palmer and Smith, 2000). This literature aimed at incorporating the dynamic nature of the decision process and considers the role of flexibility and irreversibility of investment. More recently Pertile et al. (2013) solved the dynamic problem of the economic valuation of a new health technology in the content of the optimal stopping under sequential sampling literature developed by Chernoff (1961). In their paper, in the situation where the decision maker can defer decisions and there is a certain degree of irreversibility of investment, optimal rules for technology adoption, research abandonment and continuation, are developed as functions of sample size. In this framework the decision maker sequentially observes (for example during a trial) the net incremental monetary benefit of the technology for a certain number of patients. Each observation carries a sampling cost and the DM after observing each observation updates its beliefs according to Bayes' theorem and takes a decision: either invest/abandon the project or continue monitoring the results. After each observation the posterior estimate gains precision and at the optimal stopping time the value for the net incremental monetary benefit can be inferred with enough confidence allowing for a decision of either investment or abandonment.

Forster and Pertile (2012) discuss the use of real options analysis as a way to view adoption, treatment and research decisions as a single economic project and argue that the dynamic approach to HTA can provide efficiency gains in resource

allocation. However, presently the real option approach has not been implemented in any systematic way (Meltzer and Smith, 2012).

In this paper we present a comparison between the traditional value of information framework as found in Claxton (1999) and a dynamic decision theoretic approach. We adopt a sequential value of information (S-VoI) rule (see Bregantini and Thjissen (2014)) as this helps the user to reach a decision between two hypotheses after a minimal number of experiments. This method, in contrast to Pertile et al. (2013), does not involve an estimation problem for the unknown net incremental mean benefit but specifies some bounds at which a decision can be taken. When the cumulative net incremental mean benefit hits one of the bounds the observed sample size is sufficient and the decision, either for investment or for abandonment, can be undertaken with minimal error. Additionally, Pertile et al. (2013) require a sample size specified at priori to the start of the trial, a requirement that is avoided in the S-VoI model.

In contrast to the static approach, the S-VoI model does not force a decision after observing n observations no matter the information contained in the observed sample. In particular, in the case of a fixed sample, the size can be dangerously small or redundantly large for making a reasonably good inference on which of the two hypotheses is true.

With sequential testing on the other hand, no observations are wasted. In fact, as soon as we can declare that one of the two hypotheses is true with reasonable certainty, we stop taking observations. For this reason, in the presence of sampling costs, it is clear that sequential testing is a method of testing that is less costly on average than its competitor fixed sample size testing (Poor and Hadjiliadis, 2009)).

Consistent with Claxton (1999), and in contrast to traditional sample size calculations for randomised clinical trials based on type I and type II probabilities rules that do not account for the monetary cost or making the wrong decisions, the S-VoI focuses on expected payoff and aims at maximising health benefits with minimum error probability. In the sequential setting the implication for the irrelevance of inference suggested in Claxton (1999) is that fixing the sample size ex-ante is not optimal and, as with rules based on type I and type II error minimisation, can lead to choices that do not maximise health benefits with minimum error probability.

The paper begins by outlining the static value of information approach and the sequential value of information. These are then followed by an illustrative example of research design based on simulations for the two models. Finally we report and contrast, in terms of monetary value of gained health benefit, the expected research design outcome for the Value of Information (VoI) approach to HTA found in Claxton (1999) and Claxton and Posnett (1996) and the S-VoI.

2. Static decision rules

We begin by introducing the main tools of the Value of Information approach as found in Raiffa and Shlaifer (see Pratt et al. (1995)) and adapted to the case of health technology assessment by Claxton and Posnett (1996) and Claxton (1999).

2.1. **Expected Value of Perfect Information.** Claxton (1999) propose to use the EVPI as a way to address how decision makers (DM) should interpret the results of probabilistic modelling and to address the question of whether enough evidence has been gathered. This approach mirrors the sequential nature of decision making: making an initial decision; deciding to gather evidence; revising decisions in the light of this new information; and again considering whether more information is required. It also ensures that the type of information acquired through research is driven by the objectives of the health care system and is valued in a way which is consistent with the budget constraint on service provision. In this framework, the expected cost of uncertainty is determined jointly by the probability that a decision based on current evidence will be wrong and the consequences of a wrong decision.

We take λ to be the cost-effectiveness threshold[2] and ν_0 to be the prior mean incremental net benefits of the technology. The prior standard deviation of the incremental net benefit is σ_0 and, in order to understand the possible loss that can occur when making the wrong decision, the break-even point ν_b (the value at witch the new technology would be cost-effective) is set to zero and the standardised distance is given by

$$D_0 = \frac{\mid \nu_0 - \nu_b \mid}{\sigma_0}.$$

The value of D_0 gives the standardised distance between the prior mean and the break-even point. The EVPI can be written as

$$\text{EVPI} = \lambda \sigma_0 L(D_0) \tag{1}$$

The term $L(D_0)$ indicate the normal loss integral[3] for D_0. The unit normal loss integral $L(D_0)$ represents the probability of the decision error when the incremental net benefit has a mean D_0 and $\sigma_0 = 1$. The slope of the loss function is the cost-effectiveness threshold λ and gives the monetary value placed on opportunity losses when they incurred. The probability of the decision error is determined by the distance of the prior mean incremental net benefit.

[2]This is the threshold defined by the health care system: below the threshold λ a health-care technology is deemed cost-effective, above λ the health care technology is not cost-effective and not eligible for reimbursement by the health care system.

[3]See Appendix A.

The analytic solution for the EVPI in (1) relies on the assumption that the net benefit is normally distributed. It often the case, however, for cost- effectiveness model to synthesise parameters that not only have different distributions (i.e beta distributions, logistic functions, etc) but that also come from different sources. In such cases the non-normality found in the data is a problem as highlighted in Thompson and Nixon (2005) who report that distributional assumption about data influenced the determination of cost-effectiveness. When the assumption of normality does not hold the EVPI is computed using simulation methods (See Ades et al. (2004)).

2.1.1. *Expected Value of Sample Information.* The value of information analysis can be extended in order to find the expected value of sample information for particular research design (Ades et al., 2004). In order to establish if the conditions for further research are present and to identify efficient research design there is the need to also consider the expected costs of sample information. The expected value of sample information was introduced as a decision tool for clinical trial design by Claxton and Posnett (1996) and Ades et al. (2004).

The EVPI places an upper bound on returns to further research and provides a necessary but not sufficient condition for conducting further research. If the value of EVPI exceeds the cost of further research it might be worthwhile to gather more information about the problem as a whole or on selected parameters. However, in order to establish if further research will be worthwhile (i.e. net benefits of research are positive) and to identify efficient research design there is the need to consider the marginal benefits and marginal costs of sample information.

2.1.2. *Technically efficient research design.* The EVSI can be calculated for a particular sample size from the prior information and the estimate of the sample variance (σ^2/n). The EVSI is then

$$\text{EVSI}_{(n)} = \lambda\sqrt{V(n)}\sigma_0 L(D_{0(n)}) \tag{2}$$

where λ is the cost effectiveness threshold, $D_{0(n)}$ is the standardised distance as a function of the sample variance, V is a function of the sample size and the prior variance σ_0 (See Claxton and Posnett (1996)).

It should be noted[4] that in equation (2) as $n \to \infty$, $V(n)$ goes to 1 and the EVSI approaches the EVPI. The EVPI provides an upper bound to the value of additional research as it identifies the maximum benefit that could be provided by additional information. The fact that the EVSI and the EVPI coincide when the sample n is infinite precisely indicates that the EVPI is the maximum benefit that we can get by observing an infinite trial.

[4]For details see Appendix A.

The expected net benefit of sampling is the difference between the total benefit and the total variable cost for a particular sample size:

$$\text{ENBS}_{(n)} = \text{EVSI}_{(n)} - C_{s(n)} \tag{3}$$

The subscript n indicates the step in the trial and the cost $C_{s(n)}$ is the total trial cost at step n. The optimal sample size n^* maximises $\text{ENBS}_{(n)}$. The optimal value of n is given by the following condition

$$\frac{\partial \text{EVSI}}{\partial n} = C_n \tag{4}$$

As for the EVPI, simulation methods have been proposed in order to deal with the non-linearities and non-normal distribution of the net benefit (See Ades et al. (2004). However, the solution to the decision problem in the value of information approach, as noted in equation (4), remains a static one: the maximisation of the EVSI is computed ex-ante, it computes a single value for the optimal sample size n and does not take into account any information that arises during the trial, in effect making the choice of n reasonable before the trail actually starts, but as we show in section 5.1, suboptimal at any point $n > 0$.

2.1.3. *Cost.* The EVSI does not account for cost different than those directly associated to running the trial. There are no health losses connected to delaying the decision and not treating patients with a more effective technology. The issue of forgone value of information has been introduced by Griffin et al. (2011), however, the value of information remained a static decision framework. A dynamic approach to research design has been undertaken by Claxton and Thompson (2001) where the approach found in Claxton and Posnett (1996) and Claxton (1999) are generalised to the analysis of a sequential clinical decision problem.

Claxton (1999) advocates that deciding which alternatives should be chosen, given existing information, and deciding whether more information should be required are two simultaneous but conceptually separate steps. The VoI provides a way to distinguish between these two concepts.

3. Sequential Value of Information (S-VoI)

In a sequential value of information (S-VoI) decision making model developed in Bregantini and Thjissen (2014), the DM is faced with a two-sided decision: either invest in the health care technology or abandon the health care project. The S-VoI is a quickest detection model that allows to test for the hypothesis with the minimum number of observation required, maximising payoff with minimum error probability.

In contrast to models that propose simulation based solutions, the S-VoI model uses continuous time mathematics that allows to fully understand the modelling

results. The use of continuous over discrete time modelling enables to access a mathematical toolbox that provides analytical solutions. In continuous time it is possible to apply the central limit theorem even if the sample size is not infinite. During the trial, the decision maker observes the net benefit for each patient as a sequence of outcomes. The net benefit over a small time interval is given by μdt. The decision maker however, cannot clearly observe the net benefit due to a noise element $\sqrt{\sigma}$. The evolution of the sequence is described by the following tree diagram

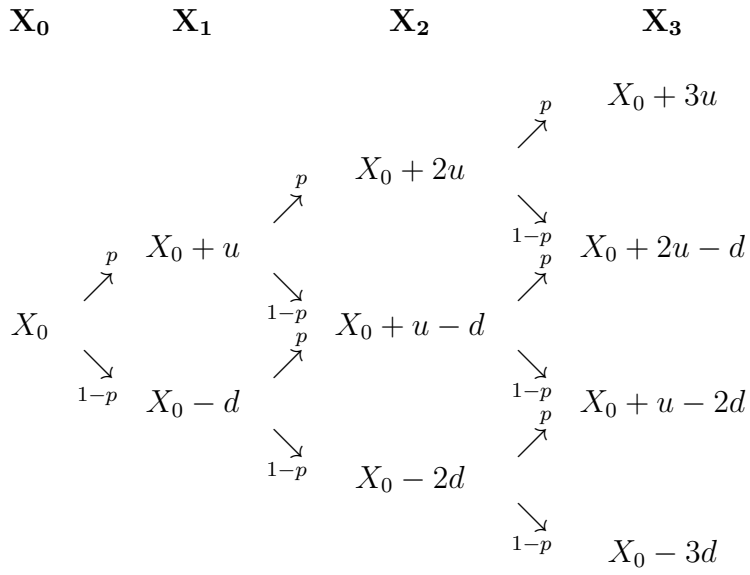

In the above tree the initial value X_0 can increase by a factor $u = \theta\mu dt + \sigma\sqrt{dt}$ or decrease by a factor $d = \theta\mu dt - \sigma\sqrt{dt}$. with probability $p = 0.5$. The term θ can be equal to 1 or 0 and will be used below for hypotheses testing. The sequence of random variable is

$$X_i = \begin{cases} \theta\mu dt + \sqrt{dt} & \text{with } pr = 1/2 \\ \theta\mu dt - \sqrt{dt} & \text{with } pr = 1/2 \end{cases}$$

At each point in the sequence the value of $X_{(n)}$ is given by $X_{(n)} = \sum_0^n X_i$ and as the time interval between steps decreases we denote $X_t = \lim_{n\to\infty} X_n(t)$ where the limit is understood to be in distribution and $n \to \infty$ implies $dt \downarrow 0$. According to the CLT, the limit X_t exists in distribution and is given by

$$X_t \sim N(\theta\mu t, \sigma^2 t)$$

implying that in the continuous time limit the process X_t follows the arithmetic Brownian motion

$$X_t = \theta\mu t + \sigma W_t. \tag{5}$$

Equation (5) describes the net benefit as a continuous time sequence of random variables. In the equation, θ represents the hypothesis that the health care technology is effective and provides the claimed net benefit $\mu > 0$. With $\theta = 1$ the technology is effective and when $\theta = 0$ the technology is no better than standard care (in which case $\theta\mu t = 0$).

We consider the case where the decision maker is interested in testing the claim from a manufacturer that seeks reimbursement for a newly developed health care technology that should provide excessive benefit μ. The claim could also be related to the minimum effectiveness required for cost-effectiveness (i.e. μ such that net incremental mean benefit (NIMB) is positive) as part a cost-effectiveness trial by a health care manufacturer. Such test will allow the manufacturer to provide stronger evidence in support for government reimbursement.

We consider research design for a project that has an irreversible fixed cost I and net present value of adoption given as function of the posterior probability π. The investment payoff is $F_I(\pi)$ and the abandonment payoff $F_A(\pi)$. At each point in the sequence, Bayes rule allow to compute a posteriori probability process π_t as a function of (i) the prior probability assigned to the likelihood of the technology being more effective than standard care and (ii) a likelihood process $\Lambda_t(X_t)$ as a function of the trial sequence X_t (For details see Bregantini and Thjissen (2014)).

In this way the posterior probability π_t that the new technology is more effective than standard care in continuously updated via Bayes rule. In order to reflect the possibility of investment when the technology is not effective (i.e. a type I error) and the possibility of abandoning the project when the technology is better than standard care (i.e. type II error) the payoffs are specified as follows:

$$F_I(\pi) = \pi P_1 - (1-\pi)P_0 - I \qquad P_1 > 0,\ P_0 > 0 \tag{6}$$

$$F_A(\pi) = -\pi P_1 \qquad P_1 > 0 \tag{7}$$

The term P_1 represents the monetary benefit to the healthcare system of investment in the new healthcare technology conditional on $\theta = 1$ and $-P_0$ represents the monetary loss of new healthcare technology conditional on $\theta = 0$.

Subject to sampling costs c and discount rate r, the problem is to find an optimal stopping time τ^* at which a decision can be taken, payoffs maximised and the value of waiting for an additional sample is zero. At the optimal stopping time τ^* the likelihood process $\Lambda_t(X_t)$ hits either the upper investment bound Λ_I or the lower abandonment bound Λ_A. The likelihood process provides evidence for hypothesis $H_1 : \theta = 1$ or $H_0 : \theta = 0$. At τ^* the DM stops sampling and an optimal

Table 1. Simulation

Simulation	σ	mode_{τ^*}	μ_{τ^*}	σ_{τ^*}	$\min_{\tau^*}$	$\max_{\tau^*}$
1	$\sigma = 0.1$	192	314	197.7	10	2619
2	$\sigma = 0.15$	195	435	321	34	3925
3	$\sigma = 2$	175	548	478	36	6803
4	$\sigma = 2.5$	179	667	660	25	8181

decision can be taken, either for investment with payoff F_I (i.e. supporting H_1) or abandonment with negative payoff F_A (i.e. supporting H_0).

In the optimal stopping model, the decision to invest/abandon or continue research, in contrast to the VoI approach, is subject to the information generated by the random variable X_t. As the trial continues, information about the net benefit X increases, and consequently uncertainty about the true net benefit decreases. The optimal decision is taken at the time τ^*, when the value of waiting for a further sample is zero.

4. Quickest detection decision rules

The S-VoI model specifies investment and abandonment bounds that aim at maximising payoff. The following simulation study shows hitting times τ^* (times at which it is optimal to make a decision) for $100,000$ simulated sample path for some hypothetical statistical values to be tested. The table 1 below reports statistics for a simulation study based on the following values: $P_1 = 130$, $P_0 = 60$, $r = 15\%$, $I = 40$, $\mu = 0.15$. The value for σ is increased in small steps for each simulation in order to show the consequence of different degrees of uncertainty on the distribution of hitting times τ^*.

Figure (1) below shows the distribution of τ^* for different values of μ. As it can be noted in Figure (1a) the distribution of τ^* is centred around the mode value of $\tau^* = 192$ with few events that occur after the $\tau^* = 1250$ region. Figure (1c) τ^* has a much thicker tail after $\tau^* = 1250$, indicating that there is a greater number of τ^* events after this value than in the previous model. For $\sigma = 0.15$ the mean $\mu_{\tau^*} = 435$ and $\sigma_{\tau^*} = 321$ with a minimum hitting time of 36 and a maximum hitting time of 3803. A substantial increase from the simulation 1.

In simulation 3, as σ increases, the statistical values for $\mu_{\tau^*}, \sigma_{\tau^*}, \min_{\tau^*}, \max_{\tau^*}$ increase. It can be noted in Figure 1b that the number of events occurring after $\tau^* = 1200$ is much greater than in simulation 1 and 2 with some extreme events occurring well inside the far right tail of the distribution with a maximum of $\tau^* = 6803$. Figure 1d shows the distribution hitting time τ^* for simulation 4. Of the simulated models, this is the most extreme case with $\sigma = 0.25$. The mean is $\mu_{\tau^*} = 667$ with $\sigma_{\tau^*} = 660$ with a the maximum $\tau^* = 8181$ with most of the events occurring before $\tau = 4000$.

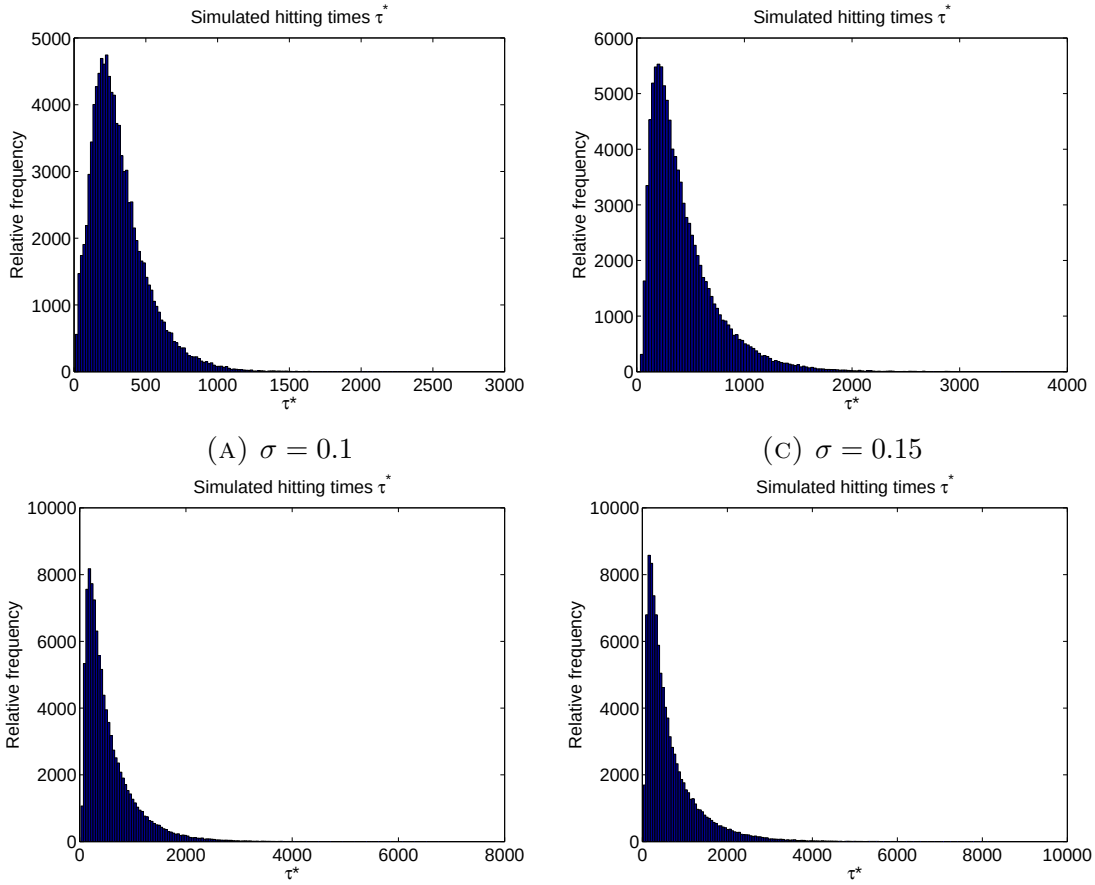

(A) $\sigma = 0.1$ (C) $\sigma = 0.15$

(B) $\sigma = 0.2$ (D) $\sigma = 0.25$

FIGURE 1. Simulated τ^* for different values of σ

The above results suggests that while the mode for the hitting times do not vary much, distribution varies considerably and the dispersion for hitting times (or decision times) varies considerably given different levels of uncertainty. A consequence of this, for models that use the total cost of new research (or stage of a trial) as a rule to determined if a new study should be undertaken (e.g. EVPI), is that when uncertainty is high it is difficult to correctly asses the cost of a new trial due to the uncertainty surrounding the optimal stopping time τ^*. For these ex-ante models, the comparison between the net benefit of new research and its costs should account for uncertainty surrounding τ^*.

Additionally, another important consequence of uncertainty is that by having a rule that specifies ex-ante a fixed sample size for a trial, decisions might be taken at points where information is not sufficient or alternatively decision might be taken later than necessary with corresponding costs for the health care system. The cost

TABLE 2. Cost-effectivness of standard vs robot-assisted laparoscopi prostaectomy

Parameter	**Description**	**Source**	**Value**
$E_1 - E_0$	Incremental QALY gain	Close *et al*	0.08 QALY
$C_1 - C_0$	Incremental cost	Close *et al*	£1412
σ	Std. deviation	Close *et al*	£1071
μ	Incremental QALY gain	Set as $NIMB > 0$	£1413
p	Prior	Assumed	0.5
r	Discount rate	Close *et al*	3.5%
c	Cost of sampling	Assumed	£50
I	Investment	Close *et al*	0
λ	QALY value	Close *et al*	£30000
n	Number of patients	Close *et al*	10000

of employing such ex-ante rules, for the specific case of EVSI, is discussed in the next section.

5. EVSI vs quickest detection rules

In this section we aim at showing, with a simple illustrative example, the difference in a research design application between the two approaches. We simulate for a number of cases the stopping time produced by the dynamic bayesian model and compare this to the optimal sample size given by the static value of information approach. For illustration purposes data is taken from Close et al. (2013)'s study of cost-effectiveness of standard vs robot-assisted laparoscopi prostaectomy. Data is shown in Table 2.

The EVSI predicts that the optimal sample size is $n^* = 91$. Figure (2) shows the relative frequency of hitting time τ^* for 10,000 simulated sample paths with sampling fixed at one new patient added per day (i.e. we assume 365 patients per year). When comparing the simulated hitting times with a static approach it can be noted that this last is likely to overestimate the sample size, nonetheless it can also underestimate the sample size for a good number of cases. In the analysis that follows τ and n represent the same values[5]. In the case of the S-VoI the mean hitting time τ^* is $\tau^* = 216$ with a mode of $\tau^* = 60$.

The EVPI describes the advantage of full information over partial information. The EVSI involves computing the opportunity loss of making a decision based on

[5]Since the same sampling scale is used we use τ to denote the dynamic optimal stopping model and retain n for the traditional EVSI approach. However, τ and n are equivalent and represent the same values.

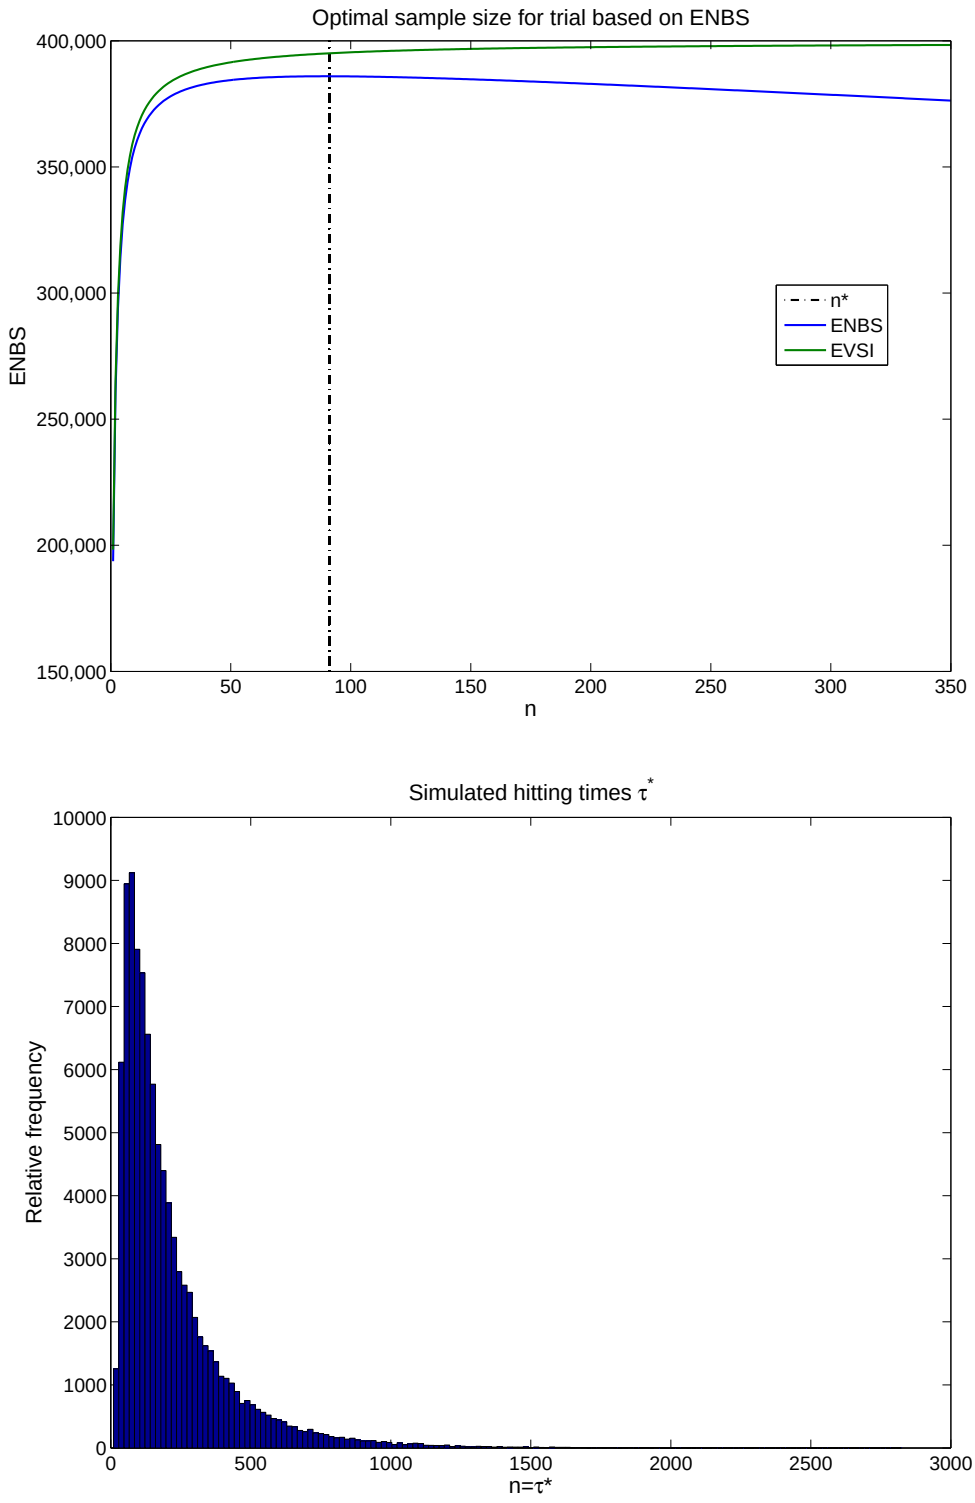

FIGURE 2. EVSI optimal sample size n^* and simulated τ^*

TABLE 3. Expected payoffs

Decision model	**Optimal sample size**	**Expected payoff £**
EVSI	Fixed (n^*)	207.7
S-VoI	Flexible (τ^*)	564.3
Health gain from S-VoI		356.6

prior information solely. In sharp contrast, in the S-VoI each decision is based on the appropriate information set generated by the random variable X_t.

The S-VoI approach provides a way to undertake the quickest decision that minimises expected opportunity loss both in terms of forgone health benefits to patients or resources allocated to the trial.

5.1. **Cost of non optimal decisions.**

For the case study above, we compare the sequential-VoI and the EVSI payoffs. This gives an estimate of the costs involved in taking decisions based on a fixed, deterministic rule versus a sequential flexible rule.

Table (3) shows the expected value obtained from 100,000 simulated paths for a trial sequence X_t based on the Close et al. (2013) prostaectomy study reported above. The expected payoff is obtained by taking the maximum payoff value at the optimal decision point in the sequence X_t. This point is given by the optimal sample size (n^*) for the EVSI and at the optimal stopping time (τ^*) for the S-VoI. The expected payoff at the optimal sample size predicted by the EVSI provides a low value, indicating that (n^*) is suboptimal when compared to the S-VoI approach that instead selects the trial size (i.e. optimal stopping time τ^*) that maximises expected payoff. The S-VoI total expected gain for the health care system when compared to the EVSI approach is £356,600 per 1000 patients.

6. CONCLUSION

Within the context of the value of information approach we compare deterministic versus dynamic rules for research design in HTA. The value of information approach selects the optimal trial length based on the prior information available and it produces a decision rule that proves to be inefficient for a great majority of cases. The reason is to be found in that under uncertainty evidence is accumulated over time and the point at which sufficient information is reached is not known at the start of the trial.

We show that this optimal decision point is reached at a random time that is optimal under some payoff based rules. As this optimal stopping time cannot be predicted at the start of the trial the research design advocated by the EVSI is inefficient brings losses to the health care system.

7. Appendix A

The $EVSI(n) = \lambda\sqrt{V(n)}\sigma_0 L\left(D_n\right)$ depends on the following factors :

$$D_{0(n)} =\mid \nu_0 - \nu_b \mid / \left(\sigma_0\sqrt{V(n)}\right)$$

$L\left(D_{0(n)}\right)$ = unit normal loss integral for standardised distance $D_{0(n)}$. The unit loss integral is computed as

$$L\left(D_{0(n)}\right) = f\left(D_{0(n)}\right) - D_{0(n)}\left(1 - \phi\left(D_{0(n)}\right)\right)$$

where $f(\cdot)$ is the density function of a standard normal distribution and $\phi(\cdot)$ is the distribution function of the standard normal distribution.

The variance is given by:

$$\begin{aligned} V(n) &= \sigma_0^2/(\sigma_0^2 + \sigma/n) \\ \sigma^2 &= \text{population variance of } v \\ \sigma^2/n &= \text{sample variance of } v \end{aligned}$$

References

Ades, A. E.; Lu, G., and Claxton, K. Expected value of sample information calculations in medical decision making. *Medical decision making*, 24(2):207–227, 2004.

Berry, D.A. and Ho, C-H. One-sided sequential stopping boundaries for clinical trials: a decision-theoretic approach. *Biometrics*, 44:219–227, 2003.

Bregantini, D. and Thjissen, J. J. J. On a simple quickest detection rule for health-care technology assessment. *University of York working paper*, 2014.

Chernoff, H. *Sequential tests for the mean of a normal distribution.* Proc. 4th Berkley Symp. Mathematical Statistics and Probability, 1961.

Claxton, K. The irrelevance of inference: a decision-making approach to the stochastic evaluation of health care technologies. *Journal of health economics*, 18:341–364, 1999.

Claxton, K. and Posnett, J. An economic approach to clinical trial design and research priority setting. *Health economics*, 5:513–524, 1996.

Claxton, K. and Thompson, K. M. A dynamic programming approach to the efficient design of clinical trials. *Journal of health economics*, 20:797–822, 2001.

Close, A.; Robertson, C.; Rushton, S.; Shirley, M.; Vale, L.; Ramsay, C., and Pickard, R. Comparative cost-effectiveness of robot-assisted and standard laparoscopic prostatectomy as alternatives to open radical prostatectomy for treatment of men with localised prostate cancer. *European urology*, 64(3):361–369, 2013.

Conti, S. and Claxton, K. Dimensions of design space: a decision theoretic approach to optimal research design. *Medical decision making*, (29):643–660, 2009.

Forster, M and Pertile, P. Optimal decision rules for hta under uncertainty: A wider, dynamic perspective. *Health Economics*, pages n/a–n/a, 2012.

Griffin, S. C.; Claxton, K. P.; Palmer, J. S., and Schulpher, J. M. Dangerous omissions: the consequences of ignoring decision uncertainty. *Health economics*, 20:212–224, 2011.

Meltzer, D. O. and Smith, P. C. *Handbook of health economics*, volume 2, chapter vii, page 433. Elsevier, 2012.

Palmer, S. and Smith, P. C. Incorporating option value into the economic evaluation of health care technologies. *Journal of Health Economics*, 19:755–766, 2000.

Pertile, P.; Forster, M., and LaTorre, D. Optimal bayesian sequential sampling rules for the economic evaluation of health technologies. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, pages n/a–n/a, 2013.

Poor, V. H. and Hadjiliadis, O. *Quickest detection*. Cambridge University Press, 2009.

Pratt, W. J.; Raiffa, H., and Schlaifer, R. *Statistical decision theory*. MIT Press, 1995.

Thompson, S. G. and Nixon, R. M. How sensitive are cost-effectivness analyses to choice of parametric distributions ? *Medical decision making*, 4(25):416–423, 2005.

Willan, A. and Kowgier, M. Determining optimal sample sizes for multi-stage randomized clinical trials using value of information methods. *Clinical Trials*, 5:289–300, 2008.

William, A. R. and M.Pinto, E. The value of information and optimal clinical trial design. *Statistics in medicine*, 24:1791–1806, 2005.