The Impact of Ambiguity over Clinical Status as a Reliable Predictor of Treatment Outcomes on Healthcare Decision Making and Management

Laura Delaney *

September 27, 2023

Abstract

This paper considers the impact of *Knightian uncertainty* or *ambiguity* about the reliability of a patient's acquired comorbidities and risk factors as predictors of treatment outcome on the optimal time to initiate treatment. I show that high levels of such ambiguity is detrimental to patient welfare. Hence, learning about the clinical state (comprised of the comorbidities and risk factors) as a predictor of treatment outcome in order to resolve, at least partially, this ambiguity is crucial to improving their welfare.

The learning is achieved via a sequential hypothesis test in which the clinician will only treat if her ambiguity about outcome is sufficiently low; i.e., below some threshold which is derived based on the cost of making a wrong decision by treating. I show that learning in this way does indeed improve patient welfare with respect to the optimal treatment (timing) strategy.

The paper concludes with a discussion on the practical considerations for clinicians, how they can use these results in managing patient care, and notes that the results support specialisation across hospitals so that certain treatments are only carried out a small number of specialist centres.

Keywords: Decision analysis; Ambiguity; Sequential hypothesis testing.

1 Introduction

It is the job of the healthcare provider (hereafter referred to as the clinician) to ensure that the decisions she makes about her patients' care are optimal. This can only be achieved if patients are receiving the right treatment at the most appropriate time. In this paper, the treatment that is being given is not in question, but rather the timing of initiating it. Determining the optimal time to treat a patient is a critical part of their care. Its importance is recognised by clinicians, but determining the timing remains a challenge for many. For example, Crossland et al. [2019] conduct a study of cardiac transplantation in a group of patients born with a

*Kings College London, Kings Business School, 30 Aldwych, London, WC2B 4BG, UK. Email: laura.delaney@kcl.ac.uk

particular congenital cardiac condition. They point out that "the optimal timing for listing and transplanting these patients is key to improving outcomes". The point is echoed in Kenny et al. [2018] who say that determining this optimal time is challenging and having a greater "understanding of this can help guide decision making" with regard to treatment. Tretter and Redington [2018] analyse the treatment of pulmonary valve replacement (PVR) in patients with a different congenital condition. They point out that for many patients with that condition, PVR is inevitable, but that "it is, of course, all about timing".

Delaney [2021] addresses the timing issue in healthcare treatment and develops a model in which two thresholds on a patient's clinical status score are derived based on real options techniques. If the patient's score is above the upper threshold, the patient is too well for treatment now, and if it is below the lower threshold, the patient is too unwell for treatment. However, if their clinical status score is between the two thresholds, the optimal strategy is to treat the patient. That paper was anchored in the example of cardiac transplantation for patients born with a cardiac condition discussed in Crossland et al. [2019]. However, the model is generalisable and is currently being trialled for use in determining the optimal time for PVR in an entirely different population of patients.

In that model (Delaney [2021]), the clinician is considering whether to provide risky life-saving treatment to the patient and the patient's acquired comorbidities and risk factors define his clinical state. On one hand, the treatment is very risky so the patient needs to be in a clinical state that is sufficiently low to justify the risk of treatment, but on the other hand, if the patient is too unwell and the clinical state is too low, the treatment outcome will most likely be bad. The clinician's objective is to determine the optimal time, with respect to clinical state, to perform the treatment such that the expectation of a good outcome is maximised.

Furthermore, the patient's clinical state serves as a signal to the clinician about his likely outcome from treatment, and it is assumed that the signal is only an accurate reflection of the true outcome with some probability less than one. Indeed, in the context of healthcare treatment, this is a natural assumption as patients often do not respond to treatment. As Tretter and Redington [2018] say, "If there were no adverse consequences to PVR, the timing question would be irrelevant".

Now say s denotes a particular clinical state and say there are a number of different hospitals all providing the treatment being considered. Moreover, say that for each hospital, retrospective data shows that in their cohort of patients who had the treatment in state s, the proportion of patients who did well from treatment varied across the hospitals. For example, Berg et al. [2017] conducted an analysis to identify a set of preoperative risk factors and comorbidities that were associated with a greater risk of post-operative mortality in patients who received a cardiac transplant at UCLH Medical Center between 1991 and 2014. They developed a scoring system based on the hazard ratio associated with each risk factor and comorbidity and patients were scored according to what risk factors and comorbidities they had acquired. This score defined the patient's clinical state and patients were stratified according to their state.

The study was replicated by Polyviou et al. [2018] who conducted the analysis on a patient cohort who also received cardiac transplant over a similar period in a UK hospital. They scored

and stratified their patients according to the same criteria as in the Berg et al. [2017] study. However, the results on postoperative mortality according to each strata (i.e., clinical state) were very heterogeneous across the two cohorts. For example, for patients in the worst clinical state, mortality in the Berg et al. [2017] study was 88%, but only 16% in the Polyviou et al. [2018] study, whereas in the group of patients in the best clinical state, the mortality was 0% and 33%, respectively.

Since it is reasonable to presume that the treatment will only be performed if the clinician is sufficiently convinced that the outcome will be successful, the probability that s is a correct reflection of the true outcome is equivalent to the success rate conditional on s. If this success rate is heterogeneous across hospitals, as in the example mentioned above, the true (conditional) success probability *cannot be known with certainty*.

In Delaney [2021], as in many standard real options models, the underlying assumption is that the decision maker is perfectly certain about the probability measure characterising the uncertainty over the future outcome. In other words, there is no distinction between *risk*, where the probabilities of the outcomes of an uncertain event can be measured, and *uncertainty*, where such probabilities cannot be measured. This distinction is emphasised by Knight [1921] who stresses that uncertainty is more realistic in decision making settings. As such, such uncertainty has become widely known in the academic literature as *Knightian uncertainty* or *ambiguity*. Indeed, as the above example illustrates, ambiguity is present in clinical contexts and, in particular, over the reliability of a patient's clinical state as a predictor of outcome from treatment (i.e., over signal quality).

In this paper, I examine the impact of this ambiguity over signal quality on the optimal treatment strategy of the clinician. The main motivation for this is clear from the example described; i.e., that the perfect certainty assumption over the signal quality is lacking in practical realism.

In the next section, I present the model with ambiguity. It is closely related to the model in Delaney [2021] but, instead, I assume that a signal arrives every period. This is to simplify the analysis sufficiently so that the effect of ambiguity on the optimal treatment threshold can be ascertained. I find that if the clinician has a high degree of ambiguity over the patient's clinical status as a predictor of treatment outcome (signal quality), the optimal strategy should be to treat him even if she does not expect him to do well from treatment. The reason is that ambiguity over signal quality reduces the value of waiting to obtain further signals on outcome and, hence, more immediate treatment is optimal before the patient's clinical state deteriorates further. However, this optimal strategy is clearly questionable and counter-intuitive. Hence, assuming the clinician adheres to the optimal treatment strategy, I examine the effect of ambiguity on patient welfare.

Importantly, I find that ambiguity has a negative impact on patient welfare because patients may be treated in very poor clinical states if the clinical state is not a reliable predictor of outcome. As such, this highlights the importance of reducing ambiguity over clinical state as a predictor of outcome if such an optimal strategy is adhered to. However, as I argue in the next paragraph, the optimal strategy itself is not to be dismissed and has the potential to be of huge

value in practical clinical decision making contexts.

The optimal strategy is one of timing and is derived using techniques from real options analysis. Real options analysis is more commonly used in corporate investment decision making contexts and is underpinned by three important characteristics (i) uncertainty over outcome, (ii) irreversibility and (iii) flexibility over timing. In the context of medical decision making and patient care management, (i) outcome is uncertain, (ii) if the patient does not do well from treatment (for example, they die after surgery), this cannot be undone, and (iii) "watchful waiting" is very important for clinicians as it "increases information upon which a clinical decision can be made" (Driffield and Smith [2007]). Therefore, it is appropriate to use such techniques in medical decision making. Driffield and Smith [2007] give examples of the application of real options in medical decision making for glue ear and small abdominal aortic aneurysms. However, they make the point that in those studies, deferral is not properly modelled because treatment cannot be initiated at any time if the patient deteriorates quickly. The model in Delaney [2021] (of which the model in this paper is a special case) does not have such stipulations over treatment timing. Most importantly, it addresses the needs called for in the medical literature for such a timing rule. As well as the points cited above, Polyviou et al. [2018] says "there is a need for tools to help guide decision-making ..." (the context being with regard to timing) and Geva et al. [2018] "Our results highlight several observations pertaining to the timing of PVR... (clinical markers) do not fully define the inflection point of individual predictors... From the clinician's perspective, it would be helpful to define the window during which the risk of poor outcomes begins to increase while the disease process has not transitioned from reversible to irreversible... it would be ideal not to implant a PVR too early in asymptomatic patients with stable disease who are at low risk of adverse outcomes". The latter cites further evidence in support of such a timing rule (in the interest of space, see references therein). The optimal timing strategy derived in Delaney [2021] is exactly as called for in Geva et al. [2018].

The model in this paper is a special case of the model in Delaney [2021]. The difference is that in this model, the patient's clinical status changes continuously. Since age is a risk factor for many treatments of serious illness, it could be argued that, indeed, a patient's clinical status is continuously changing as they are always ageing. However, in practical terms, it is difficult to apply the model with this assumption. But, on the other hand, making this restrictive assumption allows for the impact of ambiguity to be easily determined. Once its effect is known, which as pointed out above, is that it has a negative effect on patient welfare, the next step is to find a practical means of reducing ambiguity sufficiently so that the model in Delaney [2021] can be used in clinical decision-making. The second part of this paper focuses on this task.

The idea is that the clinician will only treat a patient in a particular clinical state if her ambiguity over that clinical state as a predictor of outcome is below a certain threshold. This ambiguity threshold is determined by solving for another optimal stopping problem in which the clinician is learning and researching if her ambiguity is above the threshold; i.e., it is not low. She will not treat the patient in that clinical state if she is learning to resolve ambiguity.

Contributions to the literature on ambiguity with active learning appear to be limited to

the studies of Epstein and Schneider [2007], Miao and Wang [2011] and Epstein and Ji [2020]. Epstein and Schneider [2007] present a framework of learning under ambiguity in a discrete time set-up and Miao and Wang [2011] apply the framework to a job search problem. Epstein and Ji [2020], however, adapt the framework of Epstein and Schneider [2007] to a continuous time setting, but the optimal stopping problem they address differs markedly from mine. The optimal stopping problem over learning in this paper is more closely related to Peskir and Shiryaev [2006] (Chapter 21) in that it addresses the problem of learning to resolve ambiguity as a sequential hypothesis test with the objective to minimise loss. In the context of medical decision making, this is appropriate as the loss from learning and, importantly, of making a wrong decision by treating a patient that is too well or too unwell, is a negative impact on their quality of life. I provide the technical solution for the rule and, subsequently, in the section that follows, I discuss the applicability of the rule in clinical settings.

Throughout the paper, wherever appropriate, I provide accompanying intuitive examples to support and explain the technicalities of the techniques used. The remainder of the paper is organised as follows. In the next section, I present the set-up for the model. In Section 3 I present the optimal treatment strategy in the case of ambiguity and in Section 4, I examine the effect of this ambiguity on patient welfare by adhering to the optimal treatment strategy derived. In Section 5, I address the issue of learning to resolve ambiguity and improve patient welfare and discuss how this can be achieved. Concluding remarks about the practical implications of the results and what they mean for practising clinicians are considered in Section 6. The proofs of all the main results are placed in the Appendix.

2 The Model

2.1 General Set-Up and Ambiguity-Free Dynamics

Consider a clinician with the option to treat a patient whose outcome from treatment can be Good (G), leading to an improvement in quality of life, or Bad (B), leading to a worsening in quality of life. The patient's current quality of life is $I > 0$ and this represents a cost of treatment (cf. Delaney [2021]); i.e., the patient gives up this quality of life for a better or worse state. If the outcome is G, the patient's quality of life improves by an amount of at least $V^G > 0$, but if the outcome is B, the improvement is $V^B < 0$. For expositional ease, I let $V^B = 0$ hereafter. The problem for the clinician is to determine the optimal time to treat so that the patient's expected quality of life from treatment is maximised.

Let time be continuous and indexed by $t \geq 0$. Ex ante, the clinician is uncertain about the treatment outcome but, at discrete intervals, the patient develops (or recovers from) comorbidities and risk factors. The comorbidities and risk factors that the patient has at time t defines his clinical state at time t and this is a signal of the likely outcome from treatment. The uncertainty about each outcome is modelled on a probability-statistical space $(\Omega; \mathcal{F}; P_\pi, \pi \in [0,1])$ and the process $(\pi_t)_{t\geq 0}$ represents the probability that the outcome is G. Let γ be a random variable such that $\gamma = 1$ if the true outcome is G and $\gamma = 0$ if the true outcome is B. Further let P_1 and P_0 be degenerate distributions such that $P_1(\gamma = 1) = 1$ and $P_0(\gamma = 0) = 1$; i.e., representing

the true outcomes of G and B, respectively. Therefore, the structure of P_π is as follows:

$$P_\pi = \pi P_1 + (1-\pi)P_0.$$

A patient's clinical state (signal), however, is not always an accurate reflection of what the true outcome ends up being; i.e., a signal that is interpreted as being indicative of G may actually result in B and, as such, is deemed an incorrect signal. I let λ denote the probability that the signal is a correct reflection of the true outcome; i.e., if the signal is indicative of a G (B) outcome and the true outcome is G (B), then the signal is correct. Let P_γ^j for $j = \{G, B\}$ be degenerate distributions representing the true correctness of the signal such that the outcome is γ and the signal is indicative of a j outcome. Therefore

$$\lambda = P_1^G = P_0^B \text{ and } (1-\lambda) = P_0^G = P_1^B.$$

Because time is continuous, yet signals are received discretely, we need to construct (on the probability space) a discrete-time representation of the log-likelihood ratio process of signal arrivals $(L_t)_{t\geq 0}$ and then apply a random walk approximation to derive a Brownian motion driven stochastic differential equation describing the clinician's belief process in which the signals arrive continuously.

In Appendix A I derive the dynamics of the $(L_t)_{t\geq 0}$ as being (see also Dalby et al. [2018])

$$dL_t = (2\gamma - 1)\frac{\sigma_L^2}{2}dt + \sigma_L dW_t, \tag{1}$$

where σ_L is constant and $(W_t)_{t\geq 0}$ is a standard Brownian motion under P_π. The observable process $(L_t)_{t\geq 0}$ generates the filtration $\mathcal{F}^L = (\mathcal{F}_t^L)_{t\geq 0}$ which is augmented with the P_π-null sets; i.e., it is a sub-filtration of $\mathcal{F}$. If we view $\mathcal{F}$ as representing complete information, then $\mathcal{F}^L$ represents the available information and, since it is under the available information that the clinician makes her decision, the following assumption is crucial and common when dealing with incomplete information.

Let there be a standard Brownian motion $\widetilde{W} = (\widetilde{W}_t)_{t\geq 0}$ on the filtered probability space $(\Omega; \mathcal{F}^L; P_\pi, \pi \in [0,1])$ such that the augmented natural filtration generated by $\widetilde{W}$ is identical to $\mathcal{F}^L$ (Miao [2009]). This Brownian motion is typically referred to in the literature as the *innovation process*. In essence, what this implies is that if we solve the problem with respect to the filtered probability space, such that the standard Brownian motion representing uncertainty in the dynamics of L_t is $\widetilde{W}_t$, then we can solve the problem under complete information. This is known as the separation principle in the control literature (see, for example, Flemming and Rishel [1975]). It states that optimal control problems involving incomplete information can be solved separately as two independant problems: (i) of filtering and (ii) of control under complete information (Miao [2009]).

The first step is that of filtering and this involves transforming L_t so that its dynamics are represented entirely by the filtered probability space. By applying the procedure set out in the literature on filtering (e.g. Lipster and Shiryaev [1977]), replace $(2\gamma - 1)$ by its estimate

$2(\pi 1 + (1-\pi)0) - 1 = 2\pi - 1$ in Eq. (1) so that

$$\begin{aligned} dL_t =& (2\pi_t - 1)\frac{\sigma_L^2}{2}dt + (2\gamma - 1)\frac{\sigma_L^2}{2}dt + \sigma_L dW_t - (2\pi_t - 1)\frac{\sigma_L^2}{2}dt \\ =& (2\pi_t - 1)\frac{\sigma_L^2}{2}dt + \left(dL_t - (2\pi_t - 1)\frac{\sigma_L^2}{2}dt\right) \\ =& \sigma_L^2 \pi_t dt + \sigma_L \left(\frac{1}{\sigma_L}dL_t - \sigma_L \pi_t dt\right) \\ =& \sigma_L^2 \pi_t dt + \sigma_L d\widetilde{W}_t, \end{aligned} \tag{2}$$

where $d\widetilde{W}_t = \frac{1}{\sigma_L}dL_t - \sigma_L \pi_t dt$.

Assuming that the prior $t = 0$ probability in a G outcome is $\pi_0 = 1/2$, the *a posteriori probability process* $(\pi_t)_{t\geq 0}$ is expressed as

$$\pi_t = \frac{\varphi_t}{1 + \varphi_t} \tag{3}$$

such that the *likelihood ratio process* $(\varphi_t)_{t\geq 0}$ is defined by the Radon-Nikodym derivative

$$\varphi_t = \frac{d(P_1^G | \mathcal{F}_t^L)}{d(P_0^G | \mathcal{F}_t^L)} = e^{L_t - \frac{\sigma_L^2}{2}t}. \tag{4}$$

and, by an application of Ito's lemma, is found to solve the stochastic differential equation (cf. Peskir and Shiryaev [2006] pg. 188)

$$d\pi_t = \sigma_L \pi_t (1 - \pi_t) d\widetilde{W}_t. \tag{5}$$

2.2 Dynamics under Ambiguity

So far, it has been assumed that the clinician is *perfectly certain* that the probability of the patient's clinical state being a correct reflection of the true outcome from treatment is some constant λ. However, as previously explained, this is lacking in practical realism. In this subsection, I assume that the clinician lacks certainty in the signal quality parameter and, in this way, she is *ambiguous* over the probability measure generating the likelihood process $(L_t)_{t\geq 0}$ and, thus, over the posterior belief process $(\pi_t)_{t\geq 0}$.

2.2.1 Intuitive Example

Before proceeding with the technical modelling, I provide an intuitive example as follows. Say there are three hospitals H which perform this particular treatment so that $H = \{A, B, C\}$. Each hospital has data on the outcomes for patients that had the treatment in a particular clinical state t at that H. From the data of their patient cohort, they can determine a λ_t^H denoting the probability that the clinical state t was a correct predictor of the true outcome; eg., in hospital A, 70% of their patients treated in state t had a successful outcome, in B, 30% and in C, 56%. This depicts how ambiguity may arise over the quality of the clinical state as a predictor of treatment outcome.

Now a clinician treating a patient in clinical state t has a reference probability regarding the signal quality. Say her reference probability is $\lambda_t^A = 70\%$. This λ_t^A is associated with the reference probability measure P_π and with the associated density generator $\theta_t^A = 0$. The other probabilities λ_t^B and λ_t^C are associated with different density generators for state t, θ_t^B and θ_t^C, respectively. Later I provide technical detail on how the generators arise. An example for the set of density generators is as follows:

$$\Theta = \begin{bmatrix} \theta_1^A & \theta_1^B & \theta_1^C \\ \theta_2^A & \theta_2^B & \theta_2^C \\ \theta_3^A & \theta_3^B & \theta_3^C \end{bmatrix} \tag{6}$$

where the columns pertain to the hospital and the rows to the clinical state. One generator for each state will be the clinician's reference generator and, therefore, zero. However, if the values of the other generators differ widely from the zero, then this implies the probabilities representing signal quality are also very disparate across hospitals. Hence, the extent of the deviation characterises the clinician's degree of ambiguity.

2.2.2 Technical Modelling

Assume that for some process $\theta = (\theta_t)_{0 \leq t \leq \tau} \in \Theta$, a process $(z_t^\theta)_{0 \leq t < \tau}$ is the unique solution to

$$dz_t^\theta = -z_t^\theta \theta_t d\widetilde{W}_t, \tag{7}$$

with $z_0^\theta = 1$. The process θ is a density generator if z_t^θ is a martingale under P_π (Nishimura and Ozaki [2007]). It will be a martingale under P_π if it satisfies Novikov's condition that $E^{P_\pi}\left[\exp\left(\frac{1}{2}\int_0^T \theta_s^2 ds\right)\right] < \infty$. We assume then that θ is a density generator and therefore generates another probability measure $\mathcal{Q}^\theta$ from P_π via the Radon-Nikodym derivative

$$z_\tau^\theta = \frac{d\mathcal{Q}^\theta}{dP_\pi} \tag{8}$$

and, since z_t^θ is a P_π-martingale, it is shown in Nishimura and Ozaki [2007] that $\mathcal{Q}^\theta$ is equivalent to P_π. Moreover, any probability measure that is equivalent to P_π can be generated by some density generator via Eq. (8).

Let P_π^Θ denote the set of probability measures equivalent to P_π that are generated by θ according to (8) so that the set

$$P_\pi^\Theta := \{\mathcal{Q}^\theta | \theta \in \Theta\}.$$

Now, by Girsanov's theorem (cf. Karatzas and Shreve [1991], Theorem 5.1), $W_t^\theta = \widetilde{W}_t + \int_0^t \theta_s ds$ is a standard Brownian motion with respect to $\mathcal{Q}^\theta$ so that the dynamics of the posterior belief process (5) becomes, under the probability measure $\mathcal{Q}^\theta$,

$$d\pi_t = -\sigma_L \theta_t \pi_t (1 - \pi_t) dt + \sigma_L \pi_t (1 - \pi_t) dW_t^\theta, \tag{9}$$

with θ_t varying. Using this representation, the clinician now considers the belief dynamics for all

probability measures $\mathcal{Q}^\theta \in P_\pi^\Theta$ (equivalently, for all $\theta \in \Theta$), such that each probability measure/ density generator represents a particular signal quality.

However, in order to use this set Θ of density generators to characterise her ambiguity, we assume it (and the corresponding set of probability measures it generates P_π^Θ) is *strongly rectangular*[1]. The intuition for why this is an important assumption is explained in the next section when the optimal stopping problem has been defined. However, rectangularity implies that for each outcome $X \in \mathcal{L}^2(\Omega, \mathcal{F}^L, P_\pi)$, there exists a $\mathcal{Q}^{\theta*} \in P_\pi^\Theta$ such that for all $t \geq 0$,

$$E^{\mathcal{Q}^{\theta*}}[X|\mathcal{F}_t^L] = \min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}[X|\mathcal{F}_t^L]$$

(cf. Chen and Epstein [2002] Theorem 2.1). Rectangularity ensures dynamic consistency in the sense that the global minimisation over Θ is equivalent to repeated local minimisations over Θ.

Consequently, the uncertainty about the outcome being G or B is represented by a family of filtered probability spaces defined by $(\Omega; \mathcal{F}^L; P_\pi^\Theta)$.

As in much of the related literature (see, for example, Chen and Epstein [2002]; Nishimura and Ozaki [2007]; Trojanowska and Kort [2010]), I let the measures in P_π^Θ deviate only within a small neighbourhood around the reference measure P_π. The reference measure corresponds with the generator $\theta_t = 0$ (cf. Eq.(5) vs. Eq. (9)). Hence, let

$$\Theta := \{\theta = (\theta_t)_{t\geq 0} | \theta \in [-\kappa, \kappa]\} \tag{10}$$

in which $\kappa = 0$ represents no deviation from the reference measure and, hence, zero ambiguity. Higher values of κ imply a greater degree of ambiguity. This form of ambiguity is more generally known as κ-ignorance.

The final point that needs to be addressed before proceeding to the optimal stopping problem is the following. Under κ-ignorance, the Radon-Nikodym derivative defined by Eq. (8) satisfies Novikov's condition only when the option's life horizon is finite. In other words, equivalence of measures is not ensured over an infinite horizon and Girsanov's theorem does not apply. However, it is analytically much more tractable to solve for optimal stopping problems when options are infinitely-lived. Moreover, as pointed out in Trojanowska and Kort [2010], "In practice, it does not really matter whether the option to invest exists over a sufficiently long time horizon or an infinite one so long as the project has finite life time". In my model, the project is the treatment and it is finite in the sense that once performed, the outcome is known and fixed. As such, given that the set Θ is defined by (10), we follow the literature[2] by making the assumption that for every $\theta \in [-\kappa, \kappa]$, the solution to Eq. (7) is a martingale.

[1] A full technical definition and outline of the concept in a continuous time setting is provided in Chen and Epstein [2002] pp. 1411.

[2] See, for example, Nishimura and Ozaki [2007] and references therein.

3 The Optimal Stopping Problem

3.1 The Optimal Stopping Problem Set-Up

The clinician is assumed to be ambiguity averse which has been supported by experimental evidence since the seminal study by Ellsberg [1961][3]. To characterise this ambiguity aversion, I follow the approach commonly used in the literature (devised by Gilboa and Schmeidler [1989]) by assuming her objective is to optimally treat under the worst possible outcome. Note that there are situations in which this assumption of extreme ambiguity aversion is not appropriate. For example, Heath and Tversky [1991] show that decision makers who have a high degree of self-confidence seek out ambiguous situations and Bhide [2000] states "low ambiguity aversion of the individuals who start promising businesses derives from exceptionally high levels of self-confidence". Moreover, there are studies, see for example, Schroder [2011], which examine the effect of ambiguity attitudes on the optimal stopping strategy. However, in the context of medical decision making, when patients' health and quality of life are at stake, it is not appropriate for the clinician to be anything other than highly averse to ambiguity.

The utility function must permit the distinction between risk and ambiguity and, hence, has a recursive structure in line with the Chen and Epstein [2002] notion of *recursive multiple priors utility*. As such, the optimal stopping problem takes the form

$$\begin{aligned} V^*(\pi_t) :=& \max_{\tau>t} \min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[\left(\pi_\tau(\theta_\tau)(V^G - I) + (1-\pi_\tau(\theta_t))(0-I)\right)\Big|\mathcal{F}_t^L\right] \\ =& \max_{\tau>t} \min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[\left(\pi_\tau(\theta_\tau)V^G - I\right)\Big|\mathcal{F}_t^L\right] \end{aligned} \tag{11}$$

such that $\pi_t(\theta_t)V^G - I > 0$, τ the time of treatment, the minimum operator reflects her aversion to ambiguity, and $\pi_t(\theta_t)$ denotes her probabilistic belief in a G outcome at time t conditional on θ_t which is the density generator characterising her ambiguity over signal quality at time t. The belief dynamics are defined by Eq. (9).

As stated previously, the set of density generators Θ and the corresponding set of probability measures P_π^Θ are assumed to strongly rectangular. With respect to the hospital example described above, this assumption implies that all hospitals that perform this particular treatment are elements in Θ. Otherwise, there could be some other hospital D whose data is not considered by the clinician when making the decision ($\theta^D \notin \Theta$), but whose density generator generates another probability measure equivalent to the reference measure P_π. However, θ^D could yield the worst possible outcome for at least one clinical state t. For an ambiguity-averse clinician, it is important that this is not the case and the rectangularity assumption circumvents this by assuming there is no such hospital D; in other words, all possible hospitals performing the treatment are elements of Θ. It also ensures that the required dynamic consistency across all probability measures in the continuation region is satisfied. Considering a rectangular

[3]The Ellsberg Paradox suggests agents prefer to act on known rather than unknown ambiguous probabilities. This contradicts the Bayesian paradigm of a single probability measure underlying choices (Miao and Wang [2011]).

Θ defined by, for example, (6), it is clear that $\min\Theta = \min(\min(\theta_1^H), \min(\theta_2^H), \min(\theta_3^H))$, for $H = \{A, B, C\}$.

Returning to the problem more generally, it is shown in Appendix B that the value function $V^*(\pi_t)$ solves the following Hamilton-Jacobi-Bellman equation

$$V^*(\pi_t) = \max\left\{\pi_t(\kappa)V^G - I, \left(\frac{1}{2}\sigma_L^2\pi_t^2(1-\pi_t)^2V_t'' - \sigma_L\kappa\pi_t(1-\pi_t)V_t'\right)dt + V_t\right\}, \tag{12}$$

where $V(\pi_t) \in C^2$ is assumed to be an increasing convex function such that $V(\pi_t) > \pi_t V^G - I$. A necessary condition for optimal stopping is that the value function $V^*(\pi_t)$ dominates the payoff function $\pi_t V^G - I$ for all $t \geq 0$ (Peskir and Shiryaev [2006]). Eq. (12) implies that $\pi_t(\kappa)V^G - I - V^*(\pi_t) \leq 0$, thus ensuring this condition is satisfied.

The second argument in the HJB equation (12) implies that the value in the continuation (or planning) region solves the following ODE

$$\frac{1}{2}\sigma_L^2\pi^2(1-\pi)^2V'' - \sigma_L\kappa\pi(1-\pi)V' = 0 \tag{13}$$

and this is subject to the boundary condition $\lim_{\pi\to 0} V(0) = 0$ implying that as the expected belief in a G outcome tends to zero, the value of treating is always zero.

According to the value-matching condition, at the threshold π^* above which it is optimal to treat,

$$\pi^*V^G - I = V(\pi^*), \tag{14}$$

where $V(\pi_t)$ is a general solution to (13) and, by the smooth pasting condition,

$$V^G = V'(\pi^*). \tag{15}$$

Hence, the threshold π^* satisfies the following equation:

$$\frac{V^G}{V'(\pi^*)} = \frac{\pi^*V^G - I}{V(\pi^*)}. \tag{16}$$

The second condition for optimal stopping is that the value function is superharmonic; i.e., that $V^*(\pi_t) \geq \min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}[(\pi^*V^G - I)]$. The superharmonicity condition ensures the uniqueness of the solution. It is effectively the second order condition of the optimisation problem. Eq. (16) is the first order condition and, therefore, $V^*(\pi_t)$ is superharmonic if

$$\frac{-V''(\pi^*)V^G}{V'(\pi^*)^2} - \frac{V(\pi^*)V^G - (\pi^*V^G - I)V'(\pi^*)}{V(\pi^*)^2} < 0 \iff V''(\pi^*) > 0. \tag{17}$$

This is satisfied by the convexity of $V(\pi_t)$.

The analytic solution to the optimal stopping problem (11) is then stated in the following proposition.

Proposition 1. *Let $\pi_\tau := \pi^*$ be such that $\tau := \inf\{t \geq 0 | \pi_t \geq \pi^*\}$ and let $\beta_1 = 2\kappa/\sigma_L + 1 > 0$. If the Knightian uncertainty that the ambiguity averse clinician faces is κ-ignorance, then her*

value function is equal to

$$V^*(\pi_t) = \begin{cases} \pi_t(\kappa)V^G - I & \text{for } \pi_t(\kappa) \geq \pi^* \\ \left(\pi^* V^G - I\right)\left(\frac{\pi_t(\kappa)}{\pi^*}\right)^{\beta_1}\left(\frac{\sum_{k=0}^{\infty}\pi_t^k(\kappa)a_k(\beta_1)}{\sum_{k=0}^{\infty}(\pi^*)^k a_k(\beta_1)}\right) & \text{for } \pi_t(\kappa) < \pi^*, \end{cases} \tag{18}$$

where π^* *solves*

$$\frac{\pi^* V^G}{\pi^* V^G - I} = \beta_1 + \left(\frac{\sum_{k=0}^{\infty} k(\pi^*)^k a_k(\beta_1)}{\sum_{k=0}^{\infty}(\pi^*)^k a_k(\beta_1)}\right), \tag{19}$$

for $a_0(\beta_1) = 1$*;* $a_1(\beta_1) = \frac{\beta_1}{\beta_1+1}\left(\frac{\kappa-\sigma_L(\beta-1)}{\frac{1}{2}\sigma_L\beta_1-\kappa}\right)$*, and* $a_k(\beta_1)$ *satisfies (for* $k > 1$*)*

$$\begin{aligned} &\left(\frac{1}{2}\sigma_L(\beta_1+k-1)-\kappa\right)(\beta_1+k)a_k(\beta_1) \\ =&(\beta_1+k-1)\left(\sigma_L(\beta_1+k-2)-\kappa\right)a_{k-1}(\beta_1) \\ &-\frac{1}{2}\sigma_L(\beta_1+k-2)(\beta_1+k-3)a_{k-2}(\beta_1). \end{aligned} \tag{20}$$

Proof. See Appendix C. ■

3.2 The Impact of Ambiguity on the Optimal Treatment Strategy

Proposition 2. *The optimal treatment threshold is decreasing in the extent of ambiguity.*

Proof. See Appendix D. ■

This effect is depicted in Fig. 1 for the following parameter values: $(V^G, I, \sigma_L) = (30, 10, 0.2)$.

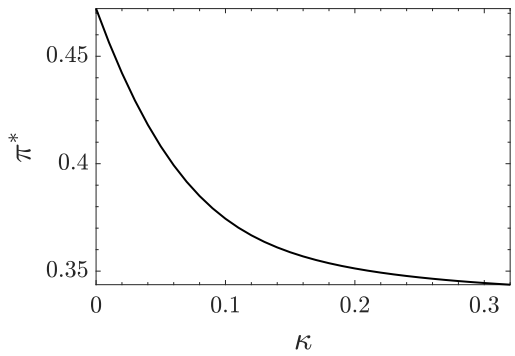

Figure 1: The impact of ambiguity on the optimal treatment threshold.

We see from Fig. 1 that higher levels of ambiguity suggest that treatment is optimal at a lower probabilistic belief in a G outcome on the part of the clinician. Since we assumed the prior probability to be 0.5, it is clear from the figure that for high levels of ambiguity, the threshold belief is below this prior and as, such, immediate treatment is optimal for these parameter values; in other words, early intervention is optimal. To understand this effect, first

consider the expected value from treating. This value is higher when the probabilistic belief in a G outcome is higher which, in turn, is higher when L_t is higher (cf. Eq. (3)). Recall that $\widetilde{W}_t = W_t^\theta - \int_0^t \theta_s ds$ and replacing for $d\widetilde{W}_t$ in (2) such that $\theta_t = \kappa$ under ambiguity aversion, we find that L_t decreases in κ; i.e., the extent of ambiguity. Therefore, the expected value from treating is lower for higher values of ambiguity which suggests that treatment should be postponed.

However, the overall effect depicted in Proposition 2 and in Fig. 1 indicates the opposite; i.e., intervention is optimal for patients when the clinician's belief in a G outcome is low; in other words, when they are in poor clinical states. This implies that the impact of ambiguity is via its impact on the value of waiting meaning that waiting longer is less valuable for higher levels of ambiguity. This is because the model is one of *incomplete information* in which the signal in the form of the patient's clinical state updates the clinican's belief in the expected outcome. If there is a high level of ambiguity over the signal quality, the value of this signal for the purpose of decision making is low so it is disregarded somewhat and the clinician will treat the sick patient promptly before he becomes his clinical state deteriorates further.

In effect, this result implies that if the clinician has a high degree of ambiguity over the reliability of her patient's acquired comorbidities and risk factors in predicting the likely outcome from treatment, she should treat him even if he is in a very poor clinical state with a low probability of a G outcome. Clearly, it is debatable whether this strategy corresponds with best clinical practice so, in the next section, I examine the effect of the optimal treatment strategy owing to ambiguity on the patient's welfare.

4 The Impact on Patient Welfare

A patient's welfare is defined as the *ex ante* expected welfare from being treated at any point in time from initial assessment $t = 0$ until some specified time $t = T$.

In the model, the critical value is measured as a belief, and not as a unit of time. However, to incorporate the time element, we need to determine the first passage time density through the threshold π^*. Following Harrison [1985] (pp. 15), and assuming that the prior belief in a G outcome is $\pi_0 = 1/2$, the probability of the patient being treated within the time interval $t = [0, T]$, is determined to be

$$\mathcal{Q}^\kappa(\sup_{0\leq t\leq T} \pi_t \geq \pi^*) = \Phi(d_1) + \left(\frac{\pi^*}{1-\pi^*}\right)^{2\kappa/\sigma_L} \Phi(d_2), \tag{21}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function and $d_1 = \frac{1}{\sigma_L\sqrt{t}} \ln\left(\frac{1-\pi^*}{\pi^*}\right) - \kappa\sigma_L\sqrt{t}$ and $d_2 = \frac{1}{\sigma_L\sqrt{t}} \ln\left(\frac{1-\pi^*}{\pi^*}\right) + \kappa\sigma_L\sqrt{t}$.

The first passage time cumulative distribution function is given by (21), and I define the associated first passage density function by $f_{\pi^*}(t)$.

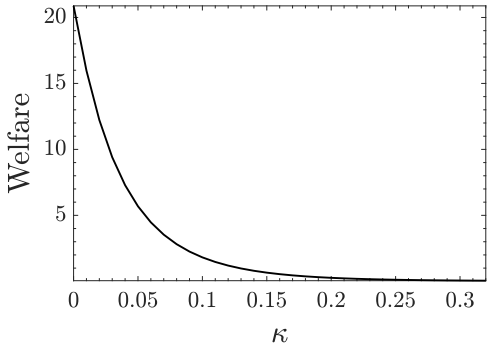

Figure 2: The effect of ambiguity on patient welfare.

Hence, his *ex ante* expected welfare from being treated at some $t \in [0, T]$ is given by

$$W_t(\kappa) = \int_0^T (\pi_t V^G - I) f_{\pi^*}(t) dt. \tag{22}$$

and Fig. 2 depicts patient welfare decreasing in the extent of ambiguity and this effect is robust to a wide choice of parameter values. Therefore, reducing the extent of ambiguity over the clinical status quality in predicting outcomes is key to improving patient welfare.

In the model, as comorbidities and risk factors are acquired (or reversed) by patients, their clinical status changes and, in this way, the clinician learns more about the patient's expected outcome from treatment. However, through this process, she does not learn about the true quality of the signals; in the context of the example provided previously, she does not learn which λ_t^H, or equivalently, which θ_t^H, is the correct one in representing the signal quality.

In the next section, I adapt the model to allow for learning about θ_t. In practical terms, this is achieved as more patients in a particular clinical state are treated at the hospitals carrying out the treatment and, in this way, the success rates of treating in a particular clinical state are updated as more patients in that state are treated. However, it is impractical for the physician and costly for the patient if the clinician will only treat when she has learned enough to resolve her ambiguity entirely. Indeed, it may not be possible to fully resolve ambiguity. Resolving ambiguity entirely is akin to reducing κ to zero over time. However, as pointed out in Epstein and Schneider [2007], ambiguity may increase owing to the acquisition of new information which is surprising relative to past experience and, hence, it is difficult to guarantee that ambiguity will eventually be resolved through learning.

As such, for the sake of patient welfare, the clinician should treat only after reducing her ambiguity about the clinical state as a predictor of outcome sufficiently, but not entirely. In particular, I set up a new optimal stopping problem determining a threshold on the optimal time to stop learning such that, to improve patient welfare, the clinician should only treat the patient in particular clinical state if her belief in a G outcome is sufficiently high *and* her ambiguity about the patient's comorbidities and risk factors is below a certain threshold.

5 Ambiguity and Learning

5.1 Background and Intuition

Until now, the signals that were used in the decision making process were indicative of the *outcome*. The signal was the patient's clinical state and was comprised of his acquired comorbidities and risk factors. Learning via a change in his clinical state corresponded to learning about outcome and the clinical state signal did not provide any information to the clinician about the true probability measure describing $(L_t)_{t\geq 0}$. By contrast, in this section, we want to learn about that true probability measure in order to reduce the ambiguity associated with signal quality for predicting outcomes from treatment in a particular clinical state.

The set of density generators Θ is associated with a set of priors which I denote by $\mathcal{M}_0$. In the analysis thus far, this set of priors remained constant and the elements were not updated over time. In this section, as more patients in a particular clinical state are treated across hospitals, the success probabilities are updated and, in this way, the clinician learns more about the *quality* of the specific clinical state as a predictor of outcome. Learning is modelled by prior by prior Bayesian updating of each prior in $\mathcal{M}_0$ and a corresponding set of posteriors $\mathcal{M}_t$ is obtained.

Once again, I return to the hospital example to provide the intuition. Consider the set Θ described by (6) for clinical states $t = \{1, 2, 3\}$ and hospitals $H = \{A, B, C\}$. Each θ_t^H generates a probability measure $\mathcal{Q}^{\theta_t^H}$ on the quality of the clinical state signal as being representative of the true outcome. However, the true measure is associated with only one of these hospitals H at each t and we want to learn which H at t. This is achieved as follows.

Recall that each generator is associated with a probability, which I denote by λ_t^H, that the patient's clinical status at t is a correct reflection on the true outcome from treatment based on the data of patients treated by hospital H. Intuitively, say patient X's clinical status at time t is denoted by z. However, another patient Y may have been in the clinical state z at some $t' < t$ and treated in that state. Past data is based on all patients treated at H at each $t' < t$ in a clinical state z. As each H treats more patients in that state z over time, more data is acquired and these λ_t^H's are updated and evolve over time.

Patient X is in clinical state z at time t. At $t = 0$, there was a prior probability, denoted here by $\mu_{z_0}^H \in \mathcal{M}_0$, that the true measure is associated with hospital H for clinical state z. The set $\mathcal{M}_0$ consists of the priors for all clinical states z. The priors specific to each state are updated over time, in a Bayesian way, as more patients are treated across the hospitals H in that clinical state. Therefore, at time t, the posterior probability that the true measure is associated with the data from H for patients in clinical state z is given by

$$\mu_{z_t}^H = \mu_{z_0}^H \frac{d(\mathcal{Q}^{\theta_z^H} | \mathcal{F}_t^L)}{d\left[\sum_H \mu_{z_0}^H (\mathcal{Q}^{\theta_z^H} | \mathcal{F}_t^L)\right]}$$

where $\mathcal{Q}^\theta$ is defined by (7) and (8) and $\mathcal{M}_t =: \{\mu_{z_t}^H | \mathcal{F}_t^L\}$.

In the analysis hereafter, the focus is on resolving ambiguity about quality of the particular

state at time t as being reflective of true outcome.

5.2 When to Stop Learning

5.2.1 Signal and Belief Dynamics

In our problem above, the ambiguity averse clinician based her decision about treatment according to the probability measure associated with the worst outcome. However, the true outcome may well be associated with a different probability measure which is associated with treatment in a better clinical state and, therefore, higher patient welfare. As such, I formulate the learning about ambiguity problem as a sequential hypothesis test such that the true probability measure is $\mathcal{Q}^\kappa$ or the reference measure P_π. In essence, I assume that the set of density generators Θ is binary across H for all t so that $\Theta = \{0, \kappa\}$, and 0 and κ are the density generators associated with $H = A$ and $H = B$, respectively.

Since the probabilities λ_t^H that the specific clinical status at t is a correct reflection of the outcomes are easily calculated, I assume these represent the signal process over the true measure. The probability statistical space is described by $(\Omega_\lambda, \mathcal{G}, P_\mu, \mu \in [0,1])$ and I let η be a random variable such that $\eta = 1$ if the true measure is associated with hospital A and $\eta = 0$ if it is associated with B; i.e., with 0 or κ, respectively. Then

$$d\lambda_t = \eta\phi dt + \sigma_\lambda dB_t \tag{23}$$

with $\phi \neq 0$ and σ_λ are constant and $(B_t)_{t\geq 0}$ is a standard Brownian motion under P_μ. Moreover, I let μ_t denote the probability that $\eta = 1$ for all t; i.e., the probability that the true measure is associated with $H = A$ and the density generator $\theta = 0$.

The filtration $\mathcal{G}$ represents the complete information but, as in Section 2.1, we need to analyse the problem in the context of available information $\mathcal{G}^\lambda$ which is the sub-filtration generated by $(\lambda_t)_{t\geq 0}$. As such, we need to apply the filtering procedure once again. To do so, we replace η in (23) with its estimate $\widehat{\eta}_t = \mu_t(1) + (1-\mu_t)(0) = \mu_t$ so that

$$d\lambda_t = \mu_t\phi dt + d\widetilde{B}_t \tag{24}$$

where $d\widetilde{B_t} = d\lambda_t - \mu_t\phi dt$ and describes the dynamics of the innovation process $(\widetilde{B}_t)$ on the *filtered* probability space $(\Omega_\lambda, \mathcal{G}^\lambda, P_\mu, \mu \in [0,1])$.

The likelihood ratio process is defined by (cf. Peskir and Shiryaev [2006] pg. 288)

$$\varphi_t^\lambda = \frac{d(P_A|\mathcal{G}_t^\lambda)}{d(P_B|\mathcal{G}_t^\lambda)} = e^{\phi(\lambda_t - \frac{\phi}{2}t)}$$

and the *a posteriori* probability process $(\mu_t)_{t\geq 0}$ is expressed as

$$\mu_t = \frac{\frac{\mu_0}{1-\mu_0}\varphi_t^\lambda}{1 + \frac{\mu_0}{1-\mu_0}\varphi_t^\lambda}. \tag{25}$$

The following dynamics describing the belief over true probability measure is thus obtained by an application of Ito's lemma:

$$d\mu_t = \phi\mu_t(1-\mu_t)d\widetilde{B}_t. \tag{26}$$

5.2.2 Optimal Stopping Problem

While the clinician is learning about the true measure, she is not treating patients. By not treating a patient in a particular clinical state, there is the chance that the patient's clinical status will deteriorate before treatment is initiated; i.e., he is treated when the clinician's probabilistic belief in a G outcome is lower and, hence, his expected outcome from treatment is also lower. As such, there is a per period cost of learning which I denote by c. Additionally, in the problem above, if κ is not the generator for the true measure, then deciding on his treatment according to κ is also costly for the patient. This is because the true density generator will be some $\theta < \kappa$ and, from Fig. 2, we infer that treating at a lower κ implies a higher patient welfare. If however, it is the true generator, then it should be used in the decision making.

By observing the process $(\lambda_t)_{t\geq 0}$, the clinician decides on what is the true measure by sequentially testing the hypotheses $H_1 : \eta = 1$ and $H_0 : \eta = 0$. By stopping learning, she accepts one of the hypotheses and her objective is to determine the optimal time to stop with minimal cost (see Peskir and Shiryaev [2006] pg. 288). Let d be the terminal decision function such that $d = 1$ if she accepts H_1 and $d = 0$ if she accepts H_0.

Say $\eta = 1$ so that the true measure is associated with $\theta = 0$. Then a Type I error is associated with accepting H_0; i.e., choosing the measure $\kappa > 0$. This implies, from Fig. 1, that the patient may be treated when he should not be; i.e., he is in too bad a clinical state. This wrong choice is costly if the patient's outcome from treatment is B; i.e., there is no improvement in their quality of life. The cost, in that case, is their loss in quality of life from a B outcome I.

If, however, $\eta = 0$ so that the true measure is associated with $\theta = \kappa$, but the clinician makes her decision according to $\theta = 0$ (i.e., accepts H_1), the clinician may not treat the patient when he should be treated. The cost to the patient from this error of not being treated at t is $\pi_t(\kappa)V^G - I + I = \pi_t(\kappa)V^G$, such that $\pi_t(\kappa)$ is constant in state t and, hence, denoted by π hereafter.

Therefore, the optimal stopping problem is then stated as follows

$$F^*(\mu_t) = \inf_{(d,\tau)} E^{P_\mu}[\mu_\tau \pi V^G + (1-\mu_t)I + c(\tau - t)\Big|\mathcal{G}_t^\lambda], \tag{27}$$

where $c(\tau - t)$ denotes the average loss due to learning and not treating for $t < \tau$.

From Shiryaev [1978] Lemma 1 pg. 166, however, $d^* = 1$ if $\mu_t\pi V^G \leq (1-\mu_t)I$ and $d^* = 0$ otherwise for all $t \geq \tau$. Therefore, the optimal stopping problem becomes

$$F^*(\mu_t) = \inf_{\tau} E^{P_\mu}[\min\left[(1-\mu_\tau)I, \mu_\tau\pi V^G\right] + c(\tau - t)\Big|\mathcal{G}_t^\lambda]. \tag{28}$$

Define a μ_d such that $(1-\mu_d)I = \mu_d \pi V^G$; i.e.,

$$\mu_d = \frac{I}{\pi V^G + I}.$$

Now for $\mu_t = 0$ or $\mu_t = 1$, the clinician is certain that the true measure is associated with $\theta = \kappa$ or $\theta = 0$, respectively; in other words, all her ambiguity is resolved. Therefore, $\mu_t = 0$ or $\mu_t = 1$ corresponds with $\kappa = 0$ in the no learning problem. Conversely, if $\mu_t = 1/2$, she is entirely ambiguous over the true measure.

The textbook treatment of this problem (see, for example, Peskir and Shiryaev [2006], Chapter 6) is to solve it as a two-sided optimal stopping problem such that there exists some $\mu^* \geq \mu_d$ and $\mu_* \leq \mu_d$ in which learning is optimal for all $\mu_t \in [\mu_*, \mu^*]$. Say that $\mu_d < 1/2$ and $\mu^* \in [\mu_d, 1/2]$. If $\mu_t \in [\mu^*, 1/2]$, then the clinician will stop learning for higher levels of ambiguity, which is contradictory to her learning objective, and the cost of stopping learning would be lower for higher levels of ambiguity. On the other hand, if $\mu_d > 1/2$, $\mu_* \in [1/2, \mu_d]$ and $\mu_t \in [1/2, \mu_*]$, then learning will also stop for higher levels of ambiguity. Thus, to ensure that the expected cost from stopping is lower for lower levels of ambiguity for all μ_t in the interval, we require that $\mu_d = 1/2$. But $\mu_d = 1/2$ implies that $\pi V^G = I$. Hence, the optimal stopping problem becomes

$$F^*(\mu_t) = \inf_{\tau} E^{P_\mu}[\min[\mu_\tau, (1-\mu_\tau)]\, I + c(\tau - t) \Big| \mathcal{G}_t^\lambda]. \tag{29}$$

Indeed, this makes sense in the context of the problem being considered. Ultimately, the decision is wrong if the patient's quality of life does not improve from treatment and, hence, he foregoes the quality of life he had prior to treatment. If, for example, H_1 is wrongly accepted, but the patient's outcome from treatment is G, accepting the wrong hypothesis is not costly.

The objective of the clinician is to stop learning when ambiguity is at least partially resolved. If $\mu_t \leq \mu_*$, the clinician is sufficiently confident that κ is the true generator and will accept H_0. However, this is a wrong decision if the true generator is $\theta = 0$; i.e., it is a wrong decision with probability μ_t. Therefore, the cost of stopping at $\mu_t \leq 1/2$ is $\mu_t I$. Conversely, the cost of stopping is $(1-\mu_t)I$ for $\mu_t > 1/2$. Clearly, since the cost of stopping increases and decreases in μ_t at the same rate for $\mu_t < 1/2$ and $\mu_t > 1/2$, respectively, and the cost functions intersect at $\mu_t = 1/2$, then $F(\mu_t) = F(1-\mu_t)$ for all μ_t and, hence, $\mu^* = 1 - \mu_*$.

The following is a verification theorem which establishes the free-boundary problem. In particular, it states the conditions which must be satisfied for the existence of a solution to (29).

Theorem 1. *Let $F(\mu_t) \in C^2$ be a concave function such that $0 \leq F(\mu_t) \leq \min(\mu_t, 1-\mu_t)I$. Suppose further that for*

1. *for $\mu_t \in (\mu_*, 1-\mu_*)$, $\frac{1}{2}\phi^2 \mu_t^2 (1-\mu_t)^2 F''(\mu_t) + c = 0$ and*
2. *for $\mu_t \notin (\mu_*, 1-\mu_*)$,*
   (a) *$\frac{1}{2}\phi^2 \mu_t^2 (1-\mu_t)^2 F''(\mu_t) + c \geq 0$ and*

(b) $F(\mu_t) = \min(\mu_t, 1-\mu_t)I$.

Then the free-boundary problem solves the optimal stopping problem; i.e., $F(\mu_t) = F^*(\mu_t)$ *and* $\tau = \inf\{t \geq 0 | \mu_t \notin (\mu_*, 1-\mu_*)\}$.

This is essentially a particular case of Oksendal [2005], Theorem 10.4.1, so its proof is omitted.

The following value matching and smooth pasting conditions must also be satisfied:

$$F(\mu_*) = \min(\mu_*, 1-\mu_*)I = \mu_* I, \tag{30}$$

since $\mu_* < 1/2$. This also implies that

$$F(1-\mu_*) = \min(\mu_*, 1-\mu_*)I = \mu_* I. \tag{31}$$

The corresponding smooth pasting conditions are as follows:

$$F'(\mu_*) = I \tag{32}$$

and

$$F'(1-\mu_*) = -I. \tag{33}$$

Now there is a function

$$g(\mu_t) = \frac{2c}{\phi^2}(1-2\mu_t)\ln\left(\frac{\mu_t}{1-\mu_t}\right) \tag{34}$$

such that

$$F(\mu_t) = g(\mu_t) - g(\mu_*) + \mu_* I + (\mu_t - \mu_*)(I - g'(\mu_*)) \tag{35}$$

satisfies condition 1 of Theorem 1 as well as Eqs. (30) and (32). It is easily verified that $g(\mu_*) = g(1-\mu*)$ implying (35) satisfies (31) also. Finally, applying condition (33) to (35) gives

$$g'(1-\mu_*) + I - g'(\mu_*) = -I.$$

Now, it is also easily verified that $g'(1-\mu_*) = -g'(\mu_*)$. Hence, μ_* satisfies the following equation:

$$g'(\mu_*) = I. \tag{36}$$

Therefore, we obtain the following result.

Proposition 3. *Let* $g(\mu_t)$ *be given by (34) and let* $\mu_* < 1/2$ *satisfy (36). Then*

$$F^*(\mu_t) = \begin{cases} \min(\mu_t, 1-\mu_t)I & \text{for } \mu_t \notin (\mu_*, 1-\mu_*) \\ g(\mu_t) - g(\mu_*) + \mu_* I + (\mu_t - \mu_*)(I - g'(\mu_*)) & \text{for } \mu_t \in (\mu_*, 1-\mu_*), \end{cases} \tag{37}$$

and

$$d^* = \begin{cases} 0 & \text{for } \mu_t \leq \mu_* \\ 1 & \text{for } \mu_t \geq 1-\mu_* \end{cases} \tag{38}$$

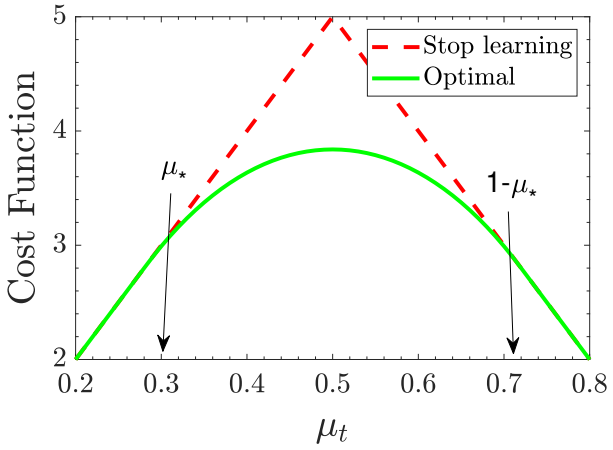

Figure 3: Ambiguity is increasing (decreasing) in μ_t for $\mu_t < 1/2$ ($\mu_t > 1/2$).

solves the optimal stopping problem (29).

Proof. See Appendix E ■

The cost functions and learning thresholds are depicted in Fig. 3 for the following parameter values: $(I, c, \phi) = (10, 0.2, 0.2)$.

Proposition 4. *The higher the per period cost of learning, the earlier the clinician will treat; i.e., for higher levels of ambiguity. However, if her belief in the true measure is very variable over time (i.e., ϕ is high), she will be more likely to wait and refrain from treating.*

Proof. See Appendix F. ■

6 Practical Considerations for Clinicians

Before discussing what the solutions mean for healthcare management, I first explain the relationship between the optimal treatment strategies in the learning and no learning cases.

6.1 No Learning versus Learning

In this section, I explain the relationship between the solutions to the optimal stopping problems defined by Eqs. (11) (no learning) and (29) (learning) as stated in Propositions 1 and 3, respectively.

Proposition 5. *Reducing ambiguity according to the optimal strategy defined in Proposition 3 improves patient welfare relative to the strategy defined in Proposition 1.*

Proof. See Appendix G. ■

To show this effect numerically, and to relate the numerical examples underpinning Figs. 3 and 1 (and, hence the welfare plot Fig. 2), I focus on the $\mu_t \in [0, 1/2]$ region in which the

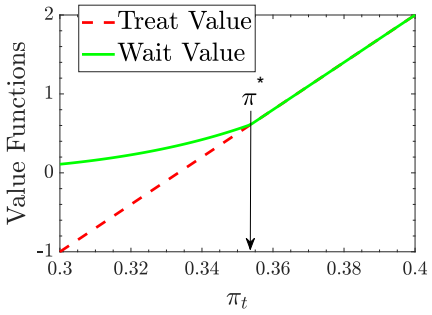

Figure 4: Value functions with no learning about ambiguity for $\kappa = 0.18$.

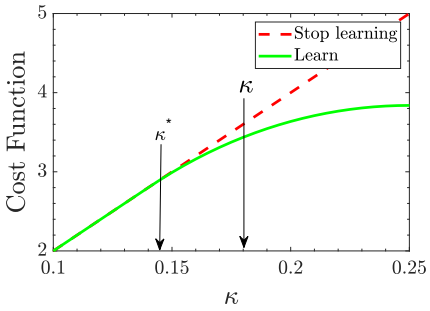

Figure 5: Cost functions for $\mu_t \in (0, 1/2]$

extent of ambiguity is decreasing in μ_t. This region corresponds with the $\kappa \in [0, \overline{\kappa}]$ ambiguity region, where $\overline{\kappa}$ denotes complete ambiguity on the part of the clinician.

Using the parameter values used to plot Figs. 1 and 2, in Fig. 4 I plot the value functions (described by Eq. (18)) which solve the optimal stopping problem in the no learning case for $\kappa = 0.18$. From Fig. 4, in which $\kappa = 0.18$, note that for some values π_t (above approx. $\pi_t \approx 0.355$), it will be optimal to treat. However, in Fig. 5 (which is a snapshot of Fig. 3 with $\overline{\kappa} = 0.25 = 2\mu_t$) we see that at $\kappa = 0.18$, learning is optimal; i.e., it is not optimal to treat.

Now, recall that π_t decreases in κ. Therefore $\pi(\kappa_* = 0.15) > \pi^*(\kappa = 0.18)$. Hence, by accounting for learning and, therefore, stopping only if $\pi_t > \pi(\kappa_*) > \pi^*$ implies treatment is optimal only if the clinician is *more* convinced of a G outcome and this leads to an improvement in patient welfare because he is treated in a better clinical state.

Note that given Prop. 4, a challenge to the result is that if the per period cost of learning is very high, the clinician will treat when she is still highly ambiguous. However, since this cost relates to a decline in a patient's clinical state, and since time is continuous, the periods are very short and, hence, the decline per period can only be very incremental. Thus, in a practical sense, c is very small implying we can accept the result in Prop. 5.

6.2 Practical Implications and Concluding Remarks

The theoretical results derived in previous sections are, for mathematicians working in operations research, not necessarily complex concepts, but they have little value for clinicians unless they can be applied in specific and practical contexts. My objective in this section is to discuss how they can be of value to practicing clinicians.

The timing model in Delaney [2021] is very applicable in a practical way to healthcare decision making. The paper discusses how it can be applied, and it is being trialled for use as a tool in assisting clinicians decide when to perform PVR. It will also intended to be trialled for use in determining the timing for cardiac transplantation. However, in using that model, the signal quality parameter, which is the probability the clinical state is a correct reflection of the true outcome, needs to be estimated from past data. This is quite straightforward but, as discussed, there is ambiguity over this value. Determining the value from different patient

cohorts with the same clinical state produces different quality values. However, the model in Delaney [2021] does not account for this ambiguity.

In this paper, I present a related model to determine the impact of ambiguity on the optimal treatment strategy. From a practical perspective, it is inferior to the Delaney [2021] model because it assumes risk factors and comorbidities are acquired every period. However, this limiting assumption simplifies the model sufficiently so that the impact of ambiguity on patient welfare, as a result of clinicians using the optimal treatment timing strategy defined in Delaney [2021], is determined. Importantly, it shows that a clinician's ambiguity over the clinical status as a predictor of treatment outcomes is detrimental to patient welfare.

Hence, from a practical perspective, clinicians should use the Delaney [2021] model to determine optimal treatment timing, *but should only treat if their ambiguity over signal quality is sufficiently low.* If, for example, that model is saying the patient should be treated, but the clinician is very ambiguous over signal quality, then she should not treat, but instead expend time researching and learning in order to reduce her ambiguity.

The question is then, of course, how to determine a threshold level on ambiguity below which treatment can be performed provided the clinician's threshold on clinical status is between the upper and lower bounds derived in Delaney [2021] which delineate the boundaries between patients that are treatable and those that are, respectively, too well and too unwell for treatment.

To determine this threshold μ_*, the following parameters need to be estimated: ϕ, I and c so that the equation

$$\frac{2c}{\phi^2}(2\mu_* - 1)\ln\left(\frac{\mu_*}{1-\mu_*}\right) = \mu_* I$$

can be solved for μ_*. The parameter I is the quality of life without treatment, and there are a number of measures used by clinicians to determine this (for example, Quality-Adjusted Life Years (QALYs), Health-Related Quality of Life (HRQoL) or Patient Reported Outcomes (PROMS)). Each period that the patient is not being treated reduces their quality of life as they deteriorate further. The learning cost c can be viewed and calculated in this way. Finally, by calculating, from past data, the probability λ_t that a specific clinical state is a correct reflection of the true outcome at various time points, the parameter ϕ can be determined as the drift rate.

While it is relatively straightforward to determine these values, analysing past data is time consuming for clinicians. The overarching point is that ambiguity over the quality of clinical status as a predictor of treatment outcome needs to be low so that the model in Delaney [2021] can be used without compromising patient welfare. In essence, the ambiguity is reduced via the access to data on a large cohort of patients. If there are many hospitals treating such conditions, each will have a relatively small number of patients in their dataset from which the clinicians base their decisions. But if there is more specialisation within the hospitals so that certain treatments are only provided at one or two specialist care centres, this provides the treating clinicians with access to much more data, as well as experience, which will naturally lead to a reduction in ambiguity. How this can be achieved, and for which conditions and illnesses, is beyond the scope of this paper, but my results support the argument for a down-sizing of the general hospital model towards more specialist care providers.

Appendix

A Limit Dynamics for $(L_t)_{t\geq 0}$

Consider a time interval $[0, T]$ and split it into n discrete steps of length Δt. In each time step, a signal arrives which is a correct reflection of the true outcome with probability λ. The log likelihood of signals received up to time t is constructed as follows.

If the true outcome is G (i.e., $\gamma = 1$), then $\Delta L_t = \ln\left(\frac{\lambda}{1-\lambda}\right)$. Hence

$$\begin{aligned} L_T =& L_{T-\Delta t} + \ln\left(\frac{d(P_1^G|\mathcal{F}_{T-\Delta t})}{d(P_1^B|\mathcal{F}_{T-\Delta t})}\right) = L_{T-\Delta t} + \ln\left(\frac{\lambda}{1-\lambda}\right) \\ =& L_{T-2\Delta t} + 2\ln\left(\frac{\lambda}{1-\lambda}\right) \\ & \vdots \\ =& L_0 + n\ln\left(\frac{\lambda}{1-\lambda}\right) \end{aligned}$$

But at $t = 0$, the clinician has no signals so that $L_0 = 0$. Thus, for $\gamma = 1$,

$$L_T = n\ln\left(\frac{\lambda}{1-\lambda}\right)$$

where $n = T/\Delta t$; i.e., the number of time steps in the $[0, T]$ interval.

On the other hand, if $\gamma = 0$,

$$L_T = -n\ln\left(\frac{\lambda}{1-\lambda}\right).$$

Overall, therefore,

$$E[L_T] = \ln\left(\frac{\lambda}{1-\lambda}\right)E[X]$$

where $X = +n$ and $X = -n$ for $\gamma = 1$ and $\gamma = 0$, respectively. If the signal over a ΔT interval is good, then the outcome expected to be successful. This is only a true representation of the actual outcome with probability λ. Therefore $E[X] = n\lambda - n(1-\lambda) = n(2\lambda - 1)$ for good outcomes and $E[X] = -n(2\lambda - 1)$ for bad outcomes.

Now

$$\begin{aligned} V[L_T] =& \left(\ln\left(\frac{\lambda}{1-\lambda}\right)\right)^2 V[X] \\ =& 4n\left(\ln\left(\frac{\lambda}{1-\lambda}\right)\right)^2 \lambda(1-\lambda). \end{aligned}$$

But in the limit, $n =\to \infty$ (or equivalently, $\Delta t \to 0$ since $n = T/(\Delta t)$). Thus, to keep the variance finite, we assume that $\frac{\left(\ln\left(\frac{\lambda}{1-\lambda}\right)\right)^2}{\Delta t} = constant := \sigma_L^2$ (cf. Dalby et al. [2018]). Now since

$\lambda = \frac{e^{\sigma_L\sqrt{\Delta t}}}{1+e^{\sigma_L\sqrt{\Delta t}}}$, $\lambda(1-\lambda) \to 1/4$ as $\Delta t \to 0$ so that $\sigma \to \sigma_L^2 T$ and

$$\lim_{n\to\infty} E[L_t] = \frac{T\sigma_L}{\sqrt{\Delta t}}\left(\frac{e^{\sigma_L\sqrt{\Delta t}}-1}{1+e^{\sigma_L\sqrt{\Delta t}}}\right) = \frac{T\sigma_L}{\sqrt{\Delta t}}\left(\frac{-1+1+\sigma_L\sqrt{\Delta t}+O(dt)}{1+1+\sigma_L\sqrt{\Delta t}+O(dt)}\right) = T\frac{\sigma_L^2}{2} \tag{A.1}$$

for good outcomes and

$$\lim_{n\to\infty} E[L_t] = -T\frac{\sigma_L^2}{2}$$

for bad outcomes.

Therefore, in the limit,

$$dL = \pm\frac{\sigma_L^2}{2}dt + \sigma_L dW = (2\gamma-1)\frac{\sigma_L^2}{2}dt + \sigma_L dW$$

since $\gamma = 1$ if outcome is G and $\gamma = 0$ if it is B.

B Derivation of Eq (12)

$$\begin{aligned} V_t =& \max_{\tau>t}\min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[e^{-\rho(\tau-t)}\left(\pi_\tau(\theta_\tau)V^G - I\right)\Big|\mathcal{F}_t\right] \\ =& \max\left\{\min_{\theta\in\Theta}\left[\pi_t(\theta_t)\right]V^G - I, J_t\right\} \end{aligned} \tag{B.1}$$

where

$$\begin{aligned} J_t =& \max_{\tau\geq t+dt}\min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[e^{-\rho(\tau-t)}\left(\pi_\tau V^G - I\right)\Big|\mathcal{F}_t\right] \\ =& e^{-\rho dt}\max_{\tau\geq t+dt}\min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[E^{\mathcal{Q}^\theta}\left[e^{-\rho(\tau-t-dt)}\left(\pi_\tau V^G - I\right)\Big|\mathcal{F}_{t+dt}\right]\Big|\mathcal{F}_t\right] \\ =& e^{-\rho dt}\max_{\tau\geq t+dt}\min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[\min_{\theta'\in\Theta} E^{\mathcal{Q}^{\theta'}}\left[e^{-\rho(\tau-t-dt)}\left(\pi_\tau V^G - I\right)\Big|\mathcal{F}_{t+dt}\right]\Big|\mathcal{F}_t\right] \\ =& e^{-\rho dt}\min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[\max_{\tau\geq t+dt}\min_{\theta'\in\Theta} E^{\mathcal{Q}^{\theta'}}\left[e^{-\rho(\tau-t-dt)}\left(\pi_\tau V^G - I\right)\Big|\mathcal{F}_{t+dt}\right]\Big|\mathcal{F}_t\right] \\ =& e^{-\rho dt}\min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[V_{t+dt}\big|\mathcal{F}_t\right] \\ =& (1-\rho dt)\left(V_t + \min_{\theta\in\Theta} E^{\mathcal{Q}^\theta}\left[dV_t\big|\mathcal{F}_t\right]\right), \end{aligned} \tag{B.2}$$

where, the second equality is due to the law of iterated expectations and the third due to the rectangularity assumption.

By Ito's lemma and Eq. (9),

$$E^{\mathcal{Q}^\theta}\left[dV_t\big|\mathcal{F}_t\right] = \left(\frac{1}{2}\sigma_L^2\pi_t^2(1-\pi_t)^2V_t'' - \sigma_L\theta_t\pi_t(1-\pi_t)V_t'\right)dt.$$

Since $V(\pi_t)$ is increasing and convex, $V'(\pi_t) > 0$, $E^{\mathcal{Q}^\theta}\left[dV_t\big|\mathcal{F}_t\right]$ is minimal for $\theta_t = \kappa$. This establishes the value of J_t; i.e., the second argument in Eq. (12).

The first argument in Eq. (12) is the value from stopping immediately. This is actually an expected value and, hence, it will be taken to be the minimum expected value with respect to

the generators in $\Theta = [-\kappa, \kappa]$. From Eq. (9), it can be inferred that $\pi_t'(\theta_t) < 0$ and, therefore, the value from stopping at t will be minimal for $\theta_t = \kappa$. Hence, the dependence of π_t on κ.

C Proof of Proposition 1

A general solution to Eq. (13) can be obtained by the method of Frobenius in the following way. Ignoring the dependence of π on t for notational convenience, let

$$V(\pi) = A\pi^{\beta} \sum_{k=0}^{\infty} \pi^k a_k(\beta)$$

where A is constant and β and $\alpha_k(\beta)$ are to be determined.

Then $V'(\pi) = A\pi^{\beta-1} \sum_{k=0}^{\infty}(\beta + k)\pi^k a_k(\beta)$ and $V''(\pi) = A\pi^{\beta-2} \sum_{k=0}^{\infty}(\beta + k - 1)\pi^k a_k(\beta)$.

Substituting for V' and V'' in Eq. (13) and rearranging and re-indexing according to the method gives

$$\begin{aligned}
&\frac{1}{2}\sigma_L^2 \sum_{k=0}^{\infty}(\beta + k)(\beta + k - 1)\pi^{\beta+k} a_k - \sigma_L^2 \sum_{k=1}^{\infty}(\beta + k - 1)(\beta + k - 2)\pi^{\beta+k} a_{k-1} \\
&+ \frac{1}{2}\sigma_L^2 \sum_{k=2}^{\infty}(\beta + k - 2)(\beta + k - 3)\pi^{\beta+k} a_{k-2} - \sigma_L\kappa \sum_{k=0}^{\infty}(\beta + k)\pi^{\beta+k} a_k \\
&+ \sigma_L\kappa \sum_{k=1}^{\infty}(\beta + k - 1)\pi^{\beta+k} a_{k-1} = 0.
\end{aligned} \tag{C.1}$$

For the first term in the series $k = 0$, we have

$$\left(\frac{1}{2}\sigma_L^2\beta(\beta - 1) - \sigma_L\kappa\beta\right) a_0 = 0. \tag{C.2}$$

Ruling out the trivial solution that $a_0 = 0$ implies that the general solution satisfies

$$V(\pi) = A_1\pi^{\beta_1} \sum_{k=0}^{\infty} \pi^k a_k(\beta_1) + A_2,$$

where A_1 and A_2 are constant and $\beta_1 = 2\kappa/\sigma_L + 1$.

However, in order for the boundary condition $V(0) = 0$ to be satisfied, I let $A_2 = 0$.

Continuing with the recursion, for $k = 1$, we have

$$a_1 = \frac{\beta_1}{\beta_1 + 1}\left(\frac{\kappa - \sigma_L(\beta - 1)}{\frac{1}{2}\sigma_L\beta_1 - \kappa}\right) a_0, \tag{C.3}$$

where we can let $a_0 = 1$ hereafter.

For the kth term in the recursion, for $k > 1$, we have

$$\begin{aligned}&\left(\frac{1}{2}\sigma_L(\beta_1+k-1)-\kappa\right)(\beta_1+k)a_k\\ =&(\beta_1+k-1)\left(\sigma_L(\beta_1+k-2)-\kappa\right)a_{k-1}\\ &-\frac{1}{2}\sigma_L(\beta_1+k-2)(\beta_1+k-3)a_{k-2}.\end{aligned} \tag{C.4}$$

The value of waiting is therefore given by

$$V(\pi_t) = A_1\pi_t^{\beta_1}\sum_{k=0}^{\infty}\pi^k a_k(\beta_1),$$

where β_1 and $a_k(\beta_1)$ are defined as above.

According to the value-matching condition, at the threshold π^*, above which it is optimal to treat,

$$A_1 = \left(\sum_{k=0}^{\infty}(\pi^*)^k a_k(\beta_1)\right)^{-1}\left(\frac{\pi^* V^G - I}{(\pi^*)^{\beta_1}}\right) \tag{C.5}$$

and (18) follows from this.

Moreover, Eq. (19) follows from (16) with the appropriate substitutions for $V(\pi_t)$ and $V'(\pi_t)$.

D Proof of Proposition 2

From Eq. (19), letting the summations range from $k = 0$ to $k = 1$, we have that

$$\frac{\pi^* V^G}{\pi^* V^G - I} \approx \beta_1 + \frac{a_1(\beta_1)\pi^*}{1 + a_1(\beta_1)\pi^*}, \tag{D.1}$$

$$a_1(\beta_1) = \frac{\beta_1}{\beta_1+1}(1-\beta_1) < 0$$

and

$$\begin{aligned}\frac{\partial a_1}{\partial \kappa} > 0 \Longleftrightarrow & 2\frac{(\beta_1+1)(1-2\beta_1)-\beta_1(1-\beta_1)}{\sigma_L(\beta_1+1)^2} > 0\\ \Longleftrightarrow & (\beta_1+1)(1-2\beta_1)-\beta_1(1-\beta_1) > 0\\ \Longleftrightarrow & 1-2\beta_1-\beta_1^2 > 0\end{aligned} \tag{D.2}$$

This cannot hold, so $\frac{\partial a_1}{\partial \kappa} < 0$.

Let

$$\mathcal{W} = \frac{\pi^* V^G}{\pi^* V^G - I} - \beta_1 + \frac{a_1\pi^*}{1 + a_1\pi^*} = 0$$

Then

$$\begin{aligned}\frac{\partial \mathcal{W}}{\partial \kappa} =& \frac{\partial \mathcal{W}}{\partial \pi^*}\frac{\partial \pi^*}{\partial \kappa} + \frac{\partial \mathcal{W}}{\partial a_1}\frac{\partial a_1}{\partial \kappa} - \frac{\partial \beta_1}{\partial \kappa} = 0 \\ \Longrightarrow & \frac{\partial \pi^*}{\partial \kappa} = \frac{2/\sigma_L - \frac{\partial \mathcal{W}}{\partial a_1}\frac{\partial a_1}{\partial \kappa}}{\frac{\partial \mathcal{W}}{\partial \pi^*}}\end{aligned} \tag{D.3}$$

Then since $\frac{\partial a_1}{\partial \kappa} < 0$ and

$$\begin{aligned}\frac{\partial \mathcal{W}}{\partial a_1} &= \frac{1}{(1+a_1\pi^*)^2} > 0 \\ \frac{\partial \pi^*}{\partial \kappa} > 0 &\Longleftrightarrow \frac{\partial \mathcal{W}}{\partial \pi^*} > 0.\end{aligned} \tag{D.4}$$

But

$$\frac{\partial \mathcal{W}}{\partial \pi^*} = -\frac{I}{(\pi^* V^G - I)^2} < 0.$$

Therefore, $\frac{\partial \pi^*}{\partial \kappa} < 0$.

E Proof of Proposition 3

Denote the function on the right hand side of Eq. (37) by $G(\cdot)$. By Dynkin's formula,

$$G(\mu_{t'}) = G(\mu_t) + \frac{\phi^2}{2}\int_t^{t'} \mu_s^2(1-\mu_s)^2 G''(\mu_s)ds + \phi\int_t^{t'} \mu_s(1-\mu_s)G'(\mu_s)d\widetilde{B}_s. \tag{E.1}$$

1. Let $\mu_{t'} \notin (\mu_*, 1-\mu_*)$. Then $G(\mu_{t'}) = \min(\mu_{t'}, 1-\mu_{t'})I$. Then, by condition 2(a) in Thm. 1, $\frac{\phi^2}{2}\mu_s^2(1-\mu_s)^2 G''(\mu_s) \geq -c$. Hence

$$F(\mu_{t'}) \geq F(\mu_t) - \int_t^{t'} cds \pm \phi I\int_t^{t'} \mu_s(1-\mu_s)d\widetilde{B}_s. \tag{E.2}$$

   It is easily verified that $\int_t^{t'} \mu_s(1-\mu_s)d\widetilde{B}_s$ is a martingale. Hence, by the optional sampling theorem (see Peskir and Shiryaev [2006], pp. 60), $E^{P_\mu}[\int_t^{t'} \mu_s(1-\mu_s)d\widetilde{B}_s] = 0$ whenever $E^{P_\mu}[t'] < \infty$. Since we take the infimum over stopping times, we can replace t' by τ in (E.2) and take the expectation over E^{P_μ} to give (cf. Eq. (29))

$$F(\mu_\tau) + c(\tau - t) \geq F(\mu_t) = F^*(\mu_t). \tag{E.3}$$

2. Let $\mu_{t'} \in (\mu_*, 1-\mu_*)$. Then by condition 1 of Thm. 1, Eq. (E.1) becomes

$$G(\mu_{t'}) = G(\mu_t) - \int_t^{t'} cds + \phi\int_t^{t'} \mu_s(1-\mu_s)G'(\mu_s)d\widetilde{B}_s. \tag{E.4}$$

   By the same reasoning as in point 1 of this proof,

$$F(\mu_\tau) + c(\tau - t) = F(\mu_t) = F^*(\mu_t). \tag{E.5}$$

Points 1 and 2 ensure that $F^*(\mu_t) \leq F(\mu_\tau) + c(\tau - t)$ (where $F(\mu_t) = \min(\mu_t, 1 - \mu_t)I$) for all $t \geq 0$, which completes the proof.

F Proof of Proposition 4

Since $\mu_* < 1/2$

$$g'(\mu_*) = \frac{2c}{\phi^2}\left[\frac{1 - 2\mu_*}{\mu_*(1 - \mu_*)} - 2\ln\left(\frac{\mu_*}{1 - \mu_*}\right) > 0\right]$$

Thus, $g'(\mu_*)$ is higher, and hence μ_* is higher for higher levels of c and lower levels of ϕ.

G Proof of Proposition 5

1. Let $\pi_t < \pi^*$: For clinical states associated with π_t, the probability of a G outcome is too low to make treatment worthwhile. Thus, the patient will never be treated so the value of learning about the clinical state is not worthwhile; i.e., the learning problem is moot. Furthermore, since the patient will not be treated for π_t in this region, the ex ante expected welfare from treatment patients in the associated clinical state is also irrelevant.

2. Therefore, the learning problem is only relevant and valuable for clinical states such that $\pi_t \geq \pi^*$. In the no learning case, treatment will be initiated with probability one. If the clinician's ambiguity is high, the clinical state of the patient is low, and the likelihood of a B outcome is higher. In other words, the ex ante expected cost from treatment at π_t is $(1 - \pi_t)I$. However, in the learning case, the probability of being treated at π_t is given by $P(\mu_t \leq \mu_*) = \frac{1 - \mu_* - \mu_t}{1 - 2\mu_*}$. This is adapted from Poor and Hadjiliadis [2009] Eq. (4.67), but I provide the derivation specific to this case at the end of the proof. Thus, the ex ante expected cost from treatment at π_t is

   $$(1 - \pi_t)IP(\mu_t \leq \mu_*) \leq (1 - \pi_t)I.$$

   Now, since μ_t denotes the extent of ambiguity, $P(\mu_t \leq \mu_*)$ is low for μ_t high. Therefore, when the extent of ambiguity is high, the ex ante expected cost from treatment is lower in the learning case than in the no learning case implying patient ex ante expected welfare is also higher in the former than in the latter.

G.1 Derivation of $P(\mu_t \leq \mu_*)$

Let $g(\mu_t) = P_\pi(\mu_t \leq \mu_*)$. In the learning region, using the SDE (26) describing the dynamics of $(\mu_t)_{t \geq 0}$, get that the following ordinary differential equation which must be satisfied:

$$\frac{\phi^2}{2}\mu_t^2(1 - \mu_t)^2 g''(\mu_t) = 0$$

such that $g(\mu^*) = 0$. A general solution to this equation is given by

$$g(\mu_t) = A_1\mu_t + A_2,$$

where A_1 and A_2 are constant. Then $A_1 = -\frac{A_2}{\mu^*}$. A further condition is that $g(\mu_*) = 1$. Hence together these conditions give

$$P(\mu_t \leq \mu_*) = \frac{\mu^* - \mu_t}{\mu^* - \mu_*}.$$

As shown and discussed, $\mu^* = 1 - \mu_*$, which yields the expression used in the proof above.

Acknowledgement

I wish to thank Tarik Driouchi and seminar participants at the 4th IMA and OR Society Conference on Mathematics of Operations Research for helpful comments and suggestions.

References

C.J. Berg, B.S. Bauer, A. Hageman, J.A. Aboulhosn, and L.C. Reardon. Mortality risk stratification in Fontan patients who underwent heart transplantation. *American Journal of Cardiology*, 119:1675–1679, 2017.

A.V. Bhide. *The Origin and Evolution of New Business*. Oxford University Press, Oxford, 2000.

Z. Chen and L. Epstein. Ambiguity, risk and asset returns in continuous time. *Econometrica*, 70:1403–1443, 2002.

D.S. Crossland, K. Jansen, and G. Parry et al. Outcome following heart transplant assessment in adults with congenital heart disease. *Heart*, 105:1741–1747, 2019.

P.A.O. Dalby, G.R. Gillerhaugen, V. Hagspiel, T. Leth-Olsen, and J.J.J. Thijssen. Green investment under policy uncertainty and Bayesian learning. *Energy*, 161:1262–1281, 2018.

L. Delaney. A decision rule for determining the optimal transplant listing window for patients with a Fontan physiology. *Medical Decision Making: Policy and Practice*, 6(2):1–13, 2021.

T. Driffield and P.C. Smith. A real options approach to watchful waiting: Theory and illustration. *Medical Decision Making*, Mar/Apr:178–188, 2007.

D. Ellsberg. Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics*, 75: 643–669, 1961.

L. Epstein and S. Ji. Optimal learning under robustness and time-consistency. *Operations Research*, forthcoming, 2020.

L. Epstein and M. Schneider. Learning under ambiguity. *Review of Economic Studies*, 74: 1275–1303, 2007.

W.H. Flemming and R.W. Rishel. *Deterministic and Stochastic Optimal Control*. Springer Verlag, 1975.

T. Geva, B. Mulder, and K. Gauvreau et al. Preoperative predictors of death and sustained ventricular tachycardia after Pulmonary Valve Replacement in patients with repaired Tetrology of Fallot enrolled in the INDICATOR cohort. *Circulation*, 138:2106–2115, 2018.

I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

M.J. Harrison. *Brownian motion and stochastic flow systems*. Wiley, New York, 1985.

C. Heath and A. Tversky. Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4(1):5–28, 1991.

I. Karatzas and S.E. Shreve. *Brownian Motion and Stochastic Calculus, second ed.* Springer, Berlin, 1991.

L.A. Kenny, F. DeRita, M. Nassar, J. Dark, L. Coats, and A. Hasan. Transplantation in the single ventricle population. *Annals of Cardiothoracic Surgery*, 7(1):152–159, 2018.

F. Knight. *Risk, Uncertainty and Profit*. Houghton Mifflin, 1921.

R. Lipster and A. Shiryaev. *Statistics of Random Processes I: General Theory*. Springer, Berlin, 1977.

J. Miao. Ambiguity, risk and portfolio choice under incomplete information. *Annals of Economics and Finance*, 10-2:257–279, 2009.

J.J. Miao and N. Wang. Risk, uncertainty and option exercise. *Journal of Economic Dynamics and Control*, 35(4):442–461, 2011.

K. Nishimura and H. Ozaki. Irreversible investment and Knightian uncertainty. *Journal of Economic Theory*, 136:668–694, 2007.

B. Oksendal. *Stochastic Differential Equations, fifth ed.* Springer, New York, 2005.

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

S. Polyviou, J. O'Sullivan, A. Hasan, and L. Coats. Mortality risk stratification in small patient cohorts: The post-Fontan heart transplant paradigm. *American Journal of Cardiology*, 122: 182–183, 2018.

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

D. Schroder. Investment under ambiguity with the best and worst in mind. *Mathematics and Financial Economics*, 4(2):107–133, 2011.

A. N. Shiryaev. *Optimal Stopping Rules*. Springer, New York- Heidelberg, 1978.

J.T. Tretter and A.N. Redington. Risk factors and biomarkers of poor outcomes. *Circulation*, 138:2116–2118, 2018.

M. Trojanowska and P. Kort. The worst case for real options. *Journal of Optimization Theory and Applications*, 146:709–734, 2010.