

This is a repository copy of *A Bayesian change-point detection approach to the economic evaluation of risky projects : an application to health-care technology assessment*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/204246/>

Version: Published Version

Article:

Bregantini, Daniele, Schmitt, Laetitia Helene Marie and Thijssen, Jacco orcid.org/0000-0001-6207-5647 (2023) A Bayesian change-point detection approach to the economic evaluation of risky projects : an application to health-care technology assessment. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*. ISSN 1467-985X

<https://doi.org/10.1093/jrssa/qnad129>

Reuse


This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

A Bayesian change-point detection approach to the economic evaluation of risky projects: an application to healthcare technology assessment

Daniele Bregantini¹, Laetitia H.M. Schmitt² and Jacco J.J. Thijssen³ 

¹RTI Health Solutions, Manchester, UK

²Clarivate, Milano, Italy

³Department of Mathematics, University of York, York YO10 5DD, UK

Address for correspondence: Jacco J.J. Thijssen, Mathematics, University of York, York YO10 5DD, UK.

Email: jacco.thijssen@york.ac.uk

Abstract

We propose a Bayesian hypothesis testing framework that allows for the assessment of evidence collected during a clinical trial about the cost-effectiveness of a healthcare technology. The model exploits a Bayesian updating rule that makes the link between the evidence collected in clinical research and the expected payoffs of adoption to the healthcare system. The framework takes into account the cost of decision errors in the payoff function, allowing the decision maker to compute the cost of taking a decision when evidence is far from the optimal decision triggers. We show, using a real-world cost-effectiveness study based on clinical trial evidence, how rules derived from a sequential adaptive design approach can lead to quicker decisions when compared to the value of information decision framework. Our application shows that a sequential approach has the potential to lead to quicker decisions, higher payoffs, and better health outcomes.

Keywords: Bayesian statistics, economic evaluation, optimal stopping, sequential analysis, value of information

1 Introduction

Randomised clinical trials (RCTs) are traditionally considered to be the gold standard for determining the safety and efficacy of healthcare technologies (HCTs; Spiegelhalter et al., 2004) and their outcomes largely determine whether new HCTs are approved by regulatory agencies. However, in recent decades, as healthcare expenditure has increased considerably, there has been great pressure on healthcare systems to provide value-for-money. In England and Wales, the National Institute for Health and Care Excellence (NICE) has been giving recommendations on adoption of healthcare technologies on the basis of clinical evidence and cost-effectiveness and when technologies display large uncertainty over their expected cost-effectiveness, NICE has made approval conditional on further research and the production of further evidence (Claxton, 1999).

Cost-effectiveness of a given intervention has long been measured using incremental cost-effectiveness ratios, that is the ratio of the difference in costs and benefits of the HCT against its comparator, such as best current clinical practice. Whilst this measure is still widely used, health economists increasingly use net incremental health effects or equivalently, net incremental monetary benefit (NIMB), as metrics of payoff to the healthcare system, as they incorporate the

Received: March 31, 2023. Revised: October 17, 2023. Accepted: October 21, 2023

© The Royal Statistical Society 2023.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

opportunity cost imposed by funding the HCT under evaluation that results from reducing the resources available for funding other technologies.

The need for a rational principle-based approach to healthcare resource allocation has led to the use of decision-analytic tools that initially relied on the traditional approach of statistical inference when assessing the cost-effectiveness of new technologies. However, whilst statistical inference is of fundamental importance to understanding whether or not a new drug is likely to be beneficial to patients, it does not necessarily answer questions about *economic* payoffs. More importantly, it does not help the decision maker (DM) assess if the level of evidence gathered is sufficient to minimise the risk of taking an incorrect decision that could lead to a net health loss at the healthcare system level. Indeed whilst a statistically significant favourable treatment effect might suggest the HCT would provide sizeable health benefits to the healthcare system, the HCT could nevertheless have very uncertain payoffs once the magnitude of the health benefits displaced elsewhere in the healthcare system by funding this intervention is taken into account.

Uncertainty around expected payoffs to the healthcare system might suggest collecting more evidence is required, for instance by extending a trial. However, running a trial is costly and there has therefore been great interest in monitoring trials' accumulated data at planned time-points and undertaking interim analysis of outcomes to ensure continuing the data collection exercise is warranted. When evidence indicates a statistically significant favourable difference in outcomes, early termination means that the HCT can be exploited sooner and conversely, in the case of a statistically significant adverse difference stopping early involve saving resources (Jennison & Turnbull, 2000). Sequential methods, both frequentist and Bayesian, when compared to traditional statistical inference, typically need a smaller sample size, and hence are quicker and cheaper. Chevret (2012) investigates the international scientific production of Bayesian clinical trials by investigating the actual development and use of Bayesian adaptive methods in Phase I and Phase II clinical trials. She reports that since 1994, the methodological and ethical advantages of Bayesian designs have been demonstrated for Phase III clinical trials (Spiegelhalter et al., 2004) and for medical devices clinical trials (Campbell, 2005). Chevret (2012) reports that since the turn of the 21st century, there has been considerable growth in the interest in Bayesian adaptive designs for clinical trials. Recently in the US, the Food and Drug Administration (FDA), responding to industry concern about speed of approval and cost, has issued non-binding guidance for early stopping study designs (FDA, 2019). In the UK, interest in adaptive trial received an impetus when Baroness Jowell mentioned the topic in her final address to the House of Lords (Hansard, 2018).

In addition, El Alili et al. (2017), in a scoping review, found that whilst numerous recommendations on how to analyse trial-based economics evaluation have been proposed, the statistical quality of trial-based economic evaluations is generally unsatisfactory, highlighting the need for better statistical tools to analyse this type of data. Over the past years, a number of methods have been proposed to include the cost of research (Berry & Ho, 1988; Jennison & Turnbull, 2000) and it has been argued that sequential methods have the potential to be incorporated in economic models with the aim at informing policy makers cost-effectiveness decisions. Pertile et al. (2014) developed a model of sequential estimation that view adoption, treatment, and research decisions as a single economic project and argued that a dynamic approach to HCT assessment could bring significant efficiency gains. Within a similar decision-theoretic framework, Chick et al. (2017) more recently developed a model of sequential experimentation in which the primary end-point is observed with delay. In addition, there have been several attempts to apply the real-option framework to clinical trial evaluation of economic data. This literature of investment under uncertainty aims at incorporating the dynamic nature of the decision process and considers the role of flexibility and irreversibility of investment. Palmer and Smith (2000) proposed the use of real options in order to handle uncertainty in health technology assessment (HTA) and to show that the degree of irreversibility of actions requires some flexibility in the timing of decisions. Driffield and Smith (2007) used real options to argue for a watchful waiting regime for diseases with slow progression. Forster and Pertile (2012) appealed to real options as a way to model irreversibility of action and advocated a greater role of this modelling framework in HTA.

Concomitantly, value of information (VoI) methods, developed by Raiffa and Schlaifer (1961) and popularised by Claxton (1999), have gained considerable traction over the years among the health economics community as a framework to address whether sufficient evidence has been gathered to support HCT investment decisions (see Fenwick et al., 2020; Rothery et al., 2020).

The VoI approach makes use of the range of possible payoffs obtained from a probabilistic sensitivity analysis of underpinning cost-effectiveness decision-analytic models (where statistical distributions are fitted to the model input parameters) to quantify the value that could be obtained by reducing the range of payoffs (up to a single estimate under value of perfect information) via further research. If the expected value of research, expressed in net health effects or net monetary benefits, is lower than the cost that would be incurred by running a new trial, than the DM may take a decision now (adopt/reject HCT) based on currently available evidence. The associated decision rule is thus sequential: trial, then review of evidence and decide whether a decision can be taken or if a further trial is needed.

Although [Raiffa and Schlaifer \(1961\)](#) advocated a fully sequential model of VoI, the mathematical formulation of such a model has proven difficult. In particular, the VoI approach is not explicitly dynamic in the sense that at each stage of the analysis it is (implicitly) assumed that the next sample will be the last one. [William and Kowgier \(2008\)](#) developed a VoI-based multi-stage adaptive design involving an early termination rule based on the expected net gain from the trial computed for each stage. It is theoretically possible to construct a purely multi-stage model that jointly determines the optimal decision values, but due to its complexity the authors suggest to proceed in two stages, where at each step the (ex ante) two-stage calculation is performed and the maximisation process is repeated. Thus, so far any VoI-based approach taken has been static, with single stage VoI estimation, not fully exploiting the cost-saving potential of sequential methods.

In this paper, we apply a hypothesis testing model developed in the spirit of the Bayesian sequential hypothesis testing framework of [Shiryayev \(1978\)](#). The model incorporates a Bayesian updating rule that links the evidence collected during a trial to the economic uncertainty that arises from such an estimate (see [Thijssen & Bregantini, 2017](#) for further details). Whilst the modelling approach is close to the [Peskir and Shiryaev \(2006\)](#) formulation, the approach resembles a real-option model as discounted monetary payoffs for adoption and abandonment are explicitly modelled and, thus, enter the health authority's objective function. The closest formulation to our model is the one of [Pertile et al. \(2014\)](#). Their approach, however, is one of sequential estimation, whilst ours is one of hypothesis testing.

Our contribution to the literature is fourfold.

- (i) We apply the Bayesian model to a recently approved HCT and show the potential of using a sequential framework in assessing whether the information obtained through a trial is enough for the DM to adopt or reject the technology. We quantify, using a real-world case study, the cost to the healthcare system of early adoption (abandonment) not supported by enough evidence. To the best of our knowledge, this is the first time such cost is identified in a sequential Bayesian setting. By using a Bayesian approach, unlike the VoI framework that quantifies the expectation over the maximum value that a new trial could generate, we quantify simply the expected value of future research. This aligns with expected utility theory, which is the dominant paradigm of decision-making under uncertainty.
- (ii) We compare our dynamic sequential model to the VoI framework. Whilst usage of the latter is now well established in both the health economics and medical statistics communities ([Steuten et al., 2013](#)), in recent years there has been more interest and applications of dynamic optimisation models to the evaluation of HCTs and medical decision-making (e.g. [Favato et al., 2013](#); [Grutters et al., 2011](#); [Pertile et al., 2014](#)). A comparison of the performances of these two approaches—which, to the best of our knowledge, has never been undertaken—is thus particularly instructive.
- (iii) We show how a sequential model can be used to formulate the research design for a trial with a focus on the payoff healthcare system.
- (iv) We simulate trial outcomes and explore the optimal stopping time distribution for the adoption and abandonment/rejection decisions. The analysis contributes to the current debate in the use of sequential methods in HCT assessment. The paper shows how a sequential adaptive approach can be developed from an existing statistical approach leading to a potential substantial reduction in decision time and decision errors leading to sizable gains to the healthcare system in terms of net health or monetary benefit.

The paper's structure is as follows. In Section 2, we provide a discrete-time set-up for the model and express the link between the evidence collected during a clinical trial and a Bayesian sequential updating device that gives rise to the posterior probability process. In Section 4, we illustrate how a Bayesian sequential hypothesis testing model can be used in practice by providing a relevant example based on a recent technology appraisal. Finally, in Section 5, a comparison is made between the evaluation obtained using the VoI approach and the sequential method.

2 Methodology

2.1 Model for sequential trial evidence under Bayesian updating

Our sequential model of inferential decision-making is related to the standard textbook case of testing the equality of the mean in two independent samples. In that model, the statistician takes one sample of patients of size n and randomly assigns patients to draw inferences on the parameter $\mu = \mu_1 - \mu_2$, where μ_i is the mean treatment effect in the group given the new HCT and μ_2 is the mean treatment effect in the group given the existing HCT.

Here, we are interested in a *dynamic* set-up where the sample information does not arrive as one database, but in smaller batches. After the arrival of every batch, the statistician then has to decide (a) whether or not to continue the trial and (b) if the decision is taken not to continue the trial, whether or not to adopt the new HCT. That is, the eventual sample size is *endogenous* and will be different in each case.

Our model thus applies a decision-theoretic approach to a basic model of sequential inference. We, therefore, have to make a clear distinction between the *payoffs* that drive the optimal decision and the *inferential* properties of the sequential trial that influence the expected values of decisions taken.

The main parameter of interest in our model is $\theta \in \{0, 1\}$ which encodes whether the new HCT is more effective than the existing HCT ($\theta = 1$) or not ($\theta = 0$). At the heart of our approach lies the sequential testing of the point hypotheses

$$H_0 : \theta = 0 \quad \text{against} \quad H_1 : \theta = 1.$$

Decisions are based on inferences about θ that are derived from trial observations. The results of these trials are reported in terms of a measure for the cost-effectiveness of the new HCT relative to the existing HCT. Here, we choose to use the *NIMB* metric, denoted X , given by

$$X = (\text{QALY}_1 - \text{QALY}_0)\lambda - (C_1 - C_0),$$

where QALY_0 and QALY_1 denote health gains, measured in quality-adjusted life years (QALY), under the null ($\theta = 0$) and alternative ($\theta = 1$) hypotheses, respectively. The terms C_0 and C_1 represent the costs of the existing and of the new technology, respectively, and λ represents the opportunity cost of spending the healthcare system's scarce resources or, in other words, how much the healthcare system can afford to pay for a QALY. The random variable X is assumed to have mean $\theta\mu$ for some $\mu > 0$. That is, the new HCT is either equally effective as the existing one, or it is better (in terms of NIMB) by, on average, an amount μ .

Our sample is now a sequence of random variables $X_n; n = 1, 2, \dots$, where X_n measures the *cumulative* evidence collected over the first n steps of the trial. It is important to recognise that this is *not* a sequence of iid random variables, because the X_n are sequentially observed. Instead the *increments* are assumed to be (conditionally) iid. Suppose that a sample of size n arrives in m batches of size $\Delta n = n/m$. Then, conditional on θ , we assume that

$$\Delta X_k := X_{(k+1)\Delta n} - X_{k\Delta n} = \begin{cases} \theta\mu\Delta n + \sigma\sqrt{\Delta n} & \text{with probability } 1/2, \\ \theta\mu\Delta n - \sigma\sqrt{\Delta n} & \text{with probability } 1/2, \end{cases} \quad 0 \leq k < m.$$

So, in expectation the cumulatively observed NIMB increases by $\theta\mu\Delta n$, i.e. by the expected NIMB of the sample batch. This mean is then subject to some random noise so that the variance of ΔX is $\sigma^2\Delta n$. So, σ can be interpreted as the standard deviation of 'one unit' of clinical trial. Note that the ΔX_k are iid.

Standard results from stochastic calculus (see, e.g. Shiryaev, 1978; Thijssen & Bregantini, 2017) give that the sequence $\{X_n; n \geq 0\}$ converges to an *arithmetic Brownian motion* as $\Delta n \downarrow 0$, which implies that for every n ,

$$X_n | \theta \sim N(\theta\mu n, \sigma^2 n).$$

This is, in fact, nothing more than a central limit theorem-like result, but at a ‘micro’ level: as Δn gets smaller the total cumulatively observed X_n gets ‘chopped’ up into ever smaller batches. For example, if X_n represents the results of n trials with 10 patients each, then a value $\Delta n = 1$ implies that sample information (on cumulative NIMB) arrives in batches pertaining to 10 patients. If $\Delta n = 1/2$, then each batch of new information, ΔX_n , pertains to five patients, etc. In the remainder of the paper, we will use this asymptotic idealisation of our sequential trial in which observations on observed NIMBs arrive continuously.

Conditional on the state of the world $\theta \in \{0, 1\}$, this sequence of observations is governed by the probability measure \mathbf{P}_θ , which are such that $\mathbf{P}_1(\theta = 1) = \mathbf{P}_0(\theta = 0) = 1$. Since we model statistical inference *à la* Bayes, it is assumed that the DM has a prior probability measure \mathbf{P}_p , where, for any $p \in (0, 1)$, $\mathbf{P}_p = p\mathbf{P}_1 + (1 - p)\mathbf{P}_0$. That is, the DM believes that, a priori, $\mathbf{P}_p(\theta = 1) = p$. If we denote by \mathcal{F}_n^X the filtration generated by the (cumulative) observations X_n , then the *likelihood ratio* is the Radon–Nikodym derivative (cf., Shiryaev, 1978)

$$\Lambda_n = \frac{d(\mathbf{P}_1 | \mathcal{F}_n^X)}{d(\mathbf{P}_0 | \mathcal{F}_n^X)} = \exp\left\{\frac{\mu}{\sigma^2} \left(X_n - \frac{\mu}{2} n\right)\right\}, \tag{1}$$

which measures at any sample size n the likelihood of the alternative hypothesis relative to the null hypothesis. Note that this ratio itself evolves over time (in fact, in the limit $\Delta n \downarrow 0$ it follows a geometric Brownian motion). From this, the *posterior* probability of the event $\{\theta = 1\}$ can be computed using Bayes’ rule (cf. Shiryaev, 1978):

$$\pi_n := \mathbf{P}(\theta = 1 | \mathcal{F}_n^X) = \left(\frac{p}{1-p} \Lambda_n\right) / \left(1 + \frac{1}{1-p} \Lambda_n\right). \tag{2}$$

Equation (2) provides the link between the evidence emerging sequentially from the trial and the expected *consequences* of the decisions that the decision maker can take. In the next section, we will make this link explicit. It should be noted that, whilst the inferential set-up uses trial units, in practical applications the DM will typically be interested in the amount of *time* that a trial takes. In the remainder of the paper we will, therefore, assume that a trial unit is given by the number of patients that can be treated in one year. The latter quantity will be denoted by N_{trial} .

2.2 Joint adoption and abandonment rules

In this section, we present a ‘two-sided’ problem where the DM has the option of either invest in the new HCT, continue trialling the technology or abandon the trial altogether. In line with the previous subsection, trial costs are measured in ‘trial units’. In the continuous-time limit, results from a trial unit arrive continuously. For simplicity, we assume that a population of $M > 0$ patients is treated every year over an infinite HCT life-time. The DM is assumed to discount the benefits of treatment, using continuous compounding, at a constant rate $r > 1$ (p.a). The payoffs accruing from adoption/rejection in the different states of nature are:

- $B_1 = M \frac{\text{QALY}_1 \lambda - C_1}{r}$: discounted stream of net benefits of the new HCT conditional on $\theta = 1$;
- $B_0 = M \frac{\text{QALY}_0 \lambda - C_1}{r}$: discounted stream of net benefits of the new HCT conditional on $\theta = 0$;
- $B = M \frac{\text{QALY}_0 \lambda - C_0}{r}$: discounted stream of standard care;
- $I > 0$: sunk cost of investing in the new health technology;
- $c > 0$: cost of a trial unit, i.e. the accumulated cost of a trial over one year;
- $L \geq 0$: loss of not using the new (better) technology (this could include the benefit forgone if $\theta = 1$).

The payoffs from investment and abandonment after n trial units are given by the conditional (on observed cumulative NIMB up to n) expectations of the payoffs resulting from these decisions. These payoffs are denoted by $F_I(\pi_n)$ and $F_A(\pi_n)$, respectively, and are given by

$$\begin{aligned} F_I(\pi_n) &= \pi_n(B_1 - I) + (1 - \pi_n)(B_0 - I) \\ &= \pi_n(B_1 - B_0) + B_0 - I \quad \text{and} \\ F_A(\pi_n) &= \pi_n(B - L) + (1 - \pi_n)B \\ &= B - \pi_n L. \end{aligned} \quad (3)$$

The payoffs in equation (3) are probability-weighted and incorporate the cost of making an incorrect inference (Type I or Type II error), as well as the benefit of making a correct inference. The investment payoff is a probability-weighted sum of the benefit of the new HCT conditional on $\theta = 1$ and the benefit of standard care conditional on $\theta = 0$ net of sunk investment costs. In the abandonment case, the expected payoff is given by the (known) benefit of standard care conditional net of the probability-adjusted loss attributed to not using the new technology if it is in fact superior ($\theta = 1$). Since the payoffs depend on the *current* beliefs in the new HCT being superior ($\theta = 1$) or not ($\theta = 0$), they evolve over time as new trial evidence emerges.

The payoffs F_I and F_A are realised at the time a decision is taken: F_I if the new technology is adopted and F_A if it is not. Obviously, if the decision is taken after the sample size has increased to n , then the DM will choose to adopt the new technology if, and only if, $F_I(\pi_n) > F_A(\pi_n)$. Note that at that time the trial is discontinued. Until then the DM incurs the running cost, c , of continuing the trial.

The DM is now confronted with the question *how long to let the trial run*. It is shown in [Thijssen and Bregantini \(2017\)](#) that, in the continuous limit, the optimal decision is to wait as long as the probability that $\theta = 1$ lies between two *triggers*, i.e. as long as $\pi_n \in (\pi_A, \pi_I)$, where the adoption trigger, π_I , and the abandonment/rejection trigger, π_A , are jointly determined as functions of the parameters of the model; cf. [Figure 1](#) for a graphical illustration.

Once these triggers are determined, the *value* of the trial is a function of the current (posterior) probability that $\theta = 1$ and the *expected* amount of discounting over payoffs that takes place until a decision is taken. The latter depends on the sample size that needs to be collected before the investment trigger π_I is reached, denoted by $n(\pi_I)$, or before the abandonment trigger π_A is reached, denoted by $n(\pi_A)$, whichever is smaller. Note that $n(\pi_I)$ and $n(\pi_A)$ are random variables: their realisation depends on the particular sample path of the posterior belief π_n that is observed. A different evolution of the observed cumulative NIMBs leads to a different evolution of the posterior belief π_n , which, in turn, typically leads to different decisions (investment or abandonment) at different sample sizes. For example, if the cumulative NIMB increases fast in the initial phases of a trial, then π_I is more likely to be reached sooner than π_A than along a sample path where the cumulative NIMB initially hovers around 0.

At any particular time the DM can look at the posterior belief in the event $\theta = 1$ and compute the expected payoff of the decision to invest in or abandon the HCT, net of the cost of the trial, given the triggers π_I and π_A , i.e.

$$F^*(\pi) = \begin{cases} F_I(\pi_I) & \text{if } \pi \geq \pi_I \\ -\frac{c}{r} + \mathbb{E}_\pi[e^{-rn(\pi_A)} \mathbf{1}_{n(\pi_A) < n(\pi_I)}] \left(F_A(\pi_A) + \frac{c}{r} \right) \\ \quad + \mathbb{E}_\pi[e^{-rn(\pi_I)} \mathbf{1}_{n(\pi_I) < n(\pi_A)}] \left(F_I(\pi_I) + \frac{c}{r} \right) & \text{if } \pi_A < \pi < \pi_I \\ F_A(\pi_A) & \text{if } \pi \leq \pi_A. \end{cases} \quad (4)$$

It turns out that the expectations in equation (4) can be computed explicitly (cf., [Thijssen & Bregantini, 2017](#) and [Appendix B](#)) and that they depend on the model's parameters. The DM now needs to determine the triggers π_I and π_A such that equation (4) is maximised for every possible $\pi \in (0, 1)$. It has been shown by [Thijssen and Bregantini \(2017\)](#) that this can be done by solving the system of non-linear equations that is given in [Appendix B](#).

The interpretation of equation (4) is as follows. If $\pi \geq \pi_I$, then it is optimal to immediately adopt the new technology and realise the (expected) payoff $F_I(\pi)$. Similarly, if $\pi \leq \pi_A$, then it is optimal to immediately abandon the trial and realise the (expected) payoff $F_A(\pi)$. In the *continuation region*,




Figure 1. Adoption and rejection decision triggers: Bayesian posterior process as evidence is gathered.

i.e. when $\pi_A < \pi < \pi_I$, it is optimal to continue the trial. The current value of the trial then consists of three components. The first component, $-c/r$, represents the (discounted) cost flow of never stopping the trial and incurring its running cost forever. This perpetual cost is corrected for by two components, representing the that, at some point in the future, either π_I or π_A will be hit. If π_I is hit before π_A , then the DM will, at that time, adopt the technology and stop the trial. Hence, the payoff at that time will be $F_I(\pi_I) + c/r$. To get the present value of this payoff, it needs to be discounted using the expected discounted factor. On the other hand, if π_A is hit before π_I , then the DM will, at time time, abandon the trial. Hence, the payoff at that time will be $F_A(\pi_I) + c/r$, which then needs to be discounted using the expected discounted factor.

2.3 Expected sample size and cost of trial

In this section, we provide simple formulas to compute the expected sample size for a trial and its related cost. We first introduce the expected discount factors, denoted by $\hat{v}_{\pi_A, \pi_I}(\pi)$ and $\check{v}_{\pi_A, \pi_I}(\pi)$, of first reaching π_I or π_A , respectively. These are given by (cf., Thijssen & Bregantini, 2017),

$$\begin{aligned} \hat{v}_{\pi_A, \pi_I}(\pi) &= \mathbb{E}_\pi[e^{-r\hat{n}(\pi)} \mid \hat{n}(\pi_I) < \check{n}(\pi_A)] \mathbf{P}_\pi(\hat{n}(\pi_I) < \hat{n}(\pi_A)) \\ &= \sqrt{\frac{\pi(1-\pi)}{\pi_I(1-\pi_I)} \left(\frac{1-\pi_A}{\pi_A} \frac{\pi}{1-\pi} \right)^\gamma - \left(\frac{\pi_A}{1-\pi_A} \frac{1-\pi}{\pi} \right)^\gamma}, \quad \text{and} \\ \check{v}_{\pi_A, \pi_I}(\pi) &= \mathbb{E}_\pi[e^{-r\check{n}(\pi)} \mid \hat{n}(\pi_I) > \check{n}(\pi_A)] \mathbf{P}_\pi(\hat{n}(\pi_I) > \hat{n}(\pi_A)) \\ &= \sqrt{\frac{\pi(1-\pi)}{\pi_A(1-\pi_A)} \left(\frac{1-\pi}{\pi} \frac{\pi_I}{1-\pi_I} \right)^\gamma - \left(\frac{\pi}{1-\pi} \frac{1-\pi_I}{\pi_I} \right)^\gamma}, \end{aligned}$$

where $\hat{n}(\pi_I)$ and $\check{n}(\pi_A)$ denote the first-hitting times of π_I from below and π_A from above, respectively, and

$$\gamma = \frac{1}{2} \sqrt{1 + 4r \left(\frac{\sigma}{\mu} \right)^2} \quad (> 1/2).$$

The expected total cost of a trial, given a current belief $P_\pi(\theta = 1) = \pi$, denoted by TC, equals (cf., Thijssen & Bregantini, 2017):

$$\mathbb{E}_\pi(\text{TC}) = \frac{c}{r} \left(1 - \hat{v}_{\pi_A, \pi_I}(\pi) - \check{v}_{\pi_A, \pi_I}(\pi) \right). \quad (5)$$

The random variable $n^*(\pi_I, \pi_A) := \hat{n}(\pi_I) \wedge \check{n}(\pi_A)$ gives the sample size at which a decision is taken, the *expected* sample size at which a decision (either for investment or for abandonment) is taken (given a current belief $P_\pi(\theta = 1) = \pi$) is given by (cf., Poor & Hadjiliadis, 2009),

$$\begin{aligned} \mathbb{E}_\pi[n^*(\pi_I, \pi_A)] &= \frac{2\sigma^2}{\mu^2} \left\{ \log \left[\left(\frac{\pi}{1-\pi} \right)^{1-2\pi} \left(\frac{1-\pi_A}{\pi_A} \right)^{1-2\pi_A} \right] \right. \\ &\quad \left. + \frac{\pi - \pi_A}{\pi_I - \pi_A} \log \left[\left(\frac{\pi_A}{1-\pi_A} \right)^{1-2\pi_A} \left(\frac{1-\pi_I}{\pi_I} \right)^{1-2\pi_I} \right] \right\}. \end{aligned} \quad (6)$$

Note that this expectation is measured in trial units, so that the expected total number of patients in the trial equals $N^* = \mathbb{E}_\pi[n^*(\pi_I, \pi_A)] \times N_{\text{trial}}$. The realised sample size is, ex ante, uncertain and depends on the results of samples as they are sequentially observed.

3 Decision-making in HCT assessment

In this section, we outline the main features of the VoI approach that will be applied in later sections of the paper. Value of information analysis (see Pratt et al., 1995) has been proposed as a framework to evaluate the value of obtaining additional evidence on the expected cost-effectiveness of new HCTs in order to reduce the risk and associated consequences of making an incorrect investment decision, expressed in net health or net monetary losses.

The cornerstone metric of the framework is the *expected value of perfect information* (EVPI) that is used to compare the expected payoff (expressed in net health benefit or net monetary benefit) of a decision made with all uncertainty around the input parameters underpinning costs and effects estimation being resolved (i.e. no decision uncertainty) and the expected payoff of a decision made with currently available evidence (Rothery et al., 2020). Since perfect information is not achievable, the EVPI metric provides an upper-bound to the value of collecting additional evidence. As a result, if EVPI is null or negligible we can already establish that uncertainty around the investment decision is low and that there is therefore no point in undertaking further research to reduce it.

Since acquiring information is costly, analysts will mostly be interested in the value of reducing uncertainty around the subset of parameters that drive decision uncertainty. This value can be quantified using the expected value of perfect parameter information (EVPPI) that evaluate the difference in payoff when having perfect information for the subset of parameters of interest (θ_i) whilst remaining complementary parameters remain uncertain and the expected payoff of a decision made with currently available evidence. EVPPI also provides an upper-bound to the value of undertaking additional studies to inform θ_i since perfect information is typically not achievable with a finite sample size. It is, however, possible to estimate the expected value of a study that will result in data D that informs (θ_i) using the same statistical method that underpins EVPPI computation and this is referred to as the *expected value of sample information* (EVSI).

Once EVSI has been established, it can be compared with the cost of undertaking the data collection exercise that is required to provide data D , and this is known as the expected net benefit of sampling (ENBS). Expected net benefit of sampling has been proposed as a metric to help support trial design (Ades et al., 2004) and, in particular, to help identify the optimal trial size (or sampling size) as it is expected to be directly related to the magnitude of reduction in uncertainty around the parameters of interest.

Note that EVSI is the difference between the expected value of a decision after N samples have been collected and the expected value of a decision made with current information. The EVSI can

be calculated for a particular sample size from the prior information and the estimate of the sample variance (σ^2/N). A sample of size N , given the uncertain parameters γ , will give a sample result D . If the sample result were known, it would be possible for the DM to choose the alternative with the maximum expected payoff.

It is possible to compute the expected NIMB by averaging over the posterior distribution of the NIMB given the sample result D :

$$\max \{ \mathbb{E}_{\gamma|D}[X], 0 \}. \tag{7}$$

As the value of D is not known in advance (i.e. the result of the sample is not known), the expected value of a decision taken with sample information is computed by averaging the maximum expected NIMB over the distribution of possible values of D .

The EVSI is the difference between the expected value of a decision made with sample information and the expected value with current information:

$$\text{EVSI} = \mathbb{E}_D \max \{ \mathbb{E}_{\gamma|D}[X], 0 \} - \max \{ \mathbb{E}_{\gamma}[X], 0 \}. \tag{8}$$

The EVSI proposed in equation (8) is for a single study design and single sample size. In order to establish the optimal sample size for a particular study these computations needs to be repeated for various sample sizes N .

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

$$\text{ENBS}_N = \text{EVSI}_N - C_N. \tag{9}$$

The optimal sample size N^* is then determined by maximising ENBS_N .

4 Using the decision model alongside cost-effectiveness studies

Whilst the model described above establishes a link between evidence accumulated during a trial and the payoffs related to investing or abandoning an HCT, in reality, we are often unable to observe interim economic data at the trial level. Additionally, trials display limitations such as truncated time horizons and failures to incorporate all evidence (Sculpher et al., 2006). Rather, we observe the treatment effect once the trial is concluded and this is used to inform economic models that estimate the related health gains and incremental costs.

Our approach tests whether μ (the mean NIMB in a ‘trial unit’) has a certain economic value (i.e. by setting λ such that μ becomes positive). In the theoretical model outlined in Section 2.1, the standard deviation is fixed and in applications of sequential hypothesis testing σ is usually derived from previous studies. In what follows, in order to obtain a measure of economic uncertainty, we estimate the standard deviation of the NIMB from a probabilistic sensitivity analysis (PSA) output of the decision-analytic model.

In order to compute triggers and posterior probabilities, the approach outlined in Section 2 requires a limited number of inputs: the mean NIMB estimate (μ) and a measure of its uncertainty (σ) obtained from an economic model, the size of the population of patients in a trial unit (N_{trial}), the (irreversible) investment cost (I), the loss of not using the best technology (L), and the size of the target population that are used to compute the payoff values (M). By explicitly incorporating the target population the model considers the cost-effectiveness for the relevant population treated by the healthcare system rather than just the trial population. In our view, this is a much more realistic analysis, closer in spirit to budget impact analyses assessed by agencies that need to consider the expenditure impact of adopting new technologies (see NICE, 2022 for a detailed discussion). In the next section we show, in an applied case, how our sequential approach can be used in conjunction with estimates obtained from a Markov decision-analytic model typically used in health economics cost-effectiveness analysis (see for example NICE, 2022 for UK economic evidence guidelines). In Section 4.1, we obtain the triggers for adoption/rejection of an HCT, whilst in

Table 1. Model inputs

Parameter	Unit	Notation	Value	Source
UK patient population	patients/yr	M	9,000	NICE (2007)
Size of trial unit	patients/yr	N_{trial}	400	Assumed
Cost of standard care	£/patient	C_0	26,546	Estimated
Cost of Omalizumab	£/patient	C_1	67,137	Estimated
QALY Standard care	yrs/patient	$QALY_0$	10.17	Estimated
QALY Omalizumab	yrs/patient	$QALY_1$	11.37	Estimated
Sampling cost	£mn/yr	c	£3.52	Estimated
Discount rate	%/yr	r	3.5%	Assumed
Investment cost	£ mn	I	0	Assumed
Loss for not using the new technology	£ mn	L	0	Assumed

Table 2. Inferential parameters, decision triggers, break-even trigger, and observed INNOVATE posterior in the base-case model

λ	Parameters		Posterior beliefs			
	$\mu/\text{patient}$	$\sigma/\text{patient}$	π_A	$\bar{\pi}$	π_1	$\pi_{1/2}$
£35k	£1,409	£11,282	0.99642	0.96645	0.96648	0.5032
£47.5k	£13,409	£14,443	0.6212	0.7121	0.7847	0.7149
£55k	£25,409	£17,648	0.4155	0.6150	0.7736	0.8435

Section 5 those triggers are used to show how a research design that allows for early stopping can lead to shorter trials and consequently savings and higher benefits to patients.

4.1 Illustrative case study: the INNOVATE clinical trial

In order to illustrate the use of the model together with a standard cost-effectiveness study, we take a clinical trial that provided evidence for a recent NICE technology appraisal. The INNOVATE clinical trial was aimed at assessing the efficacy of Omalizumab, a humanised antibody designed to reduce sensitivity to inhaled or ingested allergens. In the INNOVATE study, add-on Omalizumab significantly reduced clinically significant non-severe exacerbation rates by 26%, severe exacerbation rates by 50% and emergency visits rates by 44%, and significantly improved asthma-related quality of life compared with placebo (Buhl, 2007). Some 420 patients partook in the INNOVATE double-blind RCT and it lasted 6 months. We use these figures to approximate the number of patients who are treated with the new HCT as 400 per year, which acts as our trial unit. We denote the posterior belief in $\theta = 1$ obtained after the INNOVATE trial by $\pi_{1/2}$.

However, there has been controversy over whether or not the benefits of Omalizumab justify its cost (Brown et al., 2007; NICE, 2007). In particular, there has been concern about the great uncertainty over the net benefits of the drug. In this section, results from the INNOVATE trial are used in order to produce a economic end-point that can illustrate a Bayesian sequential assessment of the cost-effectiveness for Omalizumab. Whilst this is clearly not a complete assessment of Omalizumab, the aim is to provide an illustration of the model's potential and its use in HTA. A simple but realistic Markov model is built using values taken from Faria et al.'s (2014) cost-effectiveness analysis of Omalizumab. Table 1 displays the main parameters.

Ideally the sampling cost should be calculated from the total cost incurred during the trial. Given that this was unavailable, the cost of running the trial (c) has been computed by dividing the total drug cost by the number of patients in the trial. Given that the main cost driver for the trial is the

price of Omalizumab (totalling £1,200,000 vs £300,000 for standard care) this seems to be a reasonable illustrative measure of per-patient cost and is in line with the annual cost of £8,056 indicated by Faria et al. (2014).

Our procedure to calculate μ and σ is as follows: the economic model simulates a cohort of 1,000 patients and reports the average effectiveness (NIMB) for the cohort. We then replicate outcomes 1,000 times in a PSA and take the average NIMB as a proxy for the individual-level effectiveness and uncertainty in outcomes. In the trigger calculations, this value gets multiplied by the number of patients in the trial, providing a estimate value for the overall trial population’s NIMB. The QALYs are valued at an exogenously given price per QALY (λ). In Table 2, we record the mean NIMB and its standard deviation per patient. In our calculations these are then transformed to yearly figures by appropriately scaling up using the trial size p.a. (N_{trial}). The table also records the resulting abandonment and adoption triggers.

4.2 Decision triggers

Figure 2 shows the abandonment decision trigger for the INNOVATE trial at a value per QALY of $\lambda = \text{£}35,000$. This is the first value at which μ becomes positive and it is close to the incremental cost effectiveness ratio (ICER) estimated in Jones et al. (2009). At this value per QALY, the abandonment trigger is very high, taking a value of 0.96642. So, even if one has a posterior belief in $\theta = 1$ of, say, 95%, then the trial should still be abandoned. This happens for several reasons. First, because λ is low, the mean μ is very close to zero. It was shown by Thijssen and Bregantini (2017) that the continuation region is ‘expanding’ in μ , so that a low value of μ corresponds to a narrow continuation region. Secondly, at a low value for λ the ‘break-even’ belief (i.e. the posterior belief for which investment and abandonment have equal expected value) is very high ($\bar{\pi} = 0.96648$). Since it is always the case that $\bar{\pi} \in (\pi_A, \pi_I)$, the abandonment region is very large in this case. It should then not come as a surprise that even after the first trial this project should be abandoned.

This is an interesting case as it indicates that at the given posterior value of $\pi_{1/2} = 0.5032$ and for a neutral prior of 0.50 the technology is immediately rejected. The posterior probability process




Figure 2. Decision triggers for $\lambda = 35,000$.

is in fact close to its original value, indicating a high level of uncertainty on the technology cost-effectiveness. We should recall that abandonment (rejection) in our model means that no further research should take place. This is due to the fact that once the posterior enters the rejection region, enough information has been collected, and any further research will only bring additional costs in terms of health benefit foregone and sampling costs. Unless the DM has prior information that very strongly indicates that the technology is cost-effective, or is willing to increase the trigger value λ , the model rejects the technology.

Reaching such a sharp result just based on the INNOVATE trial seems decisive, but given the high cost of Omalizumab and the high uncertainty surrounding the cost-effectiveness estimate, it is not completely unexpected. One of the key aspects of a real-option set-up is that uncertainty decreases the value of the technology. In contrast to the VoI, where high uncertainty would increase the value of a trial, and thus further research would be recommended, the real options approach makes a balance between trial costs and NIMB expectations. Even with a positive expected NIMB, a high NIMB standard deviation indicates that the trial needs to run for a long time in order to reduce uncertainty. Given the per-period cost of running a trial, high uncertainty reduces the value of gaining further evidence, in this case leading to an outright rejection of the technology. In the next section, it can be seen that as the expected NIMB increases, further research becomes more valuable and the posterior probability falls in the continuation region.

Finally, Faria et al. (2014) find that Omalizumab ‘represents good value for money only in severe subgroups and under optimistic assumptions regarding asthma mortality and improvement in health related quality of life’. Rejecting the technology at $\lambda = £35,000$ after half a trial unit indicates that the technology is not cost-effective for the main trial population. The manufacturer could use this information, e.g. to provide further evidence on subgroups that might have higher QALY gains than the standard population. Faria et al. (2014) report that the most severe population subgroups tend to display a higher QALY gain from reducing asthma events and their severity and further analysis would have helped the manufacturer to provide more evidence on the cost-effectiveness for these groups.

4.3 Varying the value per QALY

In this section, we use different hypothetical values for λ in order to illustrate some of the features of our model. Given that the technology is rejected at $\lambda = 35,000$, we increase the trigger λ at the value of $£47,500$. This increases the payoff and force the abandonment decision trigger to decrease (see Figure 3). Two parameters are simultaneously influencing the decision triggers, the new payoff given by $\lambda = 47,500$ and the standard deviation that becomes proportionally smaller. These changes make adoption more appealing in risk-adjusted terms and this is reflected by the new adoption trigger π_1 . Similarly, a higher positive payoff and lower risk-adjusted standard deviation make abandonment less likely and this is reflected by a new lower abandonment trigger π_A .

At the value per QALY of $£47,500$, the model does not yet recommend adoption after the INNOVATE trial alone, but, rather, indicates that further research is needed. The aspect of further research deserves a separate section and will be discussed in Section 4.5.

Figure 4 shows the adoption and abandonment triggers for $£55,000$ per QALY. By increasing the value per QALY gained the cost-effectiveness estimate becomes more appealing in risk-adjusted terms and the model adopts the technology as $\pi_{1/2}$ hits the adoption trigger. In fact, at $£55,000$ per QALY, abandoning research poses serious costs to the healthcare system and this is reflected by the much lower abandonment trigger π_A . The higher observed posterior probability $\pi_{1/2}$ after the INNOVATE trial indicates adoption of Omalizumab. This is close to the finding of Faria et al. (2014) who found the ICER to be around $£57,000$ for the overall population.

4.4 Trigger sensitivity

One interesting exercise is to vary the value of λ and observe how, by increasing the expected pay-offs, the adoption triggers gradually change. This approach is in principle similar to the commonly adopted sensitivity analysis in HTA that aims at understanding how net benefits respond to different parameters.




Figure 3. Decision triggers for $\lambda = 47,500$.

Figure 5 shows the observed posterior $\pi_{1/2}$ from the INNOVATE trial and triggers π_I and π_A for different values of λ taken between £35,000 and £55,000. As the expected NIMB (conditional on $\theta = 1$) increases, the abandonment and adoption triggers vary significantly. At values in the range £35,000–£42,500, the model rejects the technology outright. It commands that more evidence is needed in the ranges £42,500–£50,000 and adopts the technology for values above £51,382.

4.5 Option value and waiting for more evidence

For the case of $\lambda = £47,500$, $\pi_{1/2}$ falls in the continuation region, which indicates that the DM should wait for more evidence. In between the triggers, the posterior, together with the payoff specification can be used to obtain the value of the investment and abandonment option. This is done by subtracting the investment (abandonment) payoff to the payoff given by immediate investment (abandonment). This gives the value of waiting for further evidence (i.e. the opportunity cost of investment with current evidence). The investment and rejection payoffs for a given posterior probability π is given by

$$\begin{aligned} F_{EI}(\pi) &= \pi_I F_I(\pi_I) - \pi F(\pi) && \text{early adoption case} \\ F_{ER}(\pi) &= \pi F(\pi) - \pi_A F_A(\pi_A) && \text{early rejection case,} \end{aligned} \tag{10}$$

where $F_{EI}(\pi_I)$ is the payoff of early investment at the investment trigger and $F_{ER}(\pi_A)$ is the payoff of early rejection of the technology at the abandonment trigger. The payoff ‘in research’ is given by $\pi F(\pi)$ for the adoption case and for the rejection case respectively. The ‘in research’ payoff is the expected investment (abandonment) payoff at any time where the posterior belief in $\theta = 1$ equals π . At the trigger, the additional value of gathering further evidence is zero, and collecting evidence beyond such point involves a loss to the healthcare system. The per-patient and overall population option values for undertaking early investment and abandonment are given in Table 3.

In the illustrative case study above, adopting Omalizumab purely on the basis of the INNOVATE trial, at £47,500 per QALY implies a yearly loss to the healthcare system of




Figure 4. Decision triggers for $\lambda = 55,000$.




Figure 5. Decision triggers and posterior $\pi_{1/2}$ for different values of λ .

£2,415 per patient in case of early adoption and £4,373 in case of early abandonment. Given that the estimated size of the population eligible for Omalizumab is $M = 9,000$, the value of waiting for more information is £21,737,606 in case of early adoption and £39,353,178 in case of early

Table 3. Option values

Adoption payoff ($\lambda = \text{£}47,500$)	Per patient	For eligible population ($M = 9,000$)
Option value early adoption	£2,415	£21,737,606
Option value early rejection	£4,373	£39,353,178




Figure 6. Option values for early adoption and early abandonment.

rejection. Figure 6 shows how the option value is computed as the difference between the current ‘in research’ payoff and the investment (abandonment) payoff due at the optimal stopping time after observing n^* trial units. Values are represented in terms of net benefit as specified by equation (3). The figure shows that great care should be taken when adopting a technology which displays large uncertainty, as the value of waiting for more evidence is greater in such cases.

4.6 Adoption payoffs

Real-option models are typically set up as ‘one-sided’ problems, i.e. only the investment decision is modelled (Dixit & Pindyck, 1994). Whilst this set-up is technically possible, the following illustration shows that it does not lead to sensible results when the methodology is extended to hypothesis testing. A one-sided adoption model has the following payoffs:

- $B_1 = M \frac{QALY_1 \lambda - C_1}{r}$: discounted stream of net benefits of the new HCT conditional on $\theta = 1$;
- $B_0 = M \frac{QALY_0 \lambda - C_1}{r}$: discounted stream of net benefits of the new HCT conditional on $\theta = 0$;
- $I > 0$: sunk cost of investing in the new health technology;
- $c > 0$: cost of a trial unit.

The expected net present value of investment is given by

$$F_I(\pi) = \pi B_1 + (1 - \pi)B_0 - I = \pi(B_1 - B_0) + (B_0 - I). \tag{11}$$

Evidence is accumulated and the posterior probability is updated by Bayes' rule as in equation (2). If the evidence, as measured by the posterior probability that $\theta = 1$, collected during the trial hits a particular (endogenously determined) trigger π_I adoption should take place and net expected benefits are maximised.

Following from the investment literature (e.g. Dixit & Pindyck, 1994; Stokey, 2008), the value of the new technology will depend on the value of π at the time of adoption and is denoted by a function $\tilde{F}^*(\pi)$. With $\pi < \tilde{\pi}_I$ the posterior is in the so-called *continuation region* and the trial should continue. The payoffs are as follows:

$$\tilde{F}^*(\pi) = \begin{cases} -\frac{c}{r} + \mathbb{E}_\pi[e^{-rn(\tilde{\pi}_I)}] \left(F_I(\tilde{\pi}_I) + \frac{c}{r} \right) & \text{if } \pi < \tilde{\pi}_I \\ F_I(\tilde{\pi}_I) & \text{if } \pi \geq \tilde{\pi}_I, \end{cases}$$

where, as above, $n(\tilde{\pi}_I)$ is the first time that $\tilde{\pi}_I$ is reached and adoption becomes optimal (for details, see Thijssen, 2013; Thijssen & Bregantini, 2017). The adoption trigger can be found by solving the following equation:

$$\phi(\pi)F_I'(\pi) = \phi'(\pi)F_I(\pi),$$

where $\phi(\pi) = \sqrt{\pi(1-\pi)}\left(\frac{1-\pi}{\pi}\right)^\gamma$, $\phi'(\pi) = \phi(\pi)\frac{1/2+\gamma-\pi}{\pi(1-\pi)}$, and $F_I'(\pi) = B_1 - B_0$. From this condition, we find that the adoption trigger equals

$$\tilde{\pi}_I = \frac{(1/2 + \gamma)(B_0 - I + c/r)}{(1/2 - \gamma)(B_1 - B_0) + (B_0 - I + c/r)}. \quad (12)$$

The value function is then given by

$$\tilde{F}^*(\pi) = \begin{cases} -\frac{c}{r} + \frac{\phi(\pi)}{\phi(\tilde{\pi}_I)} \left(F_I(\tilde{\pi}_I) + \frac{c}{r} \right) & \text{if } \pi < \tilde{\pi}_I \\ F_I(\tilde{\pi}_I) & \text{if } \pi \geq \tilde{\pi}_I. \end{cases} \quad (13)$$

Whilst it is possible to obtain an analytical solution for a one-sided problem, using the parameters obtained from the INNOVATE trial lead to a trigger $\tilde{\pi}_I$ above one. Intuitively, both Type I and Type II errors (mirrored by adoption and abandonment in our model) are needed to provide adequate decision rules as in the classical test.

4.7 Evaluating a sequence of trials

Next, we look at the situation where we wish to use the information contained in more than a trial, or a sequence of trials. Whilst the INNOVATE trial is the base case (Faria et al., 2014) by using the EXALT trial (open label non-placebo controlled trial) as an additional input, it is possible to view the decision-making process for Omalizumab as a sequence of trials, thus fully exploiting the potential of the sequential model. In this case, the posterior probability $\pi_{1/2}$ obtained after observing the INNOVATE trial at 6 months is used as the prior probability in the evaluation of the EXALT outcomes for $\lambda = \text{£}50,000$. The EXALT trial was of a similar scale as the INNOVATE trial and took 6 months, so that it can also be taken to represent half a trial unit. The posterior after observing both trials is, thus, π_1 . Figure 7 shows the posterior given by the INNOVATE and the EXALT study (labelled as $\pi_{1/2}$ and π_1 , respectively, indicating the value of the posterior after 6 months and 1 year). At the end of the INNOVATE trial, a cost-effectiveness estimate is produced and used in the model to produce a posterior probability. At $\lambda = \text{£}50,000$, the posterior probability $\pi_{1/2}$ lies in the continuation region and recommends to continue research. After another 6 months period (the length of the EXALT trial) another economic study is conducted and another cost-effectiveness study is produced. By viewing the two trials as a sequence, information about the cost-effectiveness of the drug accumulates over time. The core idea is that as the first trial was inconclusive, the DM decides to continue research and runs another trial. The model, through the Bayesian posterior,




Figure 7. Sequential decision-making in HTA. INNOVATE posterior $\pi_{1/2}$ becomes a prior probability in the subsequent study (EXALT) and produces the (accumulated) posterior probability of cost-effectiveness π_1 . HTA = health technology assessment.

captures the fact that the second trial comes about as a consequence of the first trial being inconclusive. This is an interesting case that reflects common evidence review procedures: when the studies are taken alone, none of them taken in isolation is sufficient for the DM to make a decision. However, by taking these together, the DM can ‘sum’ information obtained from the economic trial analysis and make a policy decision. In the case study presented in Figure 7, one can see that after the INNOVATE trial with $\lambda = \text{£}50,000$, the posterior probability ($\pi_{1/2}$) is close to 0.75. This is then used as a prior in the next step, leading the combined INNOVATE and EXALT trials to generate a posterior probability (π_1) above the adoption trigger of 0.85.

5 Vol and the sequential hypothesis testing approach

5.1 Optimal sample size for a research design

The decision principle in the VoI framework is as follows: at the end of the trial the EVSI associated with obtaining additional evidence on parameters of interest (i.e. those are known to be key drivers of decision uncertainty) is computed and compared to the cost of running a new trial with the difference being the ENBS. If ENBS is positive then it is worthwhile running a new trial otherwise evidence is deemed to be sufficient to make a decision. The process is repeated until it is not worthwhile to run a new trial and a decision is reached. The issue with such an approach is that there is little flexibility in early stopping and whilst sequential in principle, the VoI approach is de-facto static (the VoI framework constitutes a sequence of static optimisations, rather than a full dynamic optimisation model) and allows for little flexibility in the timing of the decision. In this section we give a proof-of-concept study that displays the advantages of a dynamic sequential framework over a static one in a research design.

Both the VoI and our Bayesian sequential approaches allow estimation of the expected sample size for a trial. Equation (6) gives the expected hitting time for a trial process and can easily be computed with the information at hand. Using the values reported in Table 1 and the resulting values of the triggers π_A and π_I for the case $\lambda = \text{£}47,500$, we obtain an expected sample size of $n^* = 230$. Note that this is with slight abuse of notation: in trial units the expected optimal stopping time is just over half a trial unit (just over 6 months). This represents 230 observations. We make this slight change to allow for one-to-one comparisons with the VoI approach.




Figure 8. Expected net benefit of sampling (ENBS). Optimal sample size N^* for a trial.




Figure 9. Relative frequency of hitting times (trial stopped due to the posterior reaching a trigger).

By using equation (9) for the ENBS, we obtain a sample size of $N^* = 296$. Figure 8 shows the EVSI and the optimal sample size N^* . The difference in notation is justified by the fact that n^* is an optimal stopping time and N^* is an optimal sample size. The two methods give sample sizes that are reasonably close. Next, we resort to a simulation study in order to check the properties of both methods.

5.2 Simulation study

Figure 9 shows the distribution of n^* for a number of simulated trials. On average the decision is made at $n = 246$. This value is similar to the one predicted above at $n^* = 230$. It is important to notice how many of the simulated trials hit either decision trigger before $n = 230$. Table 4 reports

Table 4. Simulated values decision times (in days) with cumulative frequencies

Decision time (in days)	180	120	90	60
Frequency of decisions	51%	32%	27%	8%

values indicating the percentage of trials for which a decision can be made within a given time. Sixty per cent of the simulated trials terminate before the predicted $n = 230$. For example, 51% of decisions can be made in the first 6 months. This implies that a substantial saving in terms of decision time and trial costs can be made, when compared to a research design specified using the EVSI approach. A substantial 8% of decisions can be made within 2 months simulated trials. The shortest time needed to make a decision in our simulation is 1 month (29 days). The implication is that there is great potential for many trials to be stopped earlier and for substantial efficiency gains to be made. It can also be noted that 76 simulated trials (accounting for about 7% of the total) run over 48 months, indicating that a DM would expect to make a decision within that time frame with high probability.

We believe the above example underlines the deep limitations of the current approach that sets a research design sample size N a priori with no flexibility in the timing decision and highlights the benefits of a sequential approach to decision-making. In fact, it would be desirable, in terms of costs and patients' health benefits, to stop as soon as sufficient (interim) evidence has been gathered. However, because of the static nature of the VoI approach this is not possible, as one needs to wait for the end of the trial before making an assessment, which consequently leads to losses to the healthcare system.

5.3 Limitations of our modelling approach

Our approach relies on the clinical trial outcomes to follow a normal distribution (Brownian motion). Whilst this can be shown to be a reasonable assumption when it comes to clinical trial data (see Appendix A), it does not necessarily hold for cost–benefit estimates such as the NIMB produced by an economic decision-analytic model. For example, the model simulated output for the INNOVATE and EXALT trials that are used to compute the mean NIMB and its standard deviation are assumed to be normally distributed. In practice, although with a large number of samples these parameters tend to approximate a normal distribution, we do not consider the impact of skewness or ‘fat tails’ in the distribution. Further work should be carried out to incorporate measures of uncertainty that do not conform to the normal distribution, as non-linearities and skewed distributions are the norm rather than the exception in healthcare economic decision models.

6 Conclusion

The paper contributes to the ongoing policy debate surrounding the adoption of HCTs on the basis of statistical evidence and in particular on what statistical methods should be used to assess value-for-money for new HCTs. We present a sequential Bayesian hypothesis testing model that includes discounting and that allows for the assessment of evidence gathered in a clinical trial in a decision-analytic framework. The model provides adoption and abandonment/rejection triggers that can be used to judge whether gathered evidence is enough to make an approval or rejection decision on the basis of the economic value of the new HCT balanced against the value of continuing the trial to gather further evidence. These decision triggers are computed by taking into account the benefit of adopting the technology, the loss given by terminating research with too little evidence, the cost of sampling and the uncertainty surrounding the cost-effectiveness estimate. The paper constitutes a first attempt at quantifying the opportunity cost of decisions taken at a non-optimal time in a sequential setting. We show that DMs would benefit from sequential methods when assessing the economic value of new HCTs in terms of making quicker decisions which in turn brings higher patients' benefit and reduce trial cost. In the paper, we also propose to consider economic estimates obtained from decision models as single data points in a sequence of trials and provide the foundations for future research in this area.

Further, the paper shows how the parameters obtained from a clinical trial can be used to inform an optimal research design and that there are health gains to be made by adopting a sequential approach. Our simulation study shows how a dynamic sequential approach, when compared to the VoI framework, detects a technology with better outcomes sooner and can lead to larger gains for

We term the ratio

$$t = v_n/v_N = \text{Var}(S_n)/\text{Var}(S_N) \tag{15}$$

the trial fraction (15) measures how far into the trial we are.

Denote the interim z -score $S_n/v_n^{1/2}$ at trial fraction t by $Z(t)$, and define the value

$$B(t) = \frac{S_n}{\sqrt{v_N}} = \sqrt{t}Z(t). \tag{16}$$

More generally, let $t_0 = 0, t_1 = n_1/N, \dots, t_k = n_k/N$ and let $B(t_0) = 0, B(t_1) = S_{n_1}/v_N^{1/2}, \dots, B(t_k) = S_{n_k}/v_N^{1/2}$ be interim B -values at trial fractions t_0, \dots, t_k . The increments $B(t_1) - B(t_0) = S_{n_1}/v_N^{1/2}, B(t_2) - B(t_1) = (S_{n_2} - S_{n_1})/v_N^{1/2}, \dots$ are independent as they do not involve overlapping sums. Equation (16) implies that

$$\text{Var}(B(t)) = t \times \text{Var}(Z(t)) = t.$$

The distribution of $B(t)$ has the following properties:

- $B(t_1), B(t_2), B(t_3), \dots$ have multivariate normal distribution
- $\mathbb{E}(B(t)) = 0$
- $\text{Cov}(B(t_i), B(t_j)) = t$ for $t_i \leq t_j$, where $t_j - t_i = t$.

$B(t)$ is defined at trial fractions $t = 0, 1/N, \dots, N/N$. If we take $t = \kappa(i/N) + (1 - \kappa)((i + 1)/N)$, we define $B(t)$ to be $\kappa B(i/N) + (1 - \kappa)B((i + 1)/N)$. This makes $B(t)$ continuous but not differentiable at the points $t = 0, 1, \dots, N/N$. As $N \rightarrow \infty$, the set t at which $B(t)$ is non-differentiable becomes more and more ‘dense’. In the limit, we get a standard Brownian motion $B(t) \sim N(0, t)$.

Appendix B: Numerical Methods for Solving the Decision Triggers

The decision triggers π_A and π_I are, together with two constants A and B , obtained by numerically solving the following set of equations (cf., [Thijssen & Bregantini, 2017](#)):

$$\begin{aligned} A\hat{\phi}(\pi_I) + B\check{\phi}(\pi_I) - F_I(\pi_I) - c/r &= 0 \\ A\hat{\phi}'(\pi_I) + B\check{\phi}'(\pi_I) - F'_I(\pi_I) &= 0 \\ A\hat{\phi}(\pi_A) + B\check{\phi}(\pi_A) - F_A(\pi_A) - c/r &= 0 \\ A\hat{\phi}'(\pi_A) + B\check{\phi}'(\pi_A) - F'_A(\pi_A) &= 0, \end{aligned} \tag{17}$$

where

$$\hat{\phi}(\pi) = \pi^{\frac{1}{2}+\gamma}(1 - \pi)^{\frac{1}{2}-\gamma}, \quad \check{\phi}(\pi) = \pi^{\frac{1}{2}-\gamma}(1 - \pi)^{\frac{1}{2}+\gamma},$$

and

$$\gamma = \frac{1}{2} \sqrt{1 + 4r \left(\frac{\sigma}{\mu} \right)}.$$

The value function is then given by

$$F^*(\pi) = \begin{cases} F_A(\pi) & \text{if } \pi \leq \pi_A \\ -\frac{c}{r} + A\hat{\phi}(\pi) + B\check{\phi}(\pi) & \text{if } \pi_A < \pi < \pi_I \\ F_I(\pi) & \text{if } \pi \geq \pi_I. \end{cases}$$

