
Article:
Thijssen, Jacco Johan Jacob orcid.org/0000-0001-6207-5647 and Bregantini, Daniele orcid.org/0000-0001-6802-1551 (2017) Costly Sequential Experimentation and Project Valuation with an Application to Health Technology Assessment. Journal of Economic Dynamics and Control. pp. 202-229. ISSN 0165-1889

https://doi.org/10.1016/j.jedc.2017.01.016

Costly Sequential Experimentation and Project Valuation with an Application to Health Technology Assessment

Jacco J.J. Thijssen
The York Management School
University of York
Freboys Lane
Heslington, York YO10 5GD, UK
Email: jacco.thijssen@york.ac.uk

Daniele Bregantini
University of Liverpool Management School
University of Liverpool
Chatham Street
Liverpool L69 7ZH, UK
Email: D.Bregantini@liverpool.ac.uk

December 16, 2016

Abstract

We study the optimal investment/abandonment decision for a project, where costly sequential experimentation provides information about its true profitability. We derive the optimal decision rule by appropriately extending the Bayesian framework of sequential hypothesis testing. The optimal decision time takes the form of the first exit time of a particular inaction region. We find that increased noise in the observations lowers the value of the project, and that the effect on the expected time at which a decision is taken is ambiguous. Delays in observations affect both project value and the inaction region. The model is illustrated with a health technology assessment application using data on standard versus robot-assisted laporascopic prostatectomy.

Keywords: Sequential experimentation, Optimal stopping, Health technology assessment

1 Introduction

Since its introduction by Myers (1977) and its textbook treatment by Dixit and Pindyck (1994), the real options approach has been used in many fields of economics, finance, and management to understand (investment) timing decisions and project valuation. The approach recognizes that most investment projects are (partly) irreversible and subject to uncertainty. In this paper we focus on a decision maker (DM) who knows the profitability in various states of the world, but faces uncertainty about this state. The uncertainty is resolved gradually through sequential experimentation in continuous time. Profitability in the "good" state of the world gives an incentive to invest if the DM is sufficiently sure that this state of the world prevails. The costs of sequential experimentation give an incentive to abandon the project if the DM is sufficiently sure that the "bad" state of the world prevails. Throughout, to keep matters concrete, we apply our model to a typical health technology assessment (HTA) problem that relies on information obtained from a clinical trial.

The paper extends the simplest model of sequential hypothesis testing with sampling costs, first developed by Wald (1947), analysed in a Bayesian set-up by Mikhalevich (1958) and further developed by Shiryaev (1967)[1], to a decision-theoretic setting that allows for comparisons with standard real options models. This is done by extending the Shiryaev (1967) model to a case where discounted expected payoffs are maximized, rather than non-discounted statistical losses minimized.[2] It turns out that these extensions are simple enough to allow for analytical results, while being rich enough to allow for practically relevant and economically interesting conclusions.

Sequential experimentation is modeled by assuming that the DM observes an arithmetic Brownian motion with either a positive or zero trend; these are the two states of nature. The observations provide (costly) information about these states of nature.[3] The DM wishes to find the optimal time to invest in or abandon an investment project. Optimality here is taken to mean maximization of the discounted expected value of the project – the revenues and costs of which depend on the state of nature in such a way that its net present value (NPV) is positive (negative) in case the trend is positive (zero) – net of the costs of observations.

Like Shiryaev (1967) we can prove that there exists a unique solution to this problem (Sec-

[1] See also Shiryaev (1978, Section 4.2).

[2] The Shiryaev (1967) model does not allow for discounting.

[3] We argue in Section 3 that this set-up can be seen as the continuous-time limit of a simple sequentially conducted clinical trial.

tion 3) that depends only on the process measuring the posterior belief in the trend being positive. This solution consists of a connected "inaction region" (Stokey, 2009) in the interior of the unit interval. The upper boundary of this region represents the trigger beyond which investment is optimal, whereas the lower boundary represents the trigger below which abandonment is optimal. We also obtain analytical results on comparative statics and the expected time until a decision is taken (Section 4). In addition, we extend the model to include measurement delays (Section 5). Such delays are of particular relevance to the case of HTA, where health benefits of treatments are typically only observed after significant periods of time.

We find that, first, the investment (abandonment) trigger is decreasing (increasing) in the volatility of the observations and that the value of the project decreases in volatility. These findings are the exact opposite of standard findings in the real options literature (see Dixit and Pindyck, 1994, Chapter 5 for a standard model). This is related to a difference in the source of uncertainty. In the standard real options approach the costs and/or benefits of a project are assumed to be uncertain. By setting a higher adoption threshold and, thus, typically, waiting longer for more information the decision-maker can lower the probability of erroneously adopting or abandoning the project. As a consequence higher uncertainty implies wider adoption/abandonment bounds and a higher (option) value for the project. In our model, uncertainty enters as noise in the (clinical trial) observations. There is no uncertainty over benefits and costs conditional on the true state of nature. Rather, the uncertainty is about the likelihood of each state of the world. This means that if uncertainty increases (i.e. the observations become more noisy), then the value of waiting for more information decreases. Consequently, the option value of the project is lower. We show, numerically, that this does not necessarily mean that decisions are, in expectation, taken sooner.

In addition, delays in observations have important consequences, both for the value of the project as well as for the decision bounds. Observational delays are very important in practice, and particularly in clinical trials where health benefits of treatments can take years to be fully known. In Section 5, we derive an explicit expression for the value function and show, numerically, that the inaction region shrinks if the delay gets longer. In fact, we show that there may be cases where the optimal decision rule is to take an immediate decision, irrespective of the prior. The reason for this result is that, with delay, the present value of the trial's cost before the first observation becomes available can outweigh the (more heavily discounted) informational benefits of the trial.

2 Related literature

Our paper is closely related to Kwon and Lipmann (2011), who use a similar stochastic set-up to analyse a somewhat simpler decision theoretic model than that studied here. We extend their results in several directions. First, we generalize the payoff structure of their model in two ways. Whereas they study a "new market" model, our model allows for an upgrading of an existing technology. In addition, our model can deal with both a decision theoretic as well as a purely statistical approach. The latter is still often used in the literature on clinical trials, so it is important to show that such a model is a special case of our more general set-up. At a more technical level, we obtain a slightly stronger result in that we prove uniqueness, rather than mere existence, of the investment/abandonment triggers. Secondly, in contrast to our model, Kwon and Lipmann (2011) assume that the observation of signals is costless. Thirdly, we provide more detailed analytical results on comparative statics. Finally, we extend their model to study the case, important in fields like health technology assessment, where results of experiments are observed with a delay.

Another closely related paper is Moscarini and Smith (2001), who model an experiment the volatility of which can be continuously influenced by the DM. This is modeled by assuming that at any time t the DM can decide how many observations to draw. In our model we have fixed this number to one. The main advantage of this lack of flexibility is that we are able to derive analytical comparative statics results. Also, in clinical trials the experimenter usually does not have the freedom to choose how long treatment per patient is going to take, so that the assumption of a fixed number of observations over a given time interval is not unrealistic.

Sequential hypothesis testing and bandit problems

There is a close link between the learning that happens in models of sequential hypothesis testing and multi-arm bandits. The latter are used in situations where agents try to optimise their payoffs while simultaneously gathering more information about a risky alternative to a safe option. Learning only takes place when the risky option is chosen. These models, therefore, analyse the trade-off between maximizing the current reward and the value of getting more precise information about the future reward. Typically, the decision maker can choose between using a "safe" arm, or a "risky" arm. By playing the safe arm, the agent receives an immediate payoff (initially usually lower in expectation than when playing the safe arm), but, as the process continues, the payoffs of the risky arm evolve and the player gathers more information about the distribution of returns from the risky arm. In the economics literature, bandits have been used to model search

processes (see Rothschild (1974) and Bergemann and Valimaki (2006) for an overview), learning and matching in markets such as the labour and consumer markets (see Jovanovic (1979), Sundaram (2005) and Arlotto et al. (2014)) and agent experimentation (see Bolton and Harris (1999), Keller et al. (2005), (Keller and Rady, 2010) and Heiodhues et al. (2015)).

An example of a typical bandit model is provided by Bolton and Harris (1999), who analyse a game where players have to determine, continuously, the intensity with which a safe and risky arm are used. In fact, our informational set-up is similar to theirs and their cooperative benchmark is quite close to our model. The main differences are that the decision-maker in our model does not get any payoffs until a decision about which arm to use is chosen (here choosing the safe arm could be interpreted as abandoning the project), that experimentation is costly, and that the decision to stop sampling involves sunk costs. In addition, the model by Bolton and Harris (1999) allows the decision-maker to continuously adjust the intensity with which each arm is played. In that sense, their model is more like an optimal portfolio problem, rather than an optimal stopping problem. So, while the modeling of learning in bandit problems is closely related to that of sequential hypothesis testing, the focus and, hence, the mathematical tools to analyse such problems are substantially different.[4] In particular, dynamic programming models, where the goal is the continuous optimization of a certain performance criterion flow, are solved by selecting an appropriate feedback policy (and corresponding value function). Optimal stopping, in contrast, deals with selecting a stopping time to achieve a set purpose. Optimal stopping arises naturally in situations when the timing of the decision is crucial (e.g., determining when it is optimal to sell a stock). The "bang-bang" solution, where a trigger determines the switch from one arm to the other, obtained in the cooperative benchmark of Bolton and Harris (1999) is the optimal decision rule in a large class of feedback rules. In our model, the decision rule *must* be a stopping time. In summary, the results of bandit and optimal stopping problems may look very similar (i.e., the DM acts when a certain trigger is reached), the DM's objective is different and, thus, is the solution method.

In a clinical trial setting, the Bolton and Harris (1999) model could be interpreted as follows. At each point in time a patient enters the trial and the decision maker uses a mixed strategy to determine if the patient gets the placebo (safe arm) or the new drug (risky arm). The crucial difference between this set-up and ours is that the mixed strategy of the decision-maker changes

[4]Keller et al. (2005) and Keller and Rady (2010), in a similar set-up to Bolton and Harris (1999), study a game of strategic experimentation where the risky arm distributes lump-sum payoffs according to a continuous time Poisson process, rather than a model based on Brownian motion.

over time in light of the evidence. Studying clinical trials in this way moves quite far from current practice, as the objective of the Bolton and Harris (1999) model is to maximize expected payoffs during the trial, rather than trying to learn as quickly as possible what decision (adoption or abandonment) should be taken. In particular, the Bolton and Harris (1999) could not be implemented alongside current or past clinical trials, whereas our approach can.

In the clinical trial interpretation of our model, we stay closer to real-world trials where two patients enter the trial at each time, one being given the placebo and one being given the new drug. Experimentation is then undertaken continuously until the optimal stopping time. It is only at this point that the decision maker chooses either for abandonment or adoption of the technology and obtain some pre-specified payoff. The abandonment option is typically not analysed in the bandit literature. An important feature of our model is that the adoption and abandonment bounds have to be determined simultaneously, which implies that the option values for abandonment and adoption are interdependent. We show in a comparative static exercise (Section 4) how the bounds vary when changing a number of parameters. It turns out that varying parameters that directly drive the statistical evidence process (e.g., volatility or mean expected net benefit of the technology) brings about symmetric changes in the decision bounds in the sense that they both widen or both shrink the inaction region. This is due to the fact that both adoption and abandonment payoffs are functions of the posterior process. On the other hand, we show that varying parameters that enter only the payoffs, but not the evidential process (e.g., the expected benefit from adoption if the technology is superior), moves the decision bounds in the same direction, i.e., they both increase or they both decrease. The effect on the inaction region is then ambiguous. This is to be expected as the balance of the adoption and rejection payoffs is changing in favour of one over the other. So, our model provides some novel insights that may also be of interest to the bandit literature.

Health technology assessment and sequential analysis

While our model can be applied to many investment projects with costly experimentation we choose to focus in this paper on clinical trials. One advantage is that it makes it easier to see the practical relevance of the model. Secondly, sequential methods, in the form of group-sequential analysis have long been applied to the field of medical statistics, albeit typically in discrete time.[5] Recently, due to the escalating costs of drug development and testing, there has been great interest in early stopping in clinical trials with the Food and Drug Administration (FDA) issuing

[5]Among the many references, only the seminal work by Armitage (1975) and Berry (1985) are mentioned here.

non-binding guidance for such study designs (FDA, 2010). Additionally, clinical trials evaluating medicines, medical devices and procedures have also seen, due to changes in regulatory reimbursement requirements of many health care systems, the explicit assessment of economic value of these interventions (Ramsey et al., 2005).

While randomized clinical trials are traditionally considered to be the gold standard for determining the safety and efficacy of health-care interventions (Spiegelhalter et al., 2004) and their outcomes largely determine whether new health-care technologies are approved by regulatory agencies, governments have become increasingly concerned about the value for money of such technologies. In fact, even if a technology displays superior efficacy in a controlled environment such as a clinical trial, there still remains the question of its cost-effectiveness. In this context, the example of health technology assessment is topical. In particular, especially for those interventions that display immediate effects, it is now more common for cost-effectiveness evaluations to be conducted alongside clinical trials (Meltzer and Smith, 2012)

In clinical research, typical adoption/abandonment/postponenment decisions take place after enough evidence has been provided to pass a purely inferential evidence threshold (usually stated, explicitly or implicitly, in terms of a p-value). However, such traditional sample size calculations tend to be based on arbitrary rules of inference, typically Type I and Type II error probabilities, and do not reflect the *cost* of making such errors (William and Pinto, 2005). Recently, Montazerhodjat and Lo (2015) emphasised that the traditional approach to inference can bring about thresholds that are too conservative and lead to "overly large" trials (Berry, 2006) and long approval processes for drugs intended to treat serious conditions (FDA, 2006, 2013). They suggest to incorporate the patient's perspective when calculating trials' sample sizes and critical values for a fixed-sample design. In particular, in their Bayesian framework emphasis is given to the number of patients affected by a disease and disease severity.

Motivated by the recent changes in health-care technologies approval requirements and increased criticism of standard inferential rules for clinical trials' design, our model seeks to combine the statistical evidence gathered during clinical research with cost-effectiveness analysis in order to obtain optimal economic decisions. Since the decision rule (implicitly) used in traditional inference is a special case of the decision rules that we allow, our approach is superior in terms of expected net benefits. However, while our model is mainly aimed at providing a dynamic cost-benefit analysis, it is rich enough to model purely inferential concerns of the type mentioned above.

Our decision framework is linked to a recent line of research that argues for the explicit in-

clusion of uncertainty and sunk costs. Palmer and Smith (2000), for example, apply the Dixit and Pindyck (1994) real options approach to the adjustment (under a certain degree of irreversibility and uncertainty) of the incremental cost-effectiveness ratio for a drug. Driffield and Smith (2007) use real options to argue for a watchful waiting regime for diseases with slow progression. Forster and Pertile (2013) appeal to real options to argue that flexibility and irreversibility of actions should play a much bigger role in HTA than they are hitherto assigned. However, presently the real options approach has not been implemented in any systematic way (Meltzer and Smith, 2012).

Closely related to our work is Pertile et al. (2014), who 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. Their approach, however, is one of sequential estimation, whereas ours is one of sequential hypothesis testing. The advantage of our hypothesis testing approach is that it "dichotomizes" the decision (Draper, 2013). While Draper (2013) argues that this is unnecessary, it is precisely this assumption that allows us to derive many analytical properties, even in a more complicated case with delayed observations. Papers without this feature, such as, for example, Hampson and Jennison (2013) or Chick et al. (2015), typically have to resort to numerical solutions. In short, what we lose in realism we gain in analytical clarity.

3 The Model and Main Result

Consider a decision-maker (DM) who is observing costly signals in continuous time about the profitability of an investment project. The profitability depends on the state of the world, θ, which can be "good" ($\theta = 1$) or "bad" ($\theta = 0$). To make matters concrete, we model a DM who has to value a new health technology and observes the outcomes of an ongoing clinical trial in which patients are treated sequentially. In order to make a decision about adopting or abandoning the new technology, the DM needs to carefully consider the benefits and costs of each decision, conditional on the events $\{\theta = 1\}$ and $\{\theta = 0\}$. We require that benefits and and costs are all measured in monetary units.[6] We introduce the following quantities, all of which are assumed to be known *ex ante* by the DM:

- $B_1 > 0$: the net benefit of the new health technology if $\theta = 1$;
- $B_0 \geq 0$: the net benefit of the new health technology if $\theta = 0$;

[6] In the UK, for example, health benefits are usually measured in quality adjusted life years (QALYs), which allows for comparison between different health technologies.

- $B \geq 0$: the net benefit of the existing health technology;
- $I > 0$: the sunk costs of investing in the new health technology;
- $c > 0$: per period cost of running the clinical trial;
- $L \geq 0$: the loss of not using the new technology if $\theta = 1$;
- $P \geq 0$: per period penalty, incurred during the trial, of not using the new health technology if $\theta = 1$.

Remark 1. *In a purely economic decision theoretic framework, where profits and losses, rather than statistical errors are the focus of the analysis, we will typically set $P = L = 0$. However, our framework can also be used in a purely statistical setting. In that case, I represents the loss of a Type I error and L denotes the loss of a Type II error. If one sets $B_0 = B = P = 0$ and $B_1 = I$, then a standard (Bayesian) statistical decision theoretic model appears.*

Remark 2. *The last term, P, represents a penalty for not using a new health technology during the trial that is, in fact, better than the existing technology. It has been argued in the literature (Palmer and Smith, 2000) that this should be taken into account explicitly in health technology assessment, because payoffs in an HTA exercise are not just reflecting "profits" and "losses", but represent human lives. Any patient who is outside the trial (and those in the trial but not enrolled in the arm with the new technology) effectively misses out on the incremental benefits related to the new technology (conditional on it being superior) leading to a loss to the heath-care system. This loss cumulates until the new technology is adopted (and will continue to accumulate if a superior technology is abandoned). A recent report explicitly mentions that "[d]elayed adoption leads to loss of health benefits to patients who might have benefited from earlier access to technologies which prove to be cost-effective" (YHEC, 2009, p. 5). As this is a controversial issue and no consensus exists in the literature, we will typically set $P = 0$ in numerical examples. The payoffs B_1, B, B, and L will be in terms of discounted present values over the life-time of the technology being used on the entire population of patients.*

In terms of parameter restrictions it is appropriate to assume that $B_0 - I < B \leq B_1 - I$, so that the new technology outperforms the current technology if, and only if, the state of the world is $\theta = 1$. Table 1 summarises the payoffs of investment and abandonment under different states of the world.

		State of the world	
		$\theta = 1$	$\theta = 0$
Decision	Adoption	$B_1 - I$	$B_0 - I$
	Abandonment	$B - L$	B

Table 1: Payoff matrix of a health technology decision problem.

If the DM takes a decision at time t, when the (posterior) belief in the event $\{\theta = 1\}$ is $\pi_t \in (0,1)$, the expected net present value of adoption/investment, denoted by F_I, is

$$\begin{aligned} F_I(\pi_t) :&= \pi_t(B_1 - I) + (1 - \pi_t)(B_0 - I) \\ &= \pi_t(B_1 - B_0) + B_0 - I. \end{aligned} \tag{1}$$

Similarly, if at time t the decision maker decides to abandon the new health technology, then the expected net present value is

$$\begin{aligned} F_A(\pi_t) :&= \pi_t(B - L) + (1 - \pi_t)B \\ &= B - \pi_t L. \end{aligned} \tag{2}$$

So, at time t the DM will choose to invest if, and only if,

$$\pi_t > \bar{\pi} \equiv \frac{B + I - B_0}{B_1 + L - B_0}. \tag{3}$$

Note that $B_1 > B + I > B_0$ implies that $\bar{\pi} > 0$. The problem is only interesting if $\bar{\pi} < 1$ (otherwise investment is never optimal), i.e. if

$$B_1 - I > B - L.$$

That is, the net payoff of investing in the new health technology if $\theta = 1$ must exceed the payoff of the current treatment net of the loss that is incurred if $\theta = 1$ and the new technology is not used.

Until a decision is taken, the decision maker is assumed to incur a per-period cost, $c > 0$, for conducting a clinical trial and a per-period penalty, P, for not using the new technology in case $\theta = 1$. Assuming that all payoffs and costs are discounted at a constant rate $r > 0$, the decision maker therefore needs to find a stopping time τ^* that solves the optimal stopping problem

$$\begin{aligned} F^*(\pi) &= \sup_{\tau \in \mathcal{T}} \mathsf{E}_\pi \left[-\int_0^\tau e^{-rt}(c + \pi_t P)dt + e^{-r\tau} \max\{F_I(\pi_\tau), F_A(\pi_\tau)\} \right] \\ &= \mathsf{E}_\pi \left[-\int_0^{\tau^*} e^{-rt}(c + \pi_t P)dt + e^{-r\tau^*} \max\{F_I(\pi_{\tau^*}), F_A(\pi_{\tau^*})\} \right], \end{aligned} \tag{4}$$

where $\mathcal{T}$ is the set of all stopping times with respect to an appropriate filtration defined below.

Remark 3. *The penalty of not using a superior technology during the running of the clinical trial is here assumed to be a linear function of the posterior belief in a superior technology. This restriction to linearity is not necessary. For the proof of Proposition 1 below, a more general penalty function, $P : (0,1) \to \mathbb{R}_+$, which is increasing, convex and satisfies $P(1) > 0$, would be valid. For the model with delays in measurement that is discussed in Section 5, however, linearity makes the argument much simpler. For that reason we will restrict ourselves to a linear penalty function in the remainder.*

The clinical trial provides information about the true state of nature and is used to obtain the posterior process $(\pi_t)_{t\geq 0}$. We view the optimal stopping problem (4) as one of Bayesian sequential testing of two simple hypotheses in continuous time. We use the formalization by Shiryaev (1978) as the starting point of our model. As in his setup we assume that uncertainty is modeled on a probability space $(\Omega, \mathscr{F}, \mathsf{P}_p)$ for a family of probability measures $(\mathsf{P}_p)_{p\in[0,1]}$, and that, for fixed $p \in [0,1]$, we are given a random variable $\tilde{\theta}$, representing the true state of the world and taking values in $\{0,1\}$.[7] For $p \in (0,1)$, the probability measure P_p is obtained as

$$\mathsf{P}_p = p\mathsf{P}_1 + (1-p)\mathsf{P}_0,$$

where P_1 and P_0 are degenerate distributions with $\mathsf{P}_1(\tilde{\theta}=1) = \mathsf{P}_0(\tilde{\theta}=0) = 1$.

On this probability space we construct a continuous-time model of a sequential clinical trial. Before introducing the model in continuous time, we build one in discrete time, to illustrate the main ideas. Consider a time interval $[0,t]$, for some fixed $t > 0$. During this time, assume that we observe the outcomes of the trial sequentially at n equally spaced occasions $dt = t/n$, i.e., at times $\{0 = t_0, t_1, \ldots, t_n = t\}$. The outcome of a clinical trial at time t_i is measured in terms of the *cumulative benefit* (often measured relative to some existing technology) to the patients in the trial and is denoted by X_i. We assume that observations consist of a signal and a noise component and that both components depend on the time interval between observations. This reflects the idea that more patients will be treated during longer time intervals. To keep things as simple as possible, the noise component is assumed to be binary, which implies that between times t_{i-1} and t_i the observed cumulative benefit either goes up by a factor U, or down by a factor D, both with equal probability. In particular, it is assumed that

$$U = \theta\mu dt + \sigma\sqrt{dt}, \quad \text{and} \quad D = \theta\mu dt - \sigma\sqrt{dt}.$$

Here σ measures the magnitude of the noise component, and μ measures the expected cumulative benefit per unit of time conditional on the new technology being superior. The total accumulated benefit at time t after n steps is equal to $X^{(n)}(t) := \sum_{i=1}^{n} X_i$.

[7] Since we use a Bayesian framework, the parameter θ is assumed to be the realization of a random variable $\tilde{\theta}$.

In order to obtain analytical results, we move to the continuous-time limit of this simple model. For fixed t, we do this by increasing the number of time steps at which evidence is observed. To that effect, define $X_t := \lim_{n\to\infty} X^{(n)}(t)$, where the limit is understood to be in distribution.[8] According to the central limit theorem (CLT), the conditional distribution of X_t is well-defined and is given by

$$X_t|\theta \sim \mathsf{N}(\theta\mu t, \sigma^2 t).$$

The only continuous-time stochastic process $X := (X_t)_{t\geq 0}$, that gives these distributions at each finite time t is the arithmetic Brownian motion (ABM)

$$X_t = \theta\mu t + \sigma B_t.$$

Here $(B_t)_{t\geq 0}$ is a standard P_p-Brownian motion, and p and $1-p$ play the role of *prior probabilities* of the statistical hypotheses

$$H_0 : \theta = 0, \quad \text{and} \quad H_1 : \theta = 1,$$

respectively.

Remark 4. *Applying the CLT in deriving the limit distribution of X_t only gives point-wise convergence for fixed t. A fairly straightforward application of Donsker's Invariance Principle (see, for example, Steele, 2001, Theorem 5.4) shows that for all $x \in \mathbb{R}$,*

$$\lim_{n\to\infty} P\Big(X^{(n)}(\cdot) \leq x\Big) = P(X(\cdot) \leq x).$$

Apart from theoretical importance, this result is of practical interest as it shows that the continuous-time model is a limiting case of a discrete-time model, the latter more likely being an accurate description in real-world applications.

Since we interpret the observations X as being obtained from a clinical trial, the parameter μ represents the cumulative health benefits of the new technology. In many cases, this will be a benefit relative to some existing technology that is to be expected in one time unit of trials conditional on the new health technology being superior. What is understood by "superior" can depend on the particular application. For example, one could compute the per patient break-even incremental benefit of the new health technology over the best existing treatment net of its additional cost. The parameter μ then equals this net benefit multiplied by the number of patients treated in the trial during a year. The parameter σ is a measure of the standard deviation due to

[8] Note that $n \to \infty$ implies that $dt \downarrow 0$.

random effects on health benefits from patient to patient. Another approach would be to apply the concept of "clinically important differences" as introduced by Samsa et al. (1999).

The process $(X_t)_{t\geq 0}$ generates the filtration $\mathscr{F}^X = (\mathscr{F}_t^X)_{t\geq 0}$, which is augmented with the P_p-null sets. The set of $\mathscr{F}^X$-stopping times is denoted by $\mathcal{T}$. For all $t \geq 0$, the Radon-Nikodym derivative

$$\Lambda_t := \frac{d(\mathsf{P}_1 \mid \mathscr{F}_t^X)}{d(\mathsf{P}_0 \mid \mathscr{F}_t^X)},$$

defines the *likelihood ratio* process $(\Lambda_t)_{t\geq 0}$. Shiryaev (1978) shows that the process $(\Lambda_t)_{t\geq 0}$ admits the representation

$$\Lambda_t = \exp\left(\frac{\mu}{\sigma^2}\left(X_t - \frac{\mu}{2}t\right)\right), \quad t \geq 0.^9$$

The *posterior process* $(\pi_t)_{t\geq 0}$, with $\pi_t = \mathsf{P}_p(\tilde{\theta} = 1 \mid \mathscr{F}_t^X)$, can be obtained as a function of the likelihood ratio process using Bayes' rule:

$$\pi_t = \left(\frac{p}{1-p}\Lambda_t\right) \Big/ \left(1 + \frac{p}{1-p}\Lambda_t\right).$$

From Ito's lemma it then immediately follows that $(\pi_t)_{t\geq 0}$ solves the stochastic differential equation

$$d\pi_t = \frac{\mu}{\sigma}\pi_t(1-\pi_t)d\bar{B}_t, \quad \text{with } \pi_0 = p, \mathsf{P}_p\text{-a.s.}, \tag{5}$$

where $\left(\bar{B}_t\right)_{t\geq 0}$, with

$$\bar{B}_t := \frac{1}{\sigma}\left(X_t - \mu\int_0^t \pi_s ds\right), \tag{6}$$

is a standard P_p-Brownian motion, called the *innovation process*. The process $(\pi_t)_{t\geq 0}$ is time-homogeneous and strongly Markovian under P_p. Note that, since $\mathscr{F}^\pi = \mathscr{F}^X$, the set of $\mathscr{F}^\pi$-stopping times is $\mathcal{T}$, so that problem (4) is well-defined.

With this sequential model of a clinical trial in hand we can solve the optimal stopping problem (4). In Proposition 1 below, we show that the state space $(0,1)$ can be split into three regions. The first one is a region around $\bar{\pi}$ where continuation of the trial is optimal, thence called the *continuation region*, and denoted by

$$\mathcal{C} = \{\pi \in (0,1) | F^*(\pi) > \max(F_A(\pi), F_I(\pi))\} = (\pi_A, \pi_I),$$

where π_A and π_I, with $0 < \pi_A < \bar{\pi} < \pi_I < 1$, are the abandonment and investment triggers, respectively. When π gets large enough we enter the *investment region*, where adoption of the health-care

[9] From Ito's lemma it follows that, conditional on θ, the process $(\Lambda_t)_{t\geq 0}$ follows the geometric Brownian motion (GBM) $\frac{d\Lambda}{\Lambda} = \theta\frac{\mu^2}{\sigma^2}dt + \frac{\mu}{\sigma}dB$.

technology is optimal. This region is denoted by

$$\mathcal{D}_I = \{\pi \in (0,1) | F^*(\pi) = F_I(\pi)\} = [\pi_I, 1).$$

Conversely, when π gets low enough, we enter the *abandonment region*, where abandoning the clinical trial is optimal. This region is denoted by

$$\mathcal{D}_A = \{\pi \in (0,1) | F^*(\pi) = F_A(\pi)\} = (0, \pi_A].$$

In order to derive the value function F^* in (4), henceforth called the *value of the new technology*, we introduce the parameter

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

and, for given $0 < \pi_A < \pi_I < 1$, the functions $\hat{v}_{\pi_A,\pi_I} : [\pi_A, \pi_I] \to [0,1]$ and $\check{v}_{\pi_A,\pi_I} : [\pi_A, \pi_I] \to [0,1]$, defined by

$$\hat{v}_{\pi_A,\pi_I}(\pi) := \sqrt{\frac{\pi(1-\pi)}{\pi_I(1-\pi_I)}} \frac{\left(\frac{1-\pi_A}{\pi_A}\frac{\pi}{1-\pi}\right)^\gamma - \left(\frac{\pi_A}{1-\pi_A}\frac{1-\pi}{\pi}\right)^\gamma}{\left(\frac{1-\pi_A}{\pi_A}\frac{\pi_I}{1-\pi_I}\right)^\gamma - \left(\frac{\pi_A}{1-\pi_A}\frac{1-\pi_I}{\pi_I}\right)^\gamma},$$

and

$$\check{v}_{\pi_A,\pi_I}(\pi) := \sqrt{\frac{\pi(1-\pi)}{\pi_A(1-\pi_A)}} \frac{\left(\frac{1-\pi}{\pi}\frac{\pi_I}{1-\pi_I}\right)^\gamma - \left(\frac{\pi}{1-\pi}\frac{1-\pi_I}{\pi_I}\right)^\gamma}{\left(\frac{1-\pi_A}{\pi_A}\frac{\pi_I}{1-\pi_I}\right)^\gamma - \left(\frac{\pi_A}{1-\pi_A}\frac{1-\pi_I}{\pi_I}\right)^\gamma},$$

respectively. In Appendix A we prove the following proposition.

Proposition 1. *Suppose that*

1. $B_1 - B_0 + P/r > 0$,
2. $L - P/r \geq 0$,
3. $B > B_0 - I$,
4. $B_1 - I > B - L$, *and*
5. $B + c/r > \bar{\pi}\frac{\gamma+1/2}{\bar{\pi}+\gamma-1/2}(L - P/r)$.

Then there exist unique thresholds π_A and π_I, $\pi_A < \bar{\pi} < \pi_I$, such that the optimal stopping time is the first exit time of (π_A, π_I), i.e.

$$\tau^* = \inf\{t \geq 0 | \pi_t \notin (\pi_A, \pi_I)\}.$$

In addition, the value of the new health technology, if the posterior belief in $\{\tilde{\theta} = 1\}$ *is* π*, equals*

$$F^*(\pi) = \begin{cases} F_A(\pi) & \text{if } \pi \leq \pi_A \\ -\left(\frac{c}{r} + \pi\frac{P}{r}\right) + \hat{\nu}_{\pi_A,\pi_I}(\pi)\left(F_I(\pi_I) + \frac{c}{r} + \pi_I\frac{P}{r}\right) & \\ \qquad + \check{\nu}_{\pi_A,\pi_I}(\pi)\left(F_A(\pi_A) + \frac{c}{r} + \pi_A\frac{P}{r}\right) & \text{if } \pi_A < \pi < \pi_I \\ F_I(\pi) & \text{if } \pi \geq \pi_I. \end{cases} \tag{7}$$

The triggers π_A *and* π_I *are the unique values for which* (7) *is* C^1 *for all* $\pi \in (0,1)$*.*

Remark 5. *The first four conditions in Proposition 1 are fairly natural parameter restrictions to ensure that the net present value functions of adoption and abandonment are well-behaved. Condition 5, however, may be too strong for certain applications. It can be relaxed to a condition that is more difficult to check and requires more notation; the proof of Proposition 1 gives more details. Note that, in the fairly natural case where* $L = P/r$*, the final condition is always satisfied.*

For $\pi \in (\pi_A, \pi_I)$, the values $\hat{\nu}_{\pi_A,\pi_I}(\pi)$ and $\check{\nu}_{\pi_A,\pi_I}(\pi)$ are the expected discount factors of first reaching π_I and π_A, respectively, given the current posterior probability π. So, if we denote the first hitting times of π_I and π_A from $\mathcal{C}$ by $\hat{\tau}(\pi_I)$ and $\check{\tau}(\pi_A)$, respectively, then

$$\hat{\nu}_{\pi_A,\pi_I}(\pi) = \mathsf{E}_\pi\left[e^{-r\hat{\tau}(\pi_I)}; \hat{\tau}(\pi_I) < \check{\tau}(\pi_A)\right], \tag{8}$$

and

$$\check{\nu}_{\pi_A,\pi_I}(\pi) = \mathsf{E}_\pi\left[e^{-r\check{\tau}(\pi_A)}; \hat{\tau}(\pi_I) > \check{\tau}(\pi_A)\right]. \tag{9}$$

In the region $\mathcal{C} = (\pi_A, \pi_I)$ the value of the project is, therefore, equal to the cost of running the trial forever and never taking a decision, corrected for the expected net value of adopting or abandoning the health technology at the first time $\mathcal{C}$ is exited, at π_I or π_A, respectively, discounted back to the current time using the expected discount factor.

4 Comparative Statics

In this section we investigate the sensitivity of the solution to (4) with respect to some of the model's parameters.

Proposition 2. *Suppose that the conditions of Proposition 1 hold. Let* $\tau^* = \inf\{t \geq 0 | \pi_t \notin (\pi_A, \pi_I)\}$ *be the unique solution to* (4)*. It then holds that*

$B = 30$	$B_0 = 40$	$B_1 = 60$
$I = 20$	$c = 3$	
$L = 0$	$P = 0$	

(a) Payoffs

$r = .1$	$\mu = 2$	$\sigma = 1$
$p = .5$		

(b) Parameter values

Table 2: Parameter values for a base-case numerical example.

- *π_A is increasing and π_I is decreasing in the cost of the trial, c;*
- *both π_A and π_I are decreasing in the benefits of the technology if superior, B_1;*
- *both π_A and π_I are increasing in the benefits of the current technology, B;*
- *π_A is increasing and π_I is decreasing in the volatility, σ, provided that $\pi_A < 1/2 < \pi_I$;*
- *π_A is decreasing and π_I is increasing in the expected benefits in the trial (conditional on $\theta = 1$), μ, provided that $\pi_A < 1/2 < \pi_I$.*

The proof of this proposition can be found in Appendix B.

In order to assess the quantitative effects on the bounds, the *ex ante* (prior) value of the new technology, $F^*(p)$, the expected time until a decision is made, and the expected costs of the trial, we conduct a numerical analysis. The payoffs and parameters for a base-case scenario are given in Table 2. For this case it turns out that $\pi_A = 19.99\%$ and $\pi_I = 79.44\%$.[10] For different values of the posterior belief in the event $\{\tilde{\theta} = 1\}$ the value of the new technology is depicted in Figure 1a. This figure shows the value of waiting, for beliefs between the thresholds π_A and π_I. Note that this value represents both the upside of the potential benefits if the new technology turns out to be superior as well as the downside of finding out that the new technology is not superior.

For each $\pi \in (\pi_A, \pi_I)$ the probabilities of investment and abandonment can actually be computed explicitly as (see, for example, Poor and Hadjiliadis, 2009)

$$\frac{\pi - \pi_A}{\pi_I - \pi_A}, \quad \text{and} \quad \frac{\pi_I - \pi}{\pi_I - \pi_A},$$

respectively. For the base-case scenario these probabilities of investment and abandonment are depicted in Figure 1b. The investment probability is linearly increasing, because it becomes more likely that π_I is reached before π_A as the posterior belief in $\{\tilde{\theta} = 1\}$ increases. The probability of an abandonment decision, consequently, goes in the other direction.

[10] All calculations are executed in Matlab.

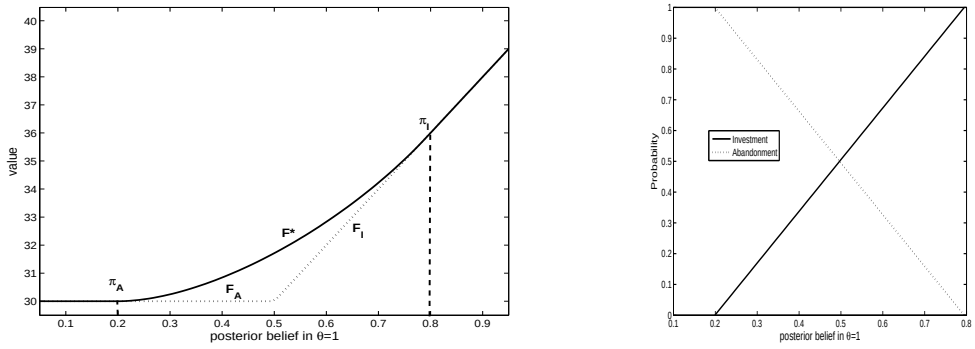

(a) Value of new technology

(b) Decision probabilities

Figure 1: Value of technology and decision probabilities for base-case scenario.

Note that the expected cost of the trial is given by

$$\mathsf{E}_p\left[\int_0^{\tau^*} ce^{-rt}dt\right] = \frac{c}{r}\left(1-\mathsf{E}_p\left[e^{-r\tau^*}\right]\right) = \frac{c}{r}\left(1-\hat{v}_{\pi_A,\pi_I}(p)-\check{v}_{\pi_A,\pi_I}(p)\right).$$

Note that $\mathsf{E}_p\left[e^{-r\tau^*}\right]$ and $\mathsf{E}_p[\tau^*]$ do not necessarily move in the same direction. The latter expectation equals

$$\begin{aligned}\mathsf{E}_\pi[\tau^*] = \frac{2\sigma^2}{\mu^2}\Bigg\{&\log\left[\left(\frac{\pi}{1-\pi}\right)^{1-2\pi}\left(\frac{1-\pi_A}{\pi_A}\right)^{1-2\pi_A}\right]\\ &+\frac{\pi-\pi_A}{\pi_I-\pi_A}\log\left[\left(\frac{\pi_A}{1-\pi_A}\right)^{1-2\pi_A}\left(\frac{1-\pi_I}{\pi_I}\right)^{1-2\pi_I}\right]\Bigg\}.\end{aligned} \tag{10}$$

The derivation of this expectation is standard (see, for example, Poor and Hadjiliadis, 2009).

The comparative statics for the noise in the clinical trial, σ, are given in Figure 2. From a real options perspective, these appear to be counterintuitive as they show that project value is lower for higher levels of noise. This is the opposite effect found in the real options literature, where more uncertainty increases the value. The reason for this difference lies in the nature of the uncertainty. Here, if σ increases, we are paying to keep a trial alive that provides less information. That means that waiting leads to less precise information, which makes waiting less valuable. As the trial is costly, one might as well decide sooner. In other words, the costs of waiting (conducting more trials) do not outweigh the benefits (the extra evidence accruing from the trials), because the trials are less informative. The effect of σ on the expected length of the trial and its expected costs is ambiguous. This happens because of two opposing effects. On the one hand, a higher volatility makes the trial less informative leading to an increasing trial length. On the other hand, the decision bounds narrow, which implies a decision is reached earlier. The relative weight of these effects changes as σ is increased.

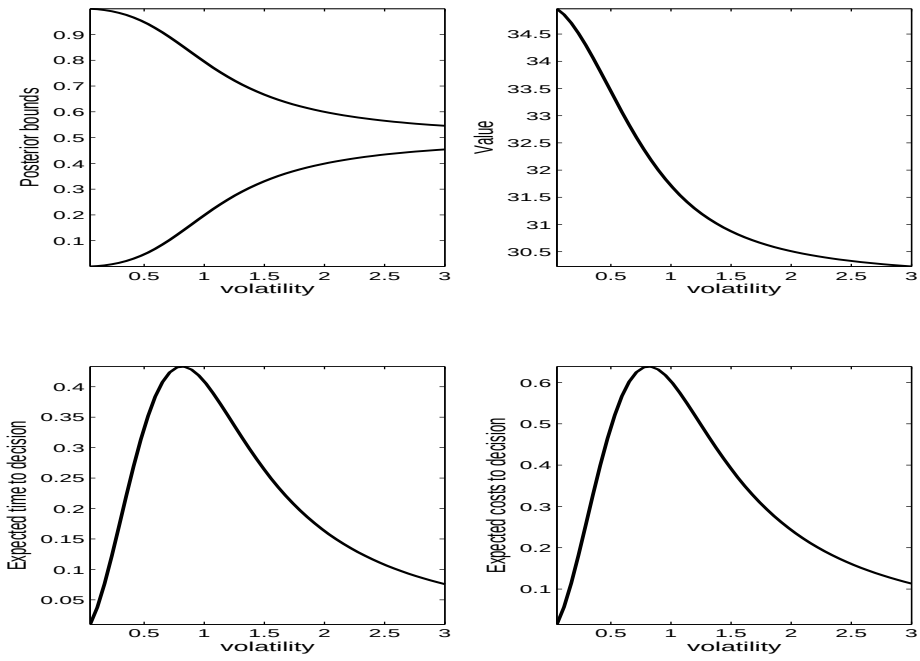

Figure 2: Comparative statics for σ.

Recall that the parameter μ measures the expected net benefit of the new technology in the trial over a unit of time, conditional on it being superior. This implies that μ and B_1 are closely related. In fact, these two parameters should only differ in so far as that B_1 refers to the *population* of treated patients after approval, whereas μ refers to the patients treated in the trial. Here we keep the benefits B_1 constant, so we are, implicitly assuming that we are treating more patients at the same cost in the trial. The comparative statics for the parameter μ are depicted in Figure 3. It should not come as a surprise that the adoption/abandonment bounds get wider, because, *ceteris paribus*, the trial becomes more informative (relative to the noise component). This increases the value of waiting for more information and, thus, the value of the new technology. The non-monotonicity in the expected time until a decision is taken follows from a similar reason as we identified for σ. On the one hand, we get more information per time period, which leads to decisions being taken sooner. On the other hand, because the value of waiting for more information increases, we want to make a decision later (the decision bounds widen).

The influence of the annual cost of conducting the trial, c, are given in Figure 4. These are as to be expected: the higher the costs of the trial the closer the thresholds are together, the lower the value of the technology, and the sooner, on average, a decision is taken. Here we see that only

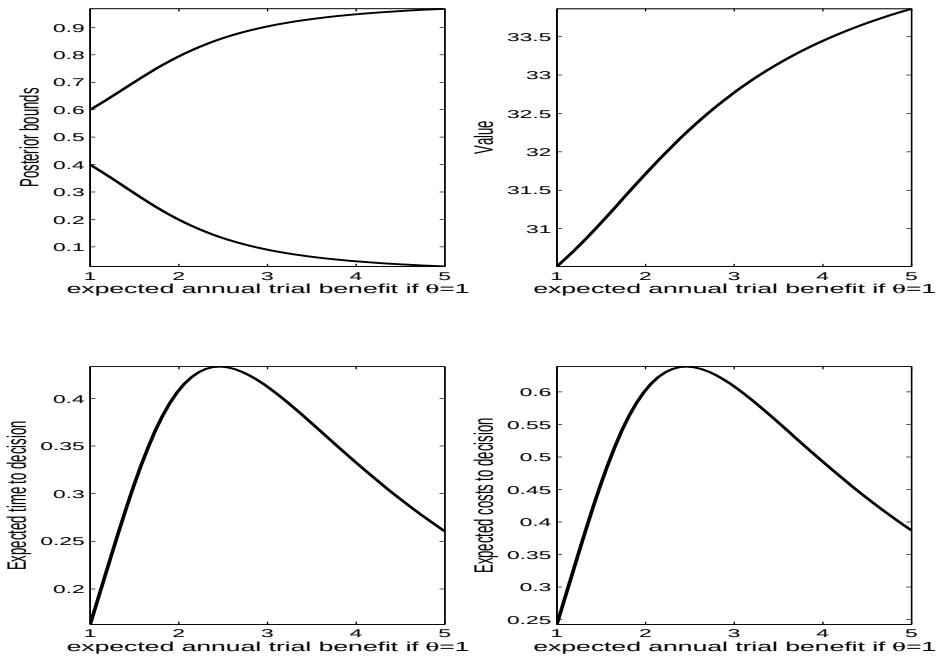

Figure 3: Comparative statics for μ.

the expected costs of the trial are non-monotonic. This happens because there are two opposing factors at work. On the one hand a decision is taken sooner, which lowers the expected costs. On the other hand, the per-period costs increase, which makes the expected costs higher.

The effects of the benefits of the new technology if superior, B_1, are clear-cut and are reported in Figure 5. The decision bounds are decreasing and initially widen, but from $B_1 \approx 72$ start narrowing again. This explains the non-monotonicity of the expected length and costs of trials, because B_1 does not influence the properties of the posterior process.

Even though it is analytically clear that both decision bounds are increasing in the benefits of the existing technology, B, Figure 6 shows that the inaction region first widens and then narrows. Also note that for small values of B a prior of $p = .5$ leads to immediate investment and, hence, no trial being started at all. Since B does not influence the prior expected net-present value of adoption, the value function is flat in this region. Conversely, for large values of B an immediate decision to abandon is taken, simply because the expected net increase in benefits if the new technology turns out to be superior over the the old technology are too small to warrant starting a trial. The non-monotonicities in expected trial length and costs are not surprising in this light: the decision bounds first widen and then narrow and since B does not influence the properties of

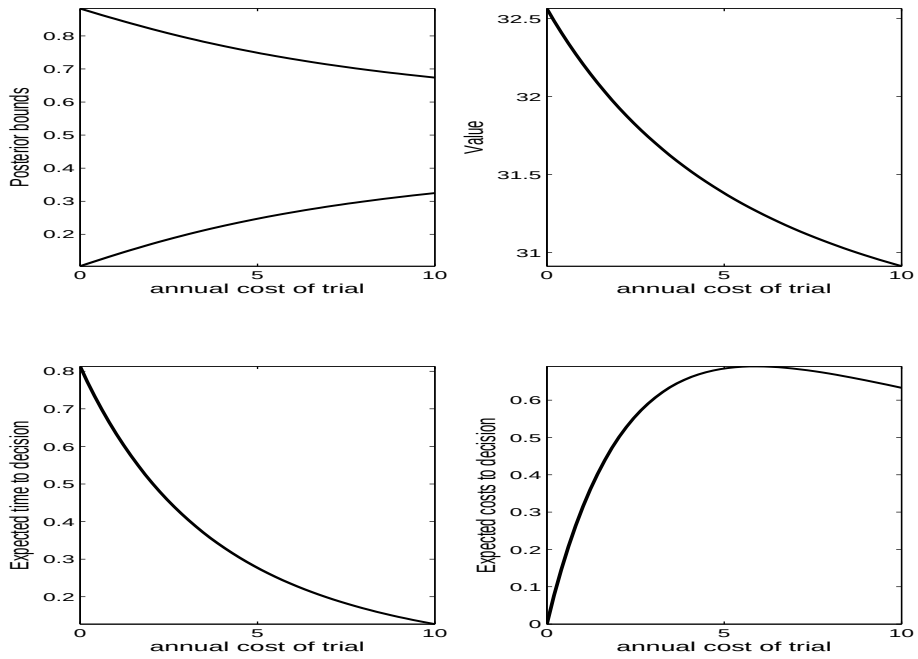

Figure 4: Comparative statics for c.

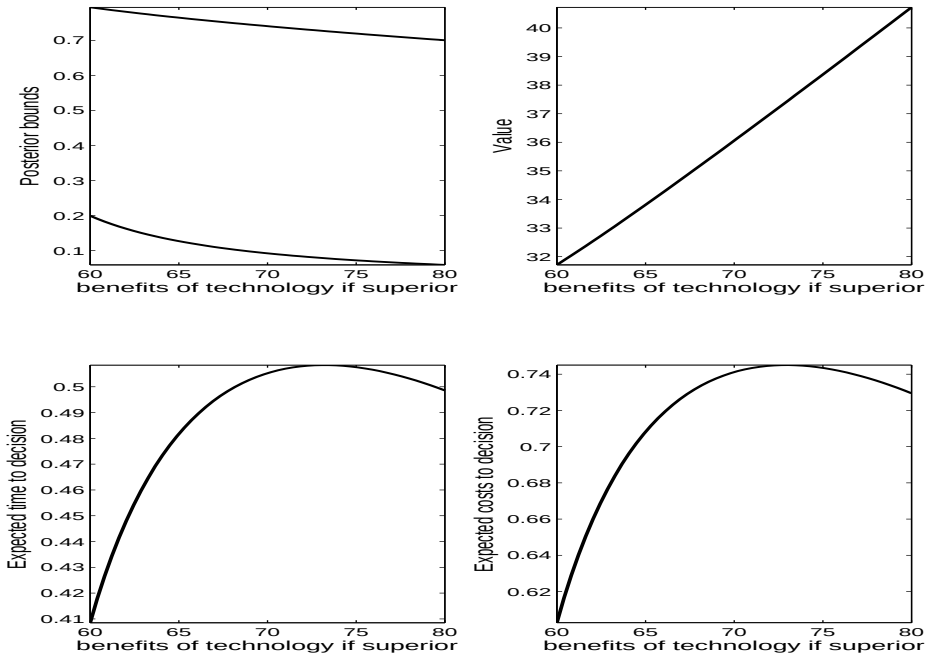

Figure 5: Comparative statics for B_1.

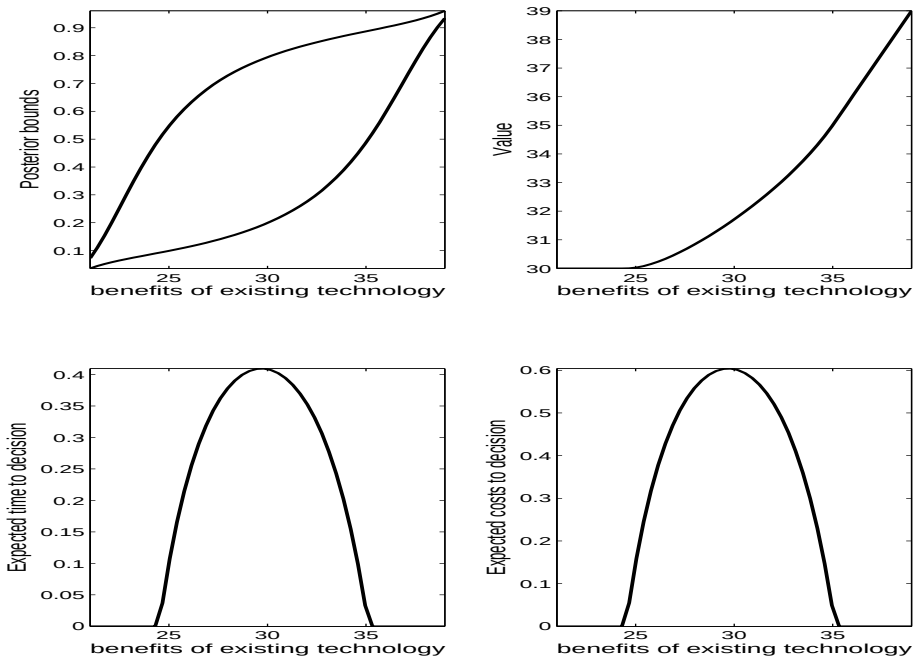

Figure 6: Comparative statics for B.

the posterior process this means that we first wait longer and then start taking decisions sooner.

Finally, in Figure 7 we depict the influence of the prior, p, on the expected time to decision and the expected trial costs. Note that the prior has no influence on the investment/abandonment thresholds, so that we can refer to Figure 1b for comparative statics of the *ex ante* investment/abandonment probabilities. Figure 7 illustrates a simple intuition: the vaguer the *a priori* information about the efficacy of the technology, the longer it will take, on average, before a decision is taken and the more, on average, it will cost to obtain enough information to make that decision. Obviously, the prior is vaguest at $p = 1/2$.

5 Delays in Measurement

An important issue in measuring health benefits as part of a clinical trial is that these are typically not observed instantaneously but, rather, with a delay. This issue has recently been investigated by Chick et al. (2015) who use a Bayesian sequential estimation method. Unlike their approach, in our sequential testing framework, incorporating delays in observations can be achieved as a straightforward extension of the basic model by following the approach introduced in Øksendal (2004).

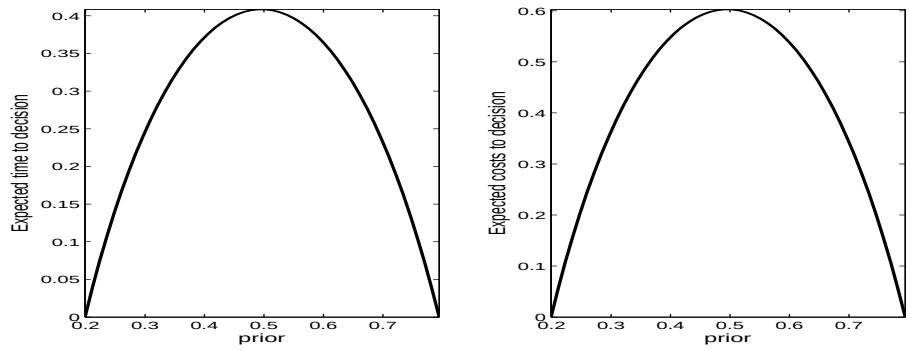

Figure 7: Comparative statics for p.

Suppose that the delay in information becoming available is denoted by $\delta > 0$. Thus, at time t the decision maker observes the information of the patient who entered the trial at time $t - \delta$. If we insist on decision rules not being anticipating, the decision maker has to restrict herself to the set $\mathcal{T}_\delta$ of *delayed stopping times*, i.e. functions $\tau_\delta : \Omega \to [0, \infty]$ such that $\{\, \omega \in \Omega \mid \tau_\delta(\omega) \leq t \,\} \in \mathcal{F}_{t-\delta}$, for all $t \geq \delta$. That is, if at time t the DM stops the trial, the last patient's treatment starts at time t. When the results of that patient come in, at time $t + \delta$, the DM decides to invest or abandon.

The value of the project now equals

$$F^*_\delta(\pi) := \sup_{\tau_\delta \in \mathcal{T}_\delta} \mathsf{E}_\pi \left[- \int_0^{\tau_\delta} e^{-rt}(c + P\pi_t)dt + e^{-r\tau_\delta} \max\left\{ F_A(\pi_{\tau_\delta}), F_I(\pi_{\tau_\delta}) \right\} \right]. \tag{11}$$

Note that $F^*_0(\pi) = F^*(\pi)$, cf. (7). We can prove the following result (for the proof see Appendix C).

Proposition 3. *If the optimal stopping problem* (11) *has a solution, then the continuation region has the form* $(\pi^\delta_A, \pi^\delta_I)$*, for some constants* $0 < \pi^\delta_A < \pi^\delta_I < 1$*. The optimal stopping time is* $\tau^*_\delta = \inf\{t \geq 0 | \pi_t \notin (\pi^\delta_A, \pi^\delta_I)\} + \delta$*. Moreover, for every* $\delta > 0$ *and* $\pi \in (\pi^\delta_A, \pi^\delta_I)$ *the value of the project is given by*

$$F^*_\delta(\pi) = -\left(\frac{c}{r} + \pi\frac{P}{r}\right) + \hat{v}_{\pi^\delta_A, \pi^\delta_I}(\pi) F_\delta(\pi^\delta_I) + \check{v}_{\pi^\delta_A, \pi^\delta_I}(\pi) F_\delta(\pi^\delta_A),$$

where

$$\begin{aligned} F_\delta(\pi) \equiv e^{-r\delta}\Big\{ & \frac{c + P\pi}{r} + (B_1 - B_0)\pi\Phi(b_+) - L\pi\Phi(d_-) \\ & + (B_0 - I)[\pi\Phi(b_+) + (1-\pi)\Phi(b_-)] + B[\pi\Phi(d_-) + (1-\pi)\Phi(d_+)]\Big\}, \\ b_\pm = \frac{\sigma}{\mu\sqrt{\delta}}\left[\log\left(\frac{(1-\bar{\pi})\pi}{\bar{\pi}(1-\pi)}\right) \pm \frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right], & \quad d_\pm = \frac{\sigma}{\mu\sqrt{\delta}}\left[\log\left(\frac{(1-\pi)\bar{\pi}}{\pi(1-\bar{\pi})}\right) \pm \frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right], \end{aligned}$$

Φ *is the cdf of the standard normal distribution, and* $\hat{v}$ *and* $\check{v}$ *are the expected discount factors defined in* (8) *and* (9)*, respectively. The triggers* π^δ_A *and* π^δ_I *are such that* F^*_δ *is* C^1 *on* $(0, 1)$*.*

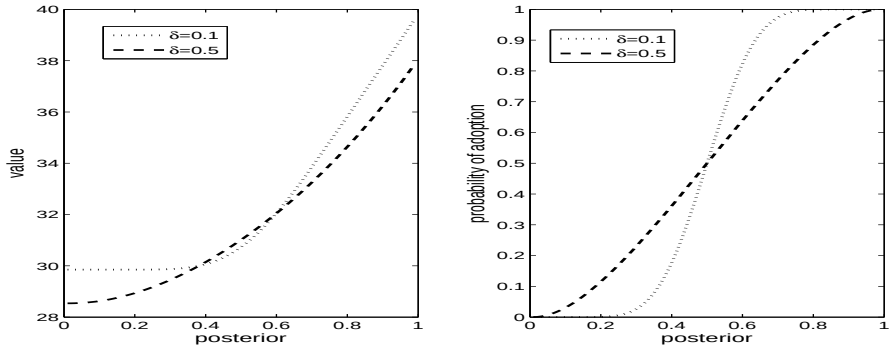

Figure 8: Net present values of an immediate decision (left-panel) and probabilities of adoption upon stopping (right-panel) for different values of δ. The parameter values are taken as in Table 2.

The function F_δ is, in a sense a probability-weighted average of F_I and F_A, delayed by δ, but computed at the time the trial is stopped. The DM knows that, when she stops the trial, she will make a decision in δ units of time. She will then adopt (abandon) if, at that time, the posterior exceeds (is below) $\bar{\pi}$. At the time she decides to stop admitting patients to the trial, she can compute the probabilities of these events, as well as the expectations of the posterior, conditional on these events. See Appendix C for details.

Figure 8 shows the net present value of an immediate adoption/abandonment decision for different values of the delay parameter δ. Note that the NPV is not unambiguously point-wise increasing or decreasing in δ, but that for small and large values of the posterior, the effect of delay is unambiguous: the longer the delay in observations, the lower the NPV. The right-panel of Figure 8 shows the probability of adoption after δ periods as a function of the posterior belief in a superior technology at the time of stopping the trial. Note that for $\delta = 0$ this probability is degenerate: 1 if $\pi > \bar{\pi}$ and 0 if $\pi < \bar{\pi}$. When $\delta > 0$ the probability of adoption is increasing in the posterior at the time the trial is stopped. For values $\pi > \bar{\pi}$, the probability of adoption is higher for smaller values of δ, as is intuitively clear. For values $\pi < \bar{\pi}$, the probability of adoption is *higher* for *longer* delays. This happens because the variance of the distribution of π_δ is larger for larger values of δ. This gives a higher probability of bigger increases over the time interval $[0, \delta]$. Since at the moment the trial is stopped the DM expects to abandon (because $\pi < \bar{\pi}$), a longer time delay essentially means taking a gamble with a larger upside.

Figure 9 shows the value function for the base-case in Table 2 with no delay (left-panel) and a delay of 0.1 years (right-panel). The value with delay is obviously lower than with delay, for every level of posterior belief in a superior technology (at the moment of stopping the trial). It

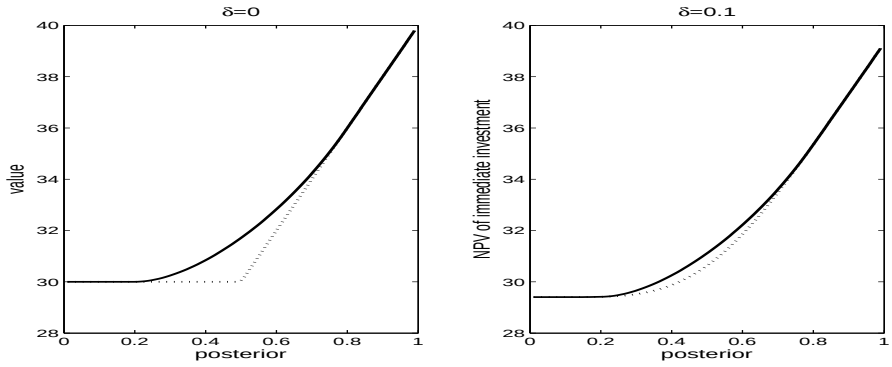

Figure 9: Project value with no delay (left-panel) and positive delay (right-panel). The parameter values are taken as in Table 2.

also appears that the continuation region $(\pi_A^\delta, \pi_A^\delta)$ is narrower with delay. This is confirmed by a numerical calculation of the bounds, which gives a continuation region of (.2063,.7860). The intuition for this result is clear: upon stopping there is another period of length δ over which new observations become available and during which costs of sampling are incurred. As a result the DM can stop the trial earlier and take a gamble on what the observations over the period δ are going to tell her. Obviously this brings extra risk and, thus, a lower value of the project.

The narrowing of the continuation region as a function of the delay δ is also confirmed by a comparative statics analysis, as shown in Figure 10. We can also see that the value of the project is monotonically decreasing in δ. The non-monotonicity of the expected time to stopping the trial and the expected costs of the trial follows from the fact that on the one hand there is a reduction in the expected time of the trial, because a decision is taken sooner, while, on the other hand, the trial duration is extended by δ. Since the effect on the expected time until a decision is taken is non-linear, the combined effect is non-monotonic.

Remark 6. *The existence of a solution to* (11) *depends, firstly, on the existence of a solution to the equations*

$$\hat{v}'_{\pi_A^\delta,\pi_I^\delta}(\pi_A^\delta)F_\delta(\pi_I^\delta) + \check{v}'_{\pi_A^\delta,\pi_I^\delta}(\pi_A^\delta)F_\delta(\pi_A^\delta) = F'_\delta(\pi_A^\delta) \tag{12}$$

and

$$\hat{v}'_{\pi_A^\delta,\pi_I^\delta}(\pi_I^\delta)F_\delta(\pi_I^\delta) + \check{v}'_{\pi_A^\delta,\pi_I^\delta}(\pi_I^\delta)F_\delta(\pi_A^\delta) = F'_\delta(\pi_I^\delta), \tag{13}$$

which ensures that the value function F_δ^* is C^1 on $(0,1)$. Unlike the case with $\delta = 0$ it is difficult to obtain analytical results on the existence and uniqueness of such triggers, because of the highly non-linear nature of F_δ. Secondly, while in the case with $\delta = 0$ we can prove sufficiency of the smooth-pasting

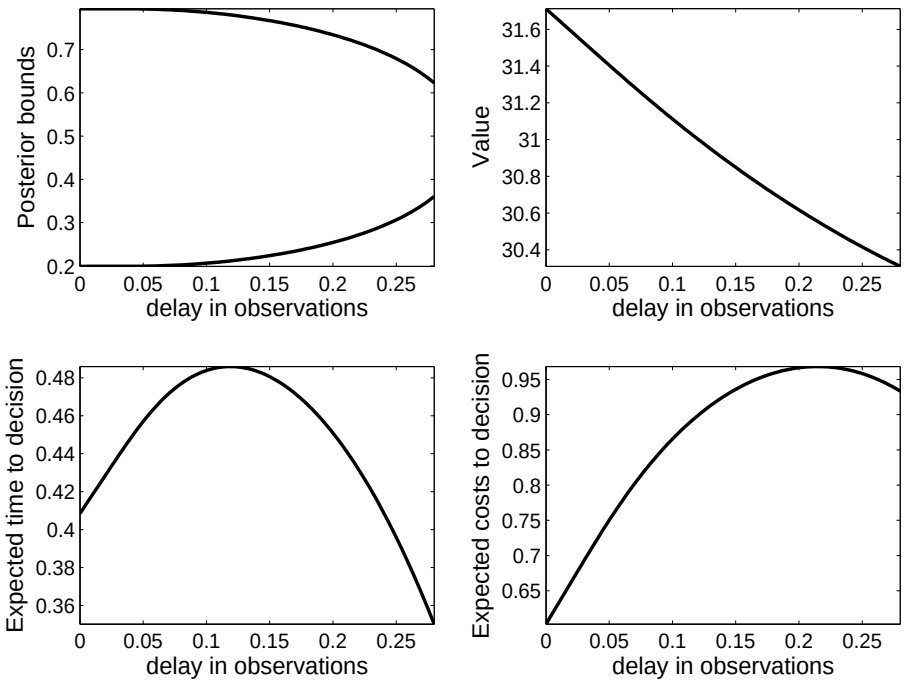

Figure 10: Comparative statics for δ.

conditions, we could not obtain analytical results for $\delta > 0$ (again because of the non-linear nature of F_δ).

Another issue may arise in the case of delay. Consider the benchmark case from Table 2 with $\delta = .3$. Here we can find thresholds that make F^*_δ a C^1 function on $(0,1)$: $\pi^\delta_A = .4316$ and $\pi^\delta_I = .5504$. For this case the smooth-pasting conditions are also sufficient. However, from Figure 11, which depicts F_δ (dashed line) and F^*_δ (solid line) on $\mathcal{C}$, we can see that both coincide, even in the continuation region $\mathcal{C}$. This would suggest that, in this case, $\mathcal{C} = \emptyset$ and, thus, that the optimal stopping time is $\tau^*_\delta = 0$. Therefore, starting the clinical trial is not optimal and one should make an adoption/abandonment decision straight away. The intuition for this result is that, if the delay is long, one incurs sampling costs for a relatively long time before any benefit occurs. Since these benefits are discounted more, the present value of the sampling costs may outweigh the informational benefit of observing trial results. It is then optimal not to incur sampling costs at all. A full investigation of this issue, however, will have to await further research.

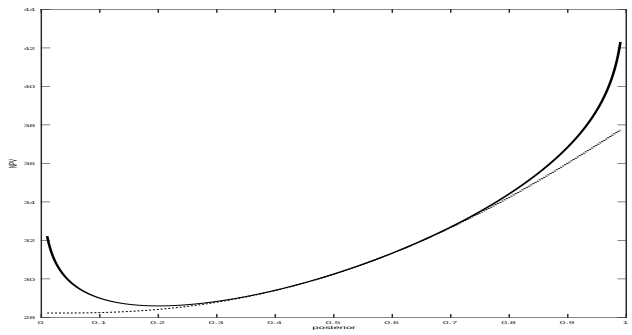

Figure 11: An example where the optimal stopping problem (11) where the continuation region seems to be empty. The parameter values are taken as in Table 2, with $\delta = .3$.

6 An Illustration: Standard vs Robot-Assisted Laparoscopic Prostatectomy

In this section we apply our model to the choice between robot-assisted and standard laparoscopic prostatectomy from the perspective of the UK national health service using data from a published study by Close et al. (2013).[11] Standard laparoscopic prostatectomy and robot-assisted laparoscopic prostatectomy are favoured over the open technique as these cause less bleeding and allow for a quicker return to activities. Robot-assisted laparoscopic prostatectomy is increasingly used compared to the standard technique. However, the high costs has led authorities to question the value of robotic-assisted procedures to patients and the health care system.

The relevant data for an optimal trial/adoption/abandonment decision are given in Table 3. These are used to compute the parameter values as given in Table 4. Here it has been assumed that:

1. the number of patients treated every year will remain constant and the technology will be used forever (to allow for straightforward computation of the discounted streams of benefits for the different scenarios);

2. the QALY level of robot-assisted laparoscopic prostatectomy is chosen such that this technique is just cost effective;

3. the total of 10,000 patients involved in the simulated trial over 10 years were spread uniformly over time in pairs;

[11] Note that Close et al. (2013) do not present the results of a clinical trial, but, rather, a simulation-based analysis. We use the study mainly to provide a realistic background to illustrate our model.

Description	Parameter	Source	Value
Benefit of standard prostatectomy (in QALY)	$QALY_0$	Close et al.	6.44
Value per QALY	$vQALY$	Claxton et al.	£13,000
Cost of standard prostatectomy	c_0	Close et al.	£7,628
Cost of robot-assisted prostatectomy	c_1	Close et al.	£9,040
Benefit of robot-assisted prostatectomy (if cost effective)	$QALY_1$	derived	7.63
Benefit of robot-assisted prostatectomy (if not cost effective)	$QALY_0$	Close et al.	6.44
Patients treated per year	$patients$	Close et al.	5,464
Length of trial		Close et al.	10 years
Patients treated in trial	$trial$	Close et al.	500

Table 3: Relevant trial/decision information for prostatectomy case.

4. the trial costs per patient are assumed to equal the cost difference between the two technologies with an added £500 for other costs;

5. the penalties for not adopting a superior technology (P and L) are set to zero;

6. the value per QALY is taken from Claxton et al. (2015).

For these values, the adoption and abandonment bounds are $\pi_I = .8882$ and $\pi_A = .0064$, respectively. Note that the lower bound is close to zero because the trial is cheap to run relative to the potential benefits to be gained if the technology turns out to be superior. With a prior belief of .5 in robot-assisted laparoscopic prostatectomy being cost effective, this implies that at the start of the trial the probability of adoption (abandonment) is .5598 (.4402). The expected trial costs are £110,600 and the expected duration of the trial is only .23 years. The value of the project at the start of the trial is £15.07bn.

Figure 12 gives a histogram of decision times based on 5,000 simulated sample paths. The right-hand panel gives decision times conditional on robot-assisted laparoscopic prostatectomy being superior. On this event decisions are taken very quickly; in fact, all sample paths lead to a decision within .085 years, i.e., after about a month. This happens because, if robot-assisted laparoscopic prostatectomy is cost effective, then its QALY gain is so much larger than that of the standard procedure that a clinical trial will pick up on this very quickly. In addition, all sample paths lead to adoption of the technique. In other words, along no sample path does a Type I error occur. The left-hand panel gives decision times conditional on robot-assisted laparoscopic prosta-

Parameter	Description	Derivation	Value
B_1	Benefit if superior	$(Qaly_1 \cdot vQaly - c_1) \cdot patients/r$	£16.42bn
B_0	Benefit if not superior	$(Qaly_0 \cdot vQaly - c_1) \cdot patients/r$	£13.60bn
B	Benefit of current treatment	$(Qaly_0 \cdot vQaly - c_0) \cdot patients/r$	£13.86bn
I	Sunk costs of adoption	assumed	£1mln
c	Trial cost flow	$(c_1 - c_0 + 500) \cdot trial$	£.38mln p.a.
r	Discount rate	Close et al.	3.5%
μ	Trend if superior	$(Qaly_1 - Qaly_0) \cdot vQaly \cdot trial$	£3.10mln p.a.
σ	Volatility of observations	assumed	£1.5mln p.a.
P	Penalty of non-adoption during trial	assumed	0
L	Penalty of non-adoption after trial	assumed	0
p	Prior	assumed	.5

Table 4: Parameter values for laparoscopic prostatectomy case.

tectomy not leading to any superior health benefits over the standard procedure. On this event decisions are taken very late: clinical trials can (theoretically) last as long as 1,000 years. Unrealistic as this may seem, it is not unexpected. Due to the low costs of conducting the trial, the fairly low number of patients involved per annum, and the high benefits if the robot-assisted technique turns out to be superior after all, it is optimal to keep the trial going. Finally, in the simulation, 87% of sample paths (conditional on $\theta = 0$) end in an abandonment decision. This implies that the (simulated) probability of a Type II error is 13%. Note that this performance is actually *better* than the usually imposed error probabilities in frequentist studies. In fact, these error probabilities can be computed exactly[12] as $\alpha = .0056$ and $\beta = .1252$, showing that our simulations are fairly accurate.

By choosing the QALY value of robot-assisted laparoscopy such that the technology is just cost effective we have biased the decision against adoption. Figure 13 shows a comparative statics exercise for this QALY level. As can be seen the value of the project increases sharply in the QLAY

[12]This uses the fact that the bounds π_I and π_A can be written as likelihood ratio bounds Λ_I and Λ_A, respectively. Since $(\Lambda_t)_{t\geq 0}$ follows a GBM conditional on θ, with $\Lambda_0 = 1$, $\mathsf{P}_\theta - a.s.$, it can easily be derived that (see, for example, Stokey, 2009):

$$\mathsf{P}_\theta(\Lambda_{\tau^*} = \Lambda_I) = \frac{1 - \Lambda_A^{1-2\theta}}{\Lambda_I^{1-2\theta} - \Lambda_A^{1-2\theta}}, \quad \text{and} \quad \mathsf{P}_\theta(\Lambda_{\tau^*} = \Lambda_A) = \frac{\Lambda_I^{1-2\theta} - 1}{\Lambda_I^{1-2\theta} - \Lambda_A^{1-2\theta}}.$$

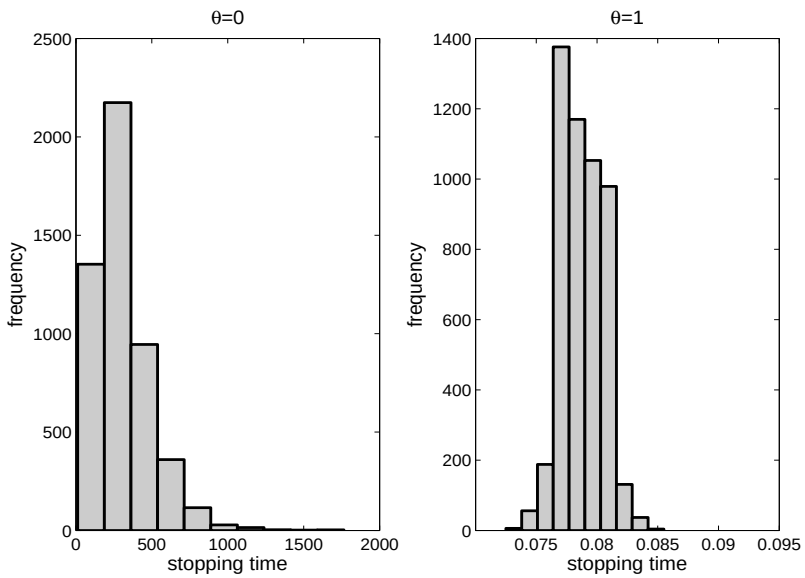

Figure 12: Stopping time distribution.

level of robot-assisted laparoscopic prostatectomy (conditional on it being superior). Since this QALY level only influences the parameter B it is no surprise to see similar comparative statics as those in Section 4, Figure 6.

7 Concluding Remarks

The Bayesian model presented in this paper was motivated by a desire to bring together sequential hypothesis testing and real options analysis to inform decision making in health technology assessment under irreversible costs. In a sequential setting, the decision maker is required to simultaneously assess (i) whether enough statistical evidence has been gathered, and (ii) if the newly proposed technology is deemed superior to the existing one.

The proposed framework deals with such requirements by specifying a decision rule that, being a function of the uncertain outcomes of a clinical trial as well as payoffs, produces inferential bounds for abandonment or adoption of the new technology. At each point during the clinical trial the Bayesian posterior process is assessed against an investment and an abandonment decision bound, indicating whether the trial should continue or not. When the trial stops, a decision is taken, either for investment or abandonment, and it is ensured that health benefits to the population are maximized.

The framework considers and combines dimensions that in present approaches are typically studied in isolation. By putting health benefits at the centre of the decision making process, and

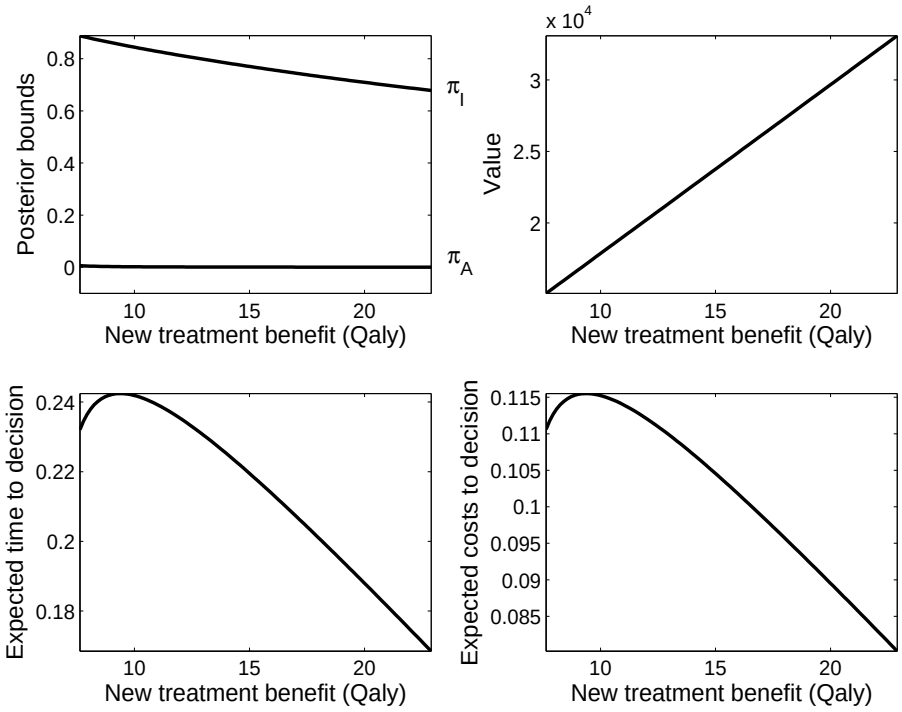

Figure 13: Comparative statics for the QALY level of robot-assisted laparoscopy.

by considering the value of gathering more evidence, the model explicitly incorporates: (i) the cost and benefits of delaying adoption/abandonment of a newly developed health-care technology, (ii) discounted population payoffs, (iii) the costs of conducting further research, and (iv) the benefits of gathering further evidence at each point during the trial. As a result, the decision rule derived in this paper is optimal (i) *ex ante*, (ii) at any time during the clinical trial, and (iii) at the time an adoption/abandonment decision is made. Delays in observing health outcomes after a trial can be incorporated and are shown to lower the value of the new technology, and to shrink the inaction region. For longer delays, it is not entirely clear what the optimal decision is, although it appears that an immediate decision might be required. A full investigation of this issue presents an avenue for future research.

Acknowledgments

We gratefully acknowledge helpful comments from Karl Claxton, Brendan McCabe, Stefan Scholz, Peter N. Smith, and seminar participants at the 2014 annual conference of the Royal Statistical Society in Sheffield, UK, and the Economics seminar at Bielefeld University, Germany. Martin Forster, in particular, is thanked for his encouragement and many helpful comments and sug-

gestions. Tomasz Zastawniak's assistance in proving Lemma 1 is greatly appreciated. Detailed, and very helpful, comments were received from two anonymous referees and the editor, Herbert Dawid. Much of the research was done when both authors were at the Department of Economics & Related Studies at the University of York, UK, which we thank for its support. The second author (Bregantini) gratefully acknowledges financial support by the ESRC through a PhD fellowship.

Appendix

A Proof of Proposition 1

1. Since $(\pi_t)_{t\geq 0}$ is a martingale, the optimal stopping problem (4) can be written as

$$\begin{aligned} F^*(\pi) &= -\left(\frac{c}{r}+\pi\frac{P}{r}\right)+\sup_\tau \mathsf{E}_\pi\left[e^{-r\tau}\max\left\{F_I(\pi_\tau)+\frac{c}{r}+\pi_\tau\frac{P}{r}, F_A(\pi_\tau)+\frac{c}{r}+\pi_\tau\frac{P}{r}\right\}\right] \\ &= -\left(\frac{c}{r}+\pi\frac{P}{r}\right)+\sup_\tau \mathsf{E}_\pi\left[e^{-r\tau}\max\left\{\pi_\tau(B_1-B_0+P/r)+B_0+\frac{c}{r}-I, B+\frac{c}{r}-\pi_\tau(L-P/r)\right\}\right] \\ &= -\left(\frac{c}{r}+\pi\frac{P}{r}\right)+\sup_\tau \mathsf{E}_\pi\left[e^{-r\tau}\max\left\{G_I(\pi_\tau), G_A(\pi_\tau)\right\}\right], \end{aligned} \tag{A.1}$$

where

$$G_I(\pi)=\pi(B_1-B_0+P/r)+B_0+\frac{c}{r}-I, \quad \text{and} \quad G_A(\pi)=B+\frac{c}{r}-\pi_\tau(L-P/r).$$

Assumptions 1 and 2 ensure that G_I and G_A are increasing and non-increasing, respectively, whereas Assumptions 3 and 4 ensure that $G_A(0)>G_I(0)$ and $G_I(1)>G_A(1)$. All these are natural assumptions. They imply that there is a unique point $\bar{x}\in(0,1)$ such that $G_I(\bar{x})=G_A(\bar{x})$. Note that $\bar{x}=\bar{\pi}$, cf. (3).

Define the function $G:(0,1)\to\mathbb{R}$ by

$$G(\pi)=1_{\pi\leq\bar{\pi}}G_A(\pi)+1_{\pi>\bar{\pi}}G_I(\pi).$$

Note that $G=G_A\vee G_I$ and that G is C^2 on $(0,1)\setminus\{\bar{\pi}\}$.

2. From Peskir and Shiryaev (2006) it follows that we need to find a function $F^*\in C^1$, with second derivatives locally bounded, which dominates G on $(0,1)$, and a set $\mathcal{C}\subset(0,1)$, that solve the free boundary problem

$$\begin{cases} \mathscr{L}F^*-rF^*=0 \quad \text{on } \mathcal{C}, \text{ and} \quad \mathscr{L}F^*-rF^*<0 \quad \text{on } (0,1)\setminus\mathcal{C} \\ F^*>G \quad \text{on } \mathcal{C}, \text{ and} \quad F^*=G \quad \text{on } (0,1)\setminus\mathcal{C} \\ \frac{\partial F^*}{\partial\pi}|_{\partial\mathcal{C}}=\frac{\partial G}{\partial\pi}|_{\partial\mathcal{C}}. \end{cases} \tag{A.2}$$

Here $\mathscr{L}$ denotes the characteristic operator of $(\pi_t)_{t\geq 0}$, i.e., for any $\varphi \in C^2$,

$$\mathscr{L}\varphi(\pi) = \frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2 \pi^2(1-\pi)^2\varphi''(\pi).$$

Note that the condition

$$\mathscr{L}F^* - rF^* < 0 \quad \text{on } (0,1)\setminus\mathcal{C},$$

is always satisfied since $G'' = 0$ on $(0,1)\setminus\{\bar{\pi}\}$. For (A.2) to make sense, we also need that $G > 0$. A necessary and sufficient condition for this to hold is that $G(\bar{\pi}) > 0$. The property follows from

$$\begin{aligned} G_A(\bar{\pi}) &= B + \frac{c}{r} - \bar{\pi}(L - P/r) \\ &> \bar{\pi}(L-P/r)\left[\frac{\gamma+1/2}{\bar{\pi}+\gamma-1/2} - 1\right] \\ &= \bar{\pi}(L-P/r)\frac{1-\bar{\pi}}{\bar{\pi}+\gamma-1/2} \geq 0, \end{aligned}$$

where the strict inequality in the second line follows from Assumption 5 and the weak inequality in the last line follows from Assumption 2 and the fact that $\gamma > 1/2$.

3. On $(0,1)$, define the functions $\hat{\varphi} : (0,1) \to \mathbb{R}_+$ and $\check{\varphi} : (0,1) \to \mathbb{R}_+$, by[13]

$$\hat{\varphi}(\pi) = \sqrt{\pi(1-\pi)}\left(\frac{\pi}{1-\pi}\right)^{\gamma}, \quad \text{and} \quad \check{\varphi}(\pi) = \sqrt{\pi(1-\pi)}\left(\frac{1-\pi}{\pi}\right)^{\gamma}. \tag{A.3}$$

Note that $\hat{\varphi}$ and $\check{\varphi}$ are the increasing and decreasing solutions to the differential equation $\mathscr{L}\varphi - r\varphi = 0$, respectively. So, any solution to $\mathscr{L}\varphi - r\varphi = 0$ is of the form

$$\varphi(\pi) = \hat{A}\hat{\varphi}(\pi) + \check{A}\check{\varphi}(\pi),$$

where $\hat{A}$ and $\check{A}$ are arbitrary constants. Furthermore, it is easily obtained that

$$\begin{aligned} \hat{\varphi}'(\pi) &= \hat{\varphi}(\pi)\frac{1/2+\gamma-\pi}{\pi(1-\pi)} > 0, \quad \check{\varphi}'(\pi) = \check{\varphi}(\pi)\frac{1/2-\gamma-\pi}{\pi(1-\pi)} < 0, \\ \hat{\varphi}''(\pi) &= \hat{\varphi}(\pi)\frac{(1/2+\gamma-\pi)(\gamma-1/2+\pi)-\pi(1-\pi)}{\pi^2(1-\pi)^2} > 0, \end{aligned}$$

and

$$\check{\varphi}''(\pi) = \check{\varphi}(\pi)\frac{(1/2-\gamma-\pi)(\pi-1/2-\gamma)-\pi(1-\pi)}{\pi^2(1-\pi)^2} > 0.$$

From this it is obtained that

$$\frac{\check{\varphi}''(\pi)}{\hat{\varphi}''(\pi)} = \frac{\check{\varphi}(\pi)}{\hat{\varphi}(\pi)} = \left(\frac{1-\pi}{\pi}\right)^{2\gamma}.$$

[13] The results in this part of the proof are standard and can be found in the literature, such as, for example, Borodin and Salminen (2002). They are collected here for ease of reference.

4. Fix $\pi_L \leq \bar{\pi}$ and define the mapping $\pi \mapsto \check{V}(\pi;\pi_L)$, by

$$\check{V}(\pi;\pi_L) = \hat{A}(\pi_L)\hat{\varphi}(\pi) + \check{A}(\pi_L)\check{\varphi}(\pi), \tag{A.4}$$

where the constants $\hat{A}(\pi_L)$ and $\check{A}(\pi_L)$ are given by

$$\hat{A}(\pi_L) = \frac{\check{\varphi}(\pi_L)G_A'(\pi_L) - \check{\varphi}'(\pi_L)G_A(\pi_L)}{\check{\varphi}(\pi_L)\hat{\varphi}'(\pi_L) - \check{\varphi}'(\pi_L)\hat{\varphi}(\pi_L)}, \quad \text{and} \quad \check{A}(\pi_L) = \frac{\hat{\varphi}'(\pi_L)G_A(\pi_L) - \hat{\varphi}(\pi_L)G_A'(\pi_L)}{\check{\varphi}(\pi_L)\hat{\varphi}'(\pi_L) - \check{\varphi}'(\pi_L)\hat{\varphi}(\pi_L)}. \tag{A.5}$$

Note that $\mathscr{L}\check{V}(\pi;\pi_L) - r\check{V}(\pi;\pi_L) = 0$ for all $\pi \in (0,1)$. In addition, the function $\check{V}$ satisfies $\check{V}(\pi_L;\pi_L) = G_A(\pi_L)$ and $\check{V}'(\pi_L;\pi_L) = G_A'(\pi_L)$.

Condition 5 implies that

$$G_A(\bar{\pi}) = B + \frac{c}{r} - \bar{\pi}(L - P/r) > \bar{\pi}\left(\frac{\gamma + \frac{1}{2}}{\bar{\pi} + \gamma - \frac{1}{2}} - 1\right)(L - P/r) = \frac{1-\bar{\pi}}{\bar{\pi} + \gamma - \frac{1}{2}}\bar{\pi}(L - P/r) > 0.$$

Since $G_A' \leq 0$, it, therefore, holds that $\check{A}(\pi_L) > 0$ for all $\pi_L \leq \bar{\pi}$.

The denominator of $\hat{A}(\pi_L)$ (and thus $\check{A}(\pi_L)$) does not change in π_L.[14] Since condition 5 ensures that $G_A(\bar{\pi}) > 0$,[15] it follows that,

$$\frac{\partial \hat{A}(\pi_L)}{\partial \pi_L} < 0, \quad \text{and} \quad \frac{\partial \check{A}(\pi_L)}{\partial \pi_L} > 0.$$

5. Condition 5 ensures that $\hat{A}(\bar{\pi}) > 0$. Therefore, $\pi \mapsto \check{V}(\pi;\bar{\pi})$ is a (strictly) convex function, which satisfies $\check{V}(\cdot;\bar{\pi}) \to \infty$ as $\pi \uparrow 1$ or $\pi \downarrow 0$. So, there is a unique point $\pi_H \in (\bar{\pi}, 1)$ such that $\check{V}'(\pi_H;\bar{\pi}) = G_I'(\pi_H)$. Since $\check{V}'(\bar{\pi};\bar{\pi}) = G_A'(\bar{\pi}) < 0$, at π_H it holds that $\check{V}(\pi_H;\bar{\pi}) < G_I(\pi_H)$. Also, for π large enough, it holds that $\check{V}(\pi;\bar{\pi}) > G_I(\pi)$.

6. Since $\hat{A}(\pi_L)$ decreases and $\check{A}(\pi_L)$ increases in π_L, the mapping $\pi \mapsto \check{V}(\pi;\pi_L)$ has the property that for every $\pi > \pi_L$ it holds that $\partial\check{V}(\pi;\pi_L)/\partial\pi_L < 0$. So, the point $\pi_H \in (\bar{\pi}, 1)$ where $\check{V}'(\pi_H;\pi_L) = G_I'(\pi_H)$ is decreasing in π_L, as is the value $\check{V}(\pi_H;\pi_L)$. Now decrease π_L from $\bar{\pi}$ to 0. There will be a unique π_A, with corresponding π_I at which $\check{V}(\pi_I;\pi_A) = G_I(\pi_I)$ and $\check{V}'(\pi_I;\pi_A) = G_I'(\pi_I)$.

7. The interval $\mathcal{C} = (\pi_A, \pi_I)$ and the proposed function F^* together solve the free-boundary problem (A.2). The fact that π_A and π_I are the unique triggers that make F^* a C^1 function on (π_A, π_I) follows by construction.

[14] It holds that $\check{\varphi}(\pi_L)\hat{\varphi}'(\pi_L) - \check{\varphi}'(\pi_L)\hat{\varphi}(\pi_L) = 2\gamma$.

[15] Note that $G_A(\pi) > 0 \iff \frac{B+c/r}{L-P/r} > \pi$. Since $\frac{\gamma+1/2}{\pi+\gamma-1/2} > 1$, it follows that $\hat{A}(\pi_L) > 0$ implies $G_A(\pi_L) > 0$.

8. The proof that a solution to (A.2) is also a solution to the optimal stopping problem, i.e., that

$$G^*(\pi) := F^*(\pi) + \left(\frac{c}{r} + \pi\frac{P}{r}\right) = \sup_\tau \mathsf{E}_\pi\left[e^{-r\tau}G(\pi_\tau)\right] =: \sup_\tau J^\tau(\pi),$$

is standard. Here we will sketch the main argument, for technical details see, e.g., Øksendal (2010). Obviously, it holds that $G^*(\pi) \leq \sup_\tau J^\tau(\pi)$. To prove the reverse inequality, take any stopping time τ. It now holds that

$$\begin{aligned} G^*(\pi) &= \mathsf{E}_\pi\left[e^{-r\tau}G^*(\pi_\tau)\right] - \mathsf{E}_\pi\left[\int_0^\tau e^{-rt}(\mathscr{L}-r)G^*(\pi_t)dt\right] \\ &\geq \mathsf{E}_\pi\left[e^{-r\tau}G^*(\pi_\tau)\right] \geq \mathsf{E}_\pi\left[e^{-r\tau}G(\pi_\tau)\right] = J^\tau(\pi), \end{aligned}$$

where the first equality follows from Dynkin's formula (Øksendal, 2010), the first inequality follows from $\mathscr{L}G^* - rG^* \leq 0$, and the second inequality follows from $G^* \geq G$. Hence, $G^*(\pi) \geq \sup_\tau J^\tau(\pi)$. ■

9. Suppose that condition 5 is not satisfied. Then there exists a unique $\pi_* \in (0, \bar{\pi})$ for which

$$\frac{B + c/r}{L - P/r} = \pi_*\frac{\gamma + 1/2}{\pi_* + \gamma - 1/2}.$$

Following previous arguments there also exists a unique π^* such that

$$\check{V}'(\pi^*; \pi_*) = G_I'(\pi^*).$$

Figure 14 illustrates the idea of the proof and from this figure it is clear that the condition needed is that

$$\check{V}(\pi^*; \pi_*) \leq G_I(\pi^*).$$

If this condition is satisfied the arguments above go through unchanged.

B Proof of Proposition 2

Let $\tau^* = \inf\{t \geq |\pi_t \notin (\pi_A, \pi_I)\}$ be the unique solution to (A.1). The constants $\hat{A}(\pi_A)$ and $\check{A}(\pi_A)$ in (A.5) can be rewritten as

$$\begin{aligned} \check{A}(\pi_A) &= \frac{\pi_A(1-\pi_A)}{2\gamma\check{\varphi}(\pi_A)}\left(\frac{\frac{1}{2}+\gamma-\pi_A}{\pi_A(1-\pi_A)}G_A(\pi_A) - G_A'(\pi_A)\right), \quad \text{and} \\ \hat{A}(\pi_A) &= \frac{\pi_A(1-\pi_A)}{2\gamma\hat{\varphi}(\pi_A)}\left(G_A'(\pi_A) - \frac{\frac{1}{2}-\gamma-\pi_A}{\pi_A(1-\pi_A)}G_A(\pi_A)\right), \end{aligned}$$

where $\hat{\varphi}$ and $\check{\varphi}$ are as in (A.3), and G_A is defined in (A.1).

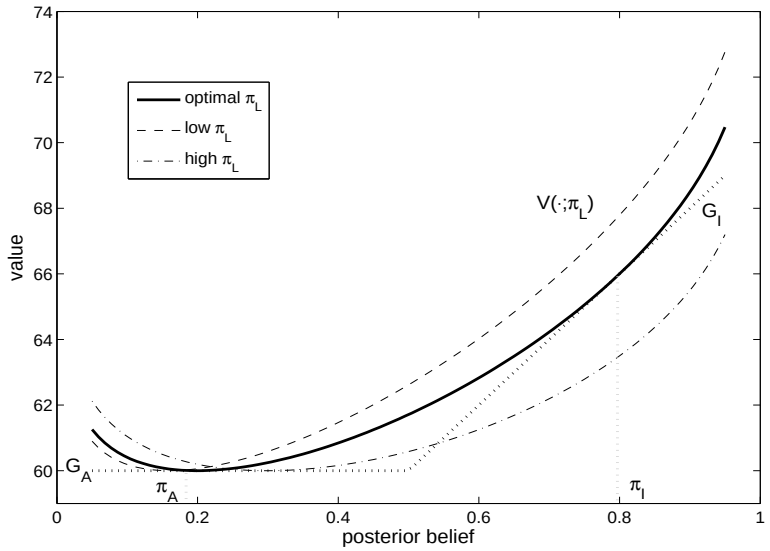

Figure 14: The function $\pi \mapsto \check{V}(\pi;\pi_L)$ for different values of π_L.

The proof of Proposition 1 was based on the fact that the mapping $\pi \mapsto \check{V}(\pi;\pi_L)$ is such that $\check{V}(\pi_L;\pi_L) = G_A(\pi_L)$ and $\check{V}'(\pi_L;\pi_L) = G_A'(\pi_L)$ for all $\pi_L \in (0,\bar{\pi})$. We could also have started by defining a mapping $\pi \mapsto \hat{V}(\pi;\pi_H)$ with the property that $\hat{V}(\pi_H;\pi_H) = G_I(\pi_H)$ and $\hat{V}'(\pi_H;\pi_H) = G_I'(\pi_H)$ for all $\pi_H \in (\bar{\pi},1)$. It is easy to see that $\hat{V}$ is given by

$$\hat{V}(\pi;\pi_H) = \hat{B}(\pi_I)\hat{\varphi}(\pi) + \check{B}(\pi_I)\check{\varphi}(\pi),$$

where

$$\check{B}(\pi_I) = \frac{\pi_I(1-\pi_I)}{2\gamma\check{\varphi}(\pi_I)}\left(\frac{\frac{1}{2}+\gamma-\pi_I}{\pi_I(1-\pi_I)}G_I(\pi_I) - G_I'(\pi_I)\right), \quad \text{and}$$
$$\hat{B}(\pi_I) = \frac{\pi_I(1-\pi_I)}{2\gamma\hat{\varphi}(\pi_I)}\left(G_I'(\pi_I) - \frac{\frac{1}{2}-\gamma-\pi_I}{\pi_I(1-\pi_I)}G_I(\pi_I)\right).$$

Using a similar reasoning as in the proof of Proposition 1 it holds that $\hat{B} > 0$, $\check{B} > 0$, and $\hat{V}$ is strictly convex.

Finally, note that at the unique solution (π_A,π_I) to the optimal stopping problem (A.1) it holds that

$$\hat{V}(\pi_A;\pi_I) = G_A(\pi_A), \quad \hat{V}'(\pi_A;\pi_I) = G_A'(\pi_A), \quad \check{V}(\pi_I;\pi_A) = G_I(\pi_I), \quad \text{and} \quad \check{V}'(\pi_I;\pi_A) = G_I'(\pi_I).$$

1. Consider a change in B (which shifts G_A up). It holds that

$$\frac{d\check{V}(\pi_I;\pi_A)}{dB} = -\frac{\frac{1}{2}-\gamma-\pi_A}{2\gamma}\frac{\hat{\varphi}(\pi_I)}{\hat{\varphi}(\pi_A)} + \frac{\frac{1}{2}+\gamma-\pi_A}{2\gamma}\frac{\check{\varphi}(\pi_I)}{\check{\varphi}(\pi_A)} > 0.$$

Since $dG_I/dB = 0$ this implies that $\check{V}(\pi_I;\pi_A) > G_I(\pi_I)$ after an increase in B. This implies that π_A must be increased to re-establish optimality.

Since $d\hat{V}(\pi_A;\pi_I)/dB = 0$ and $dG_A/dB > 0$, it holds that $\hat{V}(\pi_A;\pi_I) < G_A(\pi_A)$ after an increase in B. This implies that π_I must be increased to re-establish optimality.

2. Consider a change in B_1 (which shifts G_I up and increases its slope). It holds that

$$\frac{d\hat{V}(\pi_A;\pi_I)}{dB_1} = \frac{\pi_I}{2\gamma}\left[(\tfrac{1}{2}+\gamma)\frac{\hat{\varphi}(\pi_A)}{\hat{\varphi}(\pi_I)} + (\gamma-\tfrac{1}{2})\frac{\check{\varphi}(\pi_A)}{\check{\varphi}(\pi_I)}\right] > 0.$$

Since $dG_A/dB_1 = 0$ this implies that $\hat{V}(\pi_A;\pi_I) > G_A(\pi_A)$ after an increase in B_1. This implies that π_I must be decreased to re-establish optimality.

Since $d\check{V}(\pi_I;\pi_A)/dB_1 = 0$ and $dG_I/dB_1 > 0$, it holds that $\check{V}(\pi_I;\pi_A) < G_I(\pi_I)$ after an increase in B_1. This implies that π_A must be decreased to re-establish optimality.

3. Consider a change in c (which shifts both G_A and G_I up by the same amount). Since an increase in c increases both $\hat{A}$ and $\check{A}$ it strictly increases $\check{V}''(\pi;\pi_A)$, for all π. Since $\check{V}$ is strictly convex and G_I is linear this implies that $\check{V}(\pi_I;\pi_A) > G_I(\pi_I)$ after the change in c. To restore optimality π_A has to be increased.

Similarly, an increase in c increases both $\hat{B}$ and $\check{B}$, which strictly increases $\hat{V}''(\pi;\pi_I)$, for all π. Since $\hat{V}$ is strictly convex and G_A is linear this implies that $\hat{V}(\pi_A;\pi_I) > G_A(\pi_A)$ after the change in c. To restore optimality π_I has to be decreased.

4. Consider a change in γ. Such a change affects $\hat{A}$, $\check{A}$, $\hat{B}$, $\check{B}$, $\hat{\varphi}$, and $\check{\varphi}$. Direct computations reveal that

$$\frac{\partial\hat{\varphi}(\pi)}{\partial\gamma} = \hat{\varphi}(\pi)\log\left(\frac{\pi}{1-\pi}\right), \quad \frac{\partial\check{\varphi}(\pi)}{\partial\gamma} = \check{\varphi}(\pi)\log\left(\frac{1-\pi}{\pi}\right),$$
$$\frac{\partial\hat{A}}{\partial\gamma} = \hat{A}\left[\log\left(\frac{1-\pi_A}{\pi_A}\right) - \frac{1}{\gamma}\right] + \frac{\check{\varphi}(\pi_A)}{2\gamma}\frac{G_A(\pi_A)}{\pi_A(1-\pi_A)},$$

and

$$\frac{\partial\check{A}}{\partial\gamma} = \check{A}\left[\log\left(\frac{\pi_A}{1-\pi_A}\right) - \frac{1}{\gamma}\right] + \frac{\hat{\varphi}(\pi_A)}{2\gamma}\frac{G_A(\pi_A)}{\pi_A(1-\pi_A)}.$$

It is then easily obtained that

$$\begin{aligned}\frac{\partial\check{V}(\pi_I;\pi_A)}{\partial\gamma} =& \left(\hat{A}\hat{\varphi}(\pi_I) - \check{A}\check{\varphi}(\pi_I)\right)\log\left(\frac{1-\pi_A}{\pi_A}\frac{\pi_I}{1-\pi_I}\right) \\ &+ \frac{1}{2\gamma}\left[\frac{\hat{\varphi}(\pi_I)\check{\varphi}(\pi_A) + \check{\varphi}(\pi_I)\hat{\varphi}(\pi_A)}{\pi_A(1-\pi_A)}G_A(\pi_A) - 2G_I(\pi_I)\right].\end{aligned} \tag{B.1}$$

The first term on the right-hand side is the product of two positive terms. [The logarithmic term is positive because $\pi_A < 1/2 < \pi_I$. The term in round brackets is positive because differentiability at π_I gives that

$$\begin{aligned} &\hat{A}\hat{\varphi}'(\pi_I) + \check{A}\check{\varphi}'(\pi_I) = G_I'(\pi_I) > 0 \\ \iff &\frac{\gamma + (\frac{1}{2} - \pi_I)}{\pi_I(1-\pi_I)}\hat{A}\hat{\varphi}(\pi_I) + \frac{-\gamma + (\frac{1}{2} - \pi_I)}{\pi_I(1-\pi_I)}\check{A}\check{\varphi}(\pi_I) > 0 \\ \iff &\hat{A}\hat{\varphi}(\pi_I) > \frac{\gamma - (\frac{1}{2} - \pi_I)}{\gamma + (\frac{1}{2} - \pi_I)}\check{A}\check{\varphi}(\pi_I) > \check{A}\check{\varphi}(\pi_I), \end{aligned} \tag{B.2}$$

where the final inequality follows from $\pi_I > 1/2$.]

The term between square brackets of (B.1) can be expanded as

$$\begin{aligned} \frac{\hat{\varphi}(\pi_I)\check{\varphi}(\pi_A) + \check{\varphi}(\pi_I)\hat{\varphi}(\pi_A)}{\pi_A(1-\pi_A)} G_A(\pi_A) - 2G_I(\pi_I) &= \left(\frac{\hat{\varphi}(\pi_I)}{\hat{\varphi}(\pi_A)} + \frac{\check{\varphi}(\pi_I)}{\check{\varphi}(\pi_A)}\right) G_A(\pi_A) - 2G_I(\pi_I) \\ &= \left(\frac{\hat{\varphi}(\pi_I)}{\hat{\varphi}(\pi_A)} + \frac{\check{\varphi}(\pi_I)}{\check{\varphi}(\pi_A)}\right)\left(\hat{A}\hat{\varphi}(\pi_A) + \check{A}\check{\varphi}(\pi_A)\right) - 2\left(\hat{A}\hat{\varphi}(\pi_I) + \check{A}\check{\varphi}(\pi_I)\right) \\ &= \hat{A}\frac{\hat{\varphi}(\pi_A)\check{\varphi}(\pi_I) - \hat{\varphi}(\pi_I)\check{\varphi}(\pi_A)}{\check{\varphi}(\pi_A)} + \check{A}\frac{\hat{\varphi}(\pi_I)\check{\varphi}(\pi_A) - \hat{\varphi}(\pi_A)\check{\varphi}(\pi_I)}{\hat{\varphi}(\pi_A)} \\ &= \underbrace{(\hat{\varphi}(\pi_A)\check{\varphi}(\pi_I) - \hat{\varphi}(\pi_I)\check{\varphi}(\pi_A))}_{<0}\left[\frac{\hat{A}}{\check{\varphi}(\pi_A)} - \frac{\check{A}}{\hat{\varphi}(\pi_A)}\right] \\ &= \frac{\hat{\varphi}(\pi_A)\check{\varphi}(\pi_I) - \hat{\varphi}(\pi_I)\check{\varphi}(\pi_A)}{\check{\varphi}(\pi_A)\hat{\varphi}(\pi_A)} \underbrace{\left(\hat{A}\hat{\varphi}(\pi_A) - \check{A}\check{\varphi}(\pi_A)\right)}_{<0} \\ &> 0, \end{aligned}$$

where the last step follows from a similar procedure as in (B.2) and observing that $G_A'(\pi_A) < 0$.

So, in order to restore optimality after an increase in γ, the abandonment trigger π_A must be increased. A similar analysis shows that

$$\frac{\partial \hat{V}(\pi_A; \pi_I)}{\partial \gamma} > 0,$$

so that π_I must be decreased after an increase in γ.

The final two results now follow by observing that γ is increasing in σ and decreasing in μ. ■

C Proof of Proposition 3

For all $\pi \in (0,1)$, define

$$\begin{aligned} F(\pi) &:= \max(F_A(\pi), F_I(\pi)), \\ F^*_\delta(\pi) &:= \sup_{\tau_\delta \in \mathcal{T}_\delta} \mathsf{E}_\pi\left[-\int_0^{\tau_\delta} e^{-rt}(c + P\pi_t)dt + e^{-r\tau_\delta}F(\pi_{\tau_\delta})\right], \\ \tilde{F}_\delta(\pi) &:= \mathsf{E}_\pi\left[-\int_0^{\delta} e^{-rt}(c + P\pi_t)dt + e^{-r\delta}F(\pi_\delta)\right], \end{aligned} \tag{C.1}$$

and

$$\tilde{F}^*(\pi) := \sup_{\tau \in \mathcal{T}_0} \mathsf{E}_\pi\left[-\int_0^{\tau} e^{-rt}(c + P\pi_t)dt + e^{-r\tau}\tilde{F}_\delta(\pi_\tau)\right]. \tag{C.2}$$

From Øksendal and Sulem (2007, Proposition 2.11) it follows that $F^*_\delta = \tilde{F}^*$ and that $\tau^*_\delta \in \mathcal{T}_\delta$ is an optimal stopping time for (C.1) if, and only if, $\tau^*_\delta = \tau^* + \delta$, where $\tau^* \in \mathcal{T}_0$ is an optimal stopping time for (C.2).

It holds that

$$\begin{aligned} \tilde{F}_\delta(\pi) &= \mathsf{E}_\pi\left[-\int_0^{\delta} e^{-rt}(c + P\pi_t)dt\right] + e^{-r\delta}\mathsf{E}_\pi[F(\pi_\delta)] \\ &= e^{-r\delta}\left\{\frac{c + P\pi}{r} + \mathsf{E}_\pi\left[1_{\pi_\delta > \bar{\pi}}F_I(\pi_\delta) + 1_{\pi_\delta < 0\bar{\pi}}F_A(\pi_\delta)\right]\right\} \\ &= e^{-r\delta}\Big\{\frac{c + P\pi}{r} + (B_1 - B_0)\mathsf{E}_\pi[\pi_\delta 1_{\pi_\delta > \bar{\pi}}] - L\mathsf{E}_\pi[\pi_\delta 1_{\pi_\delta > \bar{\pi}}] \\ &\quad + (B_0 - I)\mathsf{P}_\pi(\pi_\delta > \bar{\pi}) + B\mathsf{P}_\pi(\pi_\delta < \bar{\pi})\Big\}, \end{aligned}$$

where the second equality follows because $(\pi_t)_{t\geq 0}$ is a martingale.

The expectations and probabilities in the last two lines of the above equation are derived in the following lemma, the proof of which can be found in Appendix D.[16]

Lemma 1. *Define*

$$b_\pm := \frac{\sigma}{\mu\sqrt{\delta}}\left[\log\left(\frac{(1-\bar{\pi})\pi}{\bar{\pi}(1-\pi)}\right) \pm \frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right] \quad \textit{and} \quad d_\pm := \frac{\sigma}{\mu\sqrt{\delta}}\left[\log\left(\frac{(1-\pi)\bar{\pi}}{\pi(1-\bar{\pi})}\right) \pm \frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right],$$

and denote the cdf of the standard normal distribution by Φ. *Then*

$$\begin{aligned} &E_\pi\left[\pi_\delta 1_{\pi_\delta < \bar{\pi}}\right] = \pi\Phi(d_-), \quad P_\pi(\pi_\delta < \bar{\pi}) = \pi\Phi(d_-) + (1-\pi)\Phi(d_+), \\ &E_\pi\left[\pi_\delta 1_{\pi_\delta > \bar{\pi}}\right] = \pi\Phi(b_+), \quad \textit{and} \quad P_\pi(\pi_\delta > \bar{\pi}) = \pi\Phi(b_+) + (1-\pi)\Phi(b_-). \end{aligned}$$

[16] We are grateful to Tomasz Zastawniak for his help in establish this proof.

Therefore,

$$F_\delta(\pi) = e^{-r\delta}\Big\{\frac{c+P\pi}{r} + (B_1 - B_0)\pi\Phi(b_+) - L\pi\Phi(d_-) \\ + (B_0 - I)[\pi\Phi(b_+) + (1-\pi)\Phi(b_-)] + B[\pi\Phi(d_-) + (1-\pi)\Phi(d_+)]\Big\}.$$

Using the tower property, we can then write:

$$\begin{aligned} F^*_\delta(\pi) &= \sup_{\tau\in\mathcal{T}_0} \mathsf{E}_\pi\left[-\int_0^\tau e^{-rt}(c+P\pi_t)dt + e^{-r\tau}F_\delta(\pi_\tau)\right] \\ &= -\frac{c+P\pi}{r} + \sup_{\tau\in\mathcal{T}_0}\mathsf{E}_\pi[e^{-r\tau}F_\delta(\pi_\tau)]. \end{aligned}$$

This is a standard optimal stopping problem of the form studied in Appendix A. Therefore, the same conditions for optimality apply. That is, the state space, $[0,1]$, can be spit into a continuation region $\mathcal{C}$ and a stopping region $\Gamma := [0,1]\setminus\mathcal{C}$. As mentioned in Remark 6, the "smooth-pasting" conditions (12) and (13) are necessary, but not sufficient. Sufficiency is provided by superharmonicity of the value function. In the continuation region, superharmonicity follows by construction. In the case of Proposition 1, superharmonicity on Γ follows straightforwardly from the piece-wise linear nature of the function G. In the model with delay, however, this may not be the case as we could not obtain an analytic verification due to the highly non-linear nature of F_δ. What we can say, though, is that, since[17]

$$\limsup_{\pi\downarrow 0}\frac{F_\delta(\pi)}{\breve{\varphi}(\pi)} < \infty, \quad \text{and} \quad \limsup_{\pi\uparrow 1}\frac{F_\delta(\pi)}{\hat{\varphi}(\pi)} < \infty,$$

it follows from Dayanik and Karatzas (2003, Proposition 5.14) that there exist no $\pi_L \in (0,1)$, nor $\pi_H \in (0,1)$, such that $(0,\pi_L)\subset\mathcal{C}$ or $(\pi_H,1)\subset\mathcal{C}$. It appears that, in our numerical example reported in Remark 6, we find that, $\mathcal{C}=\emptyset$ and, thus, that the optimal stopping time is $\tau^*=0$. Whether this is indeed the case will be left to future research. Suffice to say here that one has to be careful in relying solely on smooth-pasting conditions. ■

D Proof of Lemma 1

1. Fix $\pi\in(0,1)$ and consider the process $(\pi_t)_{t\geq\pi_0}$, with $\pi_0=\pi$, P_π-a.s.. To simplify notation we will drop the subscript π from now on. Consider a solution π_t of the stochastic differential equation (SDE)

$$d\pi_t = \frac{\mu}{\sigma}\pi_t(1-\pi_t)dW_t,$$

[17] It holds that $F_\delta(0) = e^{-r\delta}F(0)$ and $F_\delta(1) = e^{-r\delta}F(1)$. In addition $1/\breve{\varphi}\to 1$ as $\pi\downarrow 0$ and $1/\hat{\varphi}\to 1$ as $\pi\uparrow 1$.

where W_t is a standard Brownian motion under P. Define, for all $t \geq 0$,

$$Z_t := \pi_0 \exp\left\{ \int_0^t \frac{\mu}{\sigma}(1-\pi_s) dW_s - \frac{1}{2}\int_0^t \frac{\mu^2}{\sigma^2}(1-\pi_s)^2 ds \right\}.$$

It then follows from Ito's lemma that Z_t follows the SDE

$$\begin{aligned} dZ_t &= Z_t \left[\frac{\mu}{\sigma}(1-\pi_t) dW_t - \frac{1}{2}\frac{\mu^2}{\sigma^2}(1-\pi_t)^2 dt \right] + \frac{1}{2} Z_t \frac{\mu^2}{\sigma^2}(1-\pi_t)^2 dt \\ &= \frac{\mu}{\sigma} Z_t (1-\pi_t) dW_t, \end{aligned}$$

with $Z_0 = \pi_0$. Because of uniqueness of solutions to SDEs, it holds that $Z_t = \pi_t$, P-a.s.. Therefore, it holds that

$$\pi_t = \pi_0 \exp\left\{ \int_0^t \frac{\mu}{\sigma}(1-\pi_s) dW_s - \frac{1}{2}\int_0^t \frac{\mu^2}{\sigma^2}(1-\pi_s)^2 ds \right\}.$$

Now define an equivalent measure Q, by the Radon-Nikodym derivative

$$\frac{d\mathsf{Q}}{d\mathsf{P}} = \exp\left\{ \int_0^t \frac{\mu}{\sigma}(1-\pi_s) dW_s - \frac{1}{2}\int_0^t \frac{\mu^2}{\sigma^2}(1-\pi_s)^2 ds \right\} = \frac{\pi_t}{\pi_0},$$

and let

$$B_t := W_t - \int_0^t \frac{\mu}{\sigma}(1-\pi_s) ds.$$

It then follows from the Girsanov theorem that B_t is a standard Brownian motion under Q.

2. It, therefore, holds that

$$\mathsf{E}_\mathsf{P}[\pi_\delta 1_{\pi_\delta < \tilde{\pi}}] = \pi_0 \mathsf{E}_\mathsf{P}\left[\frac{\pi_t}{\pi_0} 1_{\pi_\delta < \tilde{\pi}} \right] = \pi_0 \mathsf{E}_\mathsf{P}\left[\frac{d\mathsf{Q}}{d\mathsf{P}} 1_{\pi_\delta < \tilde{\pi}} \right] = \pi_0 \mathsf{E}_\mathsf{Q}[1_{\pi_\delta < \tilde{\pi}}] = \mathsf{Q}(1_{\pi_\delta < \tilde{\pi}}).$$

Note that

$$d\pi_t = \frac{\mu}{\sigma}\pi_t(1-\pi_t) dB_t + \frac{\mu^2}{\sigma^2}\pi_t(1-\pi_t)^2 dt,$$

so that an application of Ito's lemma gives

$$\begin{aligned} d\log\left(\frac{\pi_t}{1-\pi_t}\right) &= d\log(\pi_t) - d\log(1-\pi_t) \\ &= \frac{1}{\pi_t} d\pi_t - \frac{1}{2}\frac{1}{\pi_t^2} d\pi_t d\pi_t + \frac{1}{1-\pi_t} d\pi_t + \frac{1}{(1-\pi_t)^2} d\pi_t d\pi_t \\ &= \frac{1}{2}\frac{\mu^2}{\sigma^2} dt + \frac{\mu}{\sigma} dB_t. \end{aligned}$$

Therefore,

$$\log\left(\frac{\pi_t}{1-\pi_t}\right) = \log\left(\frac{\pi_0}{1-\pi_0}\right) + \frac{1}{2}\frac{\mu^2}{\sigma^2} t + \frac{\mu}{\sigma} B_t.$$

It then immediately follows that

$$\frac{\pi_t}{1-\pi_t} = \frac{\pi_0}{1-\pi_0}\exp\left\{\frac{1}{2}\frac{\mu^2}{\sigma^2}t + \frac{\mu}{\sigma}B_t\right\},$$

and, thus, that

$$\pi_\delta < \bar{\pi} \iff \frac{B_\delta}{\sqrt{\delta}} < \frac{\log\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0} - \frac{1}{2}\frac{\mu^2}{\sigma^2}\delta}{(\mu/\sigma)\sqrt{\delta}}.$$

Since B_t is a Brownian motion under Q, it now follows that

$$\mathsf{E}_\mathsf{P}\left[\pi_\delta 1_{\pi_\delta<\bar{\pi}}\right] = \pi_0\mathsf{Q}(1_{\pi_\delta<\bar{\pi}}) = \pi_0\Phi\left(\frac{\log\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0} - \frac{1}{2}\frac{\mu^2}{\sigma^2}\delta}{(\mu/\sigma)\sqrt{\delta}}\right).$$

3. Using the same Radon-Nikodym derivative as before and applying the Girsanov theorem, we can write

$$\mathsf{P}(\pi_\delta < \bar{\pi}) = \mathsf{E}[1_{\pi_\delta} < \bar{\pi}] = \pi_0\mathsf{E}_\mathsf{Q}\left[\frac{1}{\pi_\delta}1_{\pi_\delta<\bar{\pi}}\right].$$

This expectation can be computed explicitly in the following steps, which only use straightforward algebra:

$$\begin{aligned}
\pi_0\mathsf{E}_\mathsf{Q}\left[\frac{1}{\pi_\delta}1_{\pi_\delta<\bar{\pi}}\right] &= \pi_0\mathsf{E}_\mathsf{Q}\left[\frac{1-\pi_0+\pi_0 e^{\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta+\frac{\mu}{\sigma}B_\delta}}{\pi_0 e^{\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta+\frac{\mu}{\sigma}B_\delta}}1_{B_\delta<\frac{\sigma}{\mu}\left(\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0}-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right)}\right] \\
&= (1-\pi_0)\mathsf{E}_\mathsf{Q}\left[e^{-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta-\frac{\mu}{\sigma}B_\delta}1_{B_\delta<\frac{\sigma}{\mu}\left(\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0}-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right)}\right] + \pi_0\mathsf{E}_\mathsf{Q}\left[1_{B_\delta<\frac{\sigma}{\mu}\left(\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0}-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right)}\right] \\
&= (1-\pi_0)\int_{-\infty}^{\frac{\sigma}{\mu}\left(\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0}-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right)} e^{-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta-\frac{\mu}{\sigma}s}\frac{1}{\sqrt{2\pi t}}e^{-s^2/2t}ds + \pi_0\Phi\left(\frac{\sigma}{\mu}\left(\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0}-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right)\right) \\
&= (1-\pi_0)\int_{-\infty}^{\frac{\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0}+\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta}{(\mu/\sigma)\sqrt{\delta}}} \frac{1}{\sqrt{2\pi}}e^{-z^2/2}dz + \pi_0\Phi\left(\frac{\sigma}{\mu}\left(\frac{\bar{\pi}(1-\pi_0)}{(1-\bar{\pi})\pi_0}-\frac{1}{2}\frac{\mu^2}{\sigma^2}\delta\right)\right) \\
&= \pi_0\Phi(d_-) + (1-\pi_0)\Phi(d_+).
\end{aligned}$$

4. A similar analysis gives the results for $\mathsf{E}_\mathsf{P}\left[\pi_\delta 1_{\pi_\delta>\bar{\pi}}\right]$ and $\mathsf{P}(\pi_\delta > \bar{\pi})$. ■

References

Arlotto, A., Chick, S., and Gans, N. (2014) Optimal hiring and retention policies for heterogeneous workers who learn. *Management Science* **60**, 110–129.

Armitage, P. (1975) *Sequential Medical Trials*. Blackwell, Oxford.

Bergemann, D. and Valimaki, J. (2006) Bandit problems. *Cowles Foundation discussion paper No 1551* 1–16.

Berry, D. (1985) Interim analyses in clinical trials: classical vs. Bayesian approaches. *Statistics in Medicine* **4**, 521–526.

Berry, D. (2006) Bayesian clinical trials. *Nature Review Drug Discovery* **5**, 27–36.

Bolton, P. and Harris, C. (1999) Strategic experimentation. *Econometrica* **67**, 349–374.

Borodin, A. and Salminen, P. (2002) *Handbook of Brownian Motion – Facts and Formulae*. Birkhäuser Verlag, Basel, second edition.

Chick, S., Forster, M., and Pertile, P. (2015) A Bayesian-theoretic model of sequential experimentation with delayed response. DERS Discussion Paper 15/09, University of York.

Claxton, K., Martin, S., Soares, M., Rice, N., Spackman, E., Hinde, S., Devlin, N., Smith, P., and Sculpher, M. (2015) Methods for the estimation of the National Institute for Health and Care Excellence cost-effectiveness threshold. *Health Technology Assessment* **19**.

Close, A., Robertson, C., Rushton, S., Shirley, M., Vale, L., Ramsay, C., and Pickard, R. (2013) Comparative Cost-Effectiveness of Robot-Assisted and Standard Laparoscopic Prostatectomy as Alternatives to Open Radical Prostatectomy for Treatment of Men with Localised Prostate Cancer: A Health Technology Assessment from the Perspective of the UK National Health Service. *European Urology* **64**, 361–369.

Dayanik, S. and Karatzas, I. (2003) On the optimal stopping problem for one-dimensional diffusions. *Stochastic Processes and their Applications* **107**, 173–212.

Dixit, A. and Pindyck, R. (1994) *Investment under Uncertainty*. Princeton University Press, Princeton.

Draper, D. (2013) Discussion of the paper by Hampson and Jennison. *Journal of the Royal Statistical Society B* **75**, 48.

Driffield, T. and Smith, P. (2007) A real option approach to watchful waiting: theory and illustration. *Medical Decision Making* **27**, 178–188.

FDA (2006) Fast track drug development programs designation, development, and application review. U.S department of Health and Human Services, Food and Drug administration. URL `http://www.fda.gov/downloads/Drugs/Guidances/ucm079736.pdf`.

FDA (2010) Adaptive design clinical trials for drugs and biologics. U.S department of Health and Human Services, Food and Drug administration. URL http://www.fda.gov/downloads/Drugs/Guidances/ucm201790.pdf.

FDA (2013) Expedited programs for serious conditions drugs and biologics. U.S department of Health and Human Services, Food and Drug administration. URL http://www.fda.gov/downloads/Drugs/GuidanceComplianceRegulatoryInformation/Guidances/UCM358301.pdf.

Forster, M. and Pertile, P. (2013) Optimal decision rules for HTA under uncertainty: a wider, dynamic perspecitve. *Health Economics* **22**, 1507–1514.

Hampson, L. and Jennison, C. (2013) Group sequential tests for delayed responses. *Journal of the Royal Statistical Society B* **75**, 3–54.

Heiodhues, P., Rady, S., and Strack, P. (2015) Strategic experimentation with private payoffs. *Journal of Economic Theory* **159**, 531–551.

Jovanovic, B. (1979) Job search and the theory of turnover. *Journal of Political Economy* **87**, 972–990.

Keller, G. and Rady, S. (2010) Strategic experimentation with Poisson bandits. *Theoretical Economics* **5**, 275–311.

Keller, G., Rady, S., and Cripps, M. (2005) Strategic experimentation with exponential bandits. *Econometrica* **73**, 39–68.

Kwon, H. and Lipmann, S. (2011) Acquisition of Project-Specific Assets with Bayesian Updating. *Operations Research* **59**, 1119–1130.

Meltzer, D. and Smith, P. (2012) Theoretical issues relevant to the economic evaluation of health technologies. In Pauly, M., Mcguire, T., and Barros, P. (eds.), *Handbook of Health Economics*, volume 2, 433–469. Elsevier, Amsterdam.

Mikhalevich, V. (1958) Bayeisan Choice between Two Hypotheses for the Mean Value of a Normal Process. *Visnik Kiiv. Univ.* **1**, 101–104. In Ukranian.

Montazerhodjat, V. and Lo, A. (2015) Is the FDA too conservative of too aggressive ? A Bayesian decision analysis of clinical trial design. NBER working paper 21499.

Moscarini, G. and Smith, L. (2001) The optimal level of experimentation. *Econometrica* **69**, 1629–1644.

Myers, S. (1977) Determinants of Corporate Borrowing. *Journal of Financial Economics* **5**, 147–176.

Øksendal, B. (2004) Optimal stopping with delayed information. *Stochastics and Dynamics* **5**, 271–280.

Øksendal, B. (2010) *Stochastic Differential Equations*. Springer, New York, fifth edition.

Øksendal, B. and Sulem, A. (2007) *Applied Stochastic Control of Jump Diffusions*. Springer–Verlag, Berlin, second edition.

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

Pertile, P., Forster, M., and LaTorre, D. (2014) Optimal Bayesian sequential sampling rules for the economic evaluation of health technologies. *Journal of the Royal Statistical Society A* **177**, 419–438.

Peskir, G. and Shiryaev, A. (2006) *Optimal Stopping and Free Boundary Problems*. Birkhäuser Verlag, Basel.

Poor, V. and Hadjiliadis, O. (2009) *Quickest Detection*. Cambridge University Press, Cambridge.

Ramsey, S., Wilkie, R., Briggs, A., Brown, R., Buxton, M., Chawla, A., Cook, J., Glick, H., Petitti, D., and Reed, S. (2005) Good research practices for cost-effectivness analysis alongside clinical trials: the ISPOR RCT-CEA task force report. *Value in Health* **8**, 521–533.

Rothschild, M. (1974) Two armed bandit theory of market pricing. *Journal of Economic Theory* **9**, 185–202.

Samsa, G., Edelman, D., Rothman, M., Williams, G., Lipscomb, J., and Matchar, D. (1999) Determining clinically important differences in health status measures: a general approach with illustration to the Health Utilities Index Mark II. *Pharmacoeconomics* **15**, 141–155.

Shiryaev, A. (1967) On Two Problems of Sequential Analysis. *Kibernetika* **2**, 79–80.

Shiryaev, A. (1978) *Optimal Stopping Rules*. Springer–Verlag, Berlin.

Spiegelhalter, D., Abrams, K., and Myles, J. (2004) *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. John Wiley & Sons, New York.

Steele, J. (2001) *Stochastic Calculus and Financial Applications*. Springer–Verlag, New York.

Stokey, N. (2009) *The Economics of Inaction*. Princeton University Press, Princeton.

Sundaram, R. (2005) *Generalized bandit problems*. Springer–Verlag, Berlin.

Wald, A. (1947) *Sequential Analysis*. John Wiley & Sons, London.

William, A. and Pinto, E. (2005) The value of information and optimal clinical trial design. *Statistics in Medicine* **24**, 1791–1806.

YHEC (2009) Organisational and Behavioural Barriers to Medical Technology Adoption. Report by York Health Economics Consortium for NHS Institute for Innovation and Improvement.