

BAYESIAN SEQUENTIAL TESTING OF THE DRIFT OF A BROWNIAN MOTION

ERIK EKSTRÖM¹ AND JUOZAS VAICENAVICIUS

ABSTRACT. We study a classical Bayesian statistics problem of sequentially testing the sign of the drift of an arithmetic Brownian motion with the 0-1 loss function and a constant cost of observation per unit of time for general prior distributions. The statistical problem is reformulated as an optimal stopping problem with the current conditional probability that the drift is non-negative as the underlying process. The volatility of this conditional probability process is shown to be non-increasing in time, which enables us to prove monotonicity and continuity of the optimal stopping boundaries as well as to characterize them completely in the finite-horizon case as the unique continuous solution to a pair of integral equations. In the infinite-horizon case, the boundaries are shown to solve another pair of integral equations and a convergent approximation scheme for the boundaries is provided. Also, we describe the dependence between the prior distribution and the long-term asymptotic behaviour of the boundaries.

1. INTRODUCTION

One of the classical questions in Sequential Analysis concerns the testing of two simple hypotheses about the sign of the drift of an arithmetic Brownian motion. More precisely, suppose that an observed process X_t is an arithmetic Brownian motion

$$X_t = Bt + W_t,$$

where the constant B is unknown and W is a standard driftless Brownian motion. Based on observations of the process X , one wants to test sequentially the hypotheses $H_0 : B < 0$ and $H_1 : B \geq 0$. In the Bayesian formulation of this sequential testing problem, the drift B is a random variable with distribution μ , corresponding to the hypothesis tester's prior belief about the likeliness of the different values B may take. Moreover, it is assumed that B and W are independent. In this article, we consider a classical formulation of the problem in which the accuracy and urgency of a decision is governed by a 0-1 loss function together with a constant cost $c > 0$ of

Date: July 24, 2015.

2000 Mathematics Subject Classification. Primary 62L10, 60G40; Secondary 62C10.

Key words and phrases. Bayesian analysis; sequential hypothesis testing; optimal stopping.

¹ Support from the Swedish Research Council (VR) is gratefully acknowledged.

observation per unit time. The 0-1 loss function means that the tester gains nothing for a right decision but pays a penalty of size 1 for being wrong. The overall goal is to find a decision rule minimising the expected total cost (provided such a decision rule exists). If the decision is required to be made before a fixed predetermined time $T > 0$, the problem is said to have a finite horizon, and if there is no upper bound on the decision time, an infinite horizon.

In the classical literature [5] by Chernoff and [2] by Bather on Bayesian sequential testing procedures for the sign of a drift, the special case of a normal prior distribution is studied. While Bather considers the 0-1 loss function described above as well as a few other alternatives, Chernoff deals with a different penalty function, which equals the magnitude $|B|$ of the error. In these papers, it is argued that the sequential analysis problem reduces to a free-boundary problem for a function of time and the current value of the observation process, but, as in the case of most time-dependent free-boundary problems, the free-boundary problem lacks an explicit solution. Instead, the focus of these and many follow-up articles in the area, including [3], [4], [6], [7], and [15] to mention a few, is on asymptotic approximations for optimal stopping boundaries (for more references, see the survey article [16]). Only recently, [22] has characterised the optimal stopping boundaries for the original Chernoff's problem in terms of an integral equation, which can be solved numerically.

In [21], the sequential testing problem is solved explicitly for a two-point prior distribution by utilising the connection with a time-homogeneous free-boundary problem. Notably, the natural spatial variable in this free-boundary problem is not the value of the observation process, but the conditional probability of the drift taking one of the two possible values. (Since there is a one-to-one correspondence between these two processes at each fixed time, the free-boundary problem could be transformed into one based on the observation process instead, but that formulation would introduce time-dependencies and thus make the explicit solution more difficult to find.)

The fact that the problem can be solved in a very special case of a two-point prior, raises a natural question – can the sequential testing problem be solved for a more general prior distribution? In this article, we investigate the sequential testing problem under a general prior distribution. Since this introduces time-dependencies in the problem, there is generally no hope for explicit solutions. Nevertheless, additional structure is found, which enables us to arrive at a fairly satisfactory answer.

To explain in some further detail, following standard arguments, the statistical problem is shown to admit an equivalent formulation as an optimal stopping problem, which we study to characterise optimal decision rules. The underlying process of the optimal stopping problem is chosen to be the current probability, conditional on observations of X , that the drift is non-negative, i.e.

$$\Pi_t := \mathbb{P}(B \geq 0 | \mathcal{F}_t^X).$$

The pay-off function of the associated optimal stopping problem is then concave in Π , so general results about preservation of concavity for optimal stopping problems may be employed to derive structural properties of the continuation region. Moreover, the volatility of the underlying process Π can be shown to be decreasing in time (except for the two-point distribution discussed above, in which case it is constant). These important facts allow us to show that the optimal stopping boundaries are monotone, so techniques from the theory of free-boundary problems with monotone boundaries can be applied. In particular, the monotonicity of the boundaries enables us to prove the smooth-fit condition and the continuity of the boundaries, as well as to study the corresponding integral equations. In the finite-horizon case, we characterise the optimal boundaries as a unique continuous solution to a pair of integral equations. In the infinite-horizon case, the situation turns out to be more subtle. The boundaries are shown to solve another pair of integral equations, but whether the system admits a unique solution remains unanswered. Instead, we provide a converging approximation scheme for the optimal stopping boundaries, establishing that the optimal boundaries of the finite-horizon problem converge pointwise to the optimal boundaries of the infinite-horizon problem. Also, we determine the long-term asymptotes of the boundaries and describe their dependence on the prior distribution.

From a technical perspective, we tackle a number of issues stemming from the infinite-dimensionality of the parameter space of the underlying process Π , the particular form of the unbounded payoff function, as well as the presence of time-dependent infinite-horizon boundaries. Filtering and analytic techniques are used to understand the behaviour of the conditional probability Π , with a particular focus on the properties that are invariant under any prior distribution. Also, the generality of the prior makes the verification of the smooth-fit condition more involved than in standard situations. Moreover, the specific form of the payoff function with the additive unbounded time term requires some additional effort to prove optimality of the hitting time in the infinite-horizon case. Our approach to approximate the optimal infinite-horizon boundaries could possibly be utilised in other similar situations.

Although we manage to solve this classical sequential hypothesis testing problem for the sign of the drift of a Brownian motion in the general prior case, further generalisations appear to be challenging. The assumption that the observation process is a Brownian motion gives a Markovian structure to the problem in the sense that the posterior distribution then only depends on the current value of the observation process. In contrast, if the observation process was a more general diffusion process, possibly involving time and level-dependent coefficients, then the posterior distribution of the drift would depend on the whole observation path, and thus our one-dimensional Markovian structure would be lost. Another noticeable limitation is the fixed 0-1 loss function, which allows us to use a convenient process Π as the underlying diffusion. For more general loss functions, the natural process

to express the stopping problem in terms of is the underlying observation process itself. However, then the corresponding boundaries are not necessarily monotone, so the corresponding free-boundary problem is less tractable. For Chernoff's loss function and the particular case of a normal prior distribution, [22] finds an appropriate scaling and time-change of X in that the transformation guarantees the monotonicity of the corresponding stopping boundaries. The extension to more general prior distributions and other loss functions remains an important and challenging open problem.

The paper is organised as follows. In Section 2, the sequential testing problem is formulated and reduced to an optimal stopping problem. In Section 3, filtering techniques are applied to find an expression for Π in terms of the observation process X , and its dynamics in terms of the innovation process are determined. We also study the volatility function of Π , and it is shown that this function is non-increasing in time. In Section 4, the optimal stopping problem is studied together with the corresponding free-boundary problem, and it is shown that the optimal stopping boundaries are continuous. In Section 5, integral equations for the boundaries are determined, and uniqueness of solutions is established in the finite-horizon case. The long-term asymptotic behaviour of the infinite-horizon boundaries is presented in Section 6. Finally, Section 7 is devoted to a special case of the normal prior distribution.

Acknowledgement. We are most grateful to Ioannis Karatzas for his suggestion to generalise the setting of an early version of the current paper, and for sharing his unpublished notes on a related problem with us. We also thank two anonymous referees and the associate editor for their careful reading of the paper.

2. PROBLEM FORMULATION AND REDUCTION TO AN OPTIMAL STOPPING PROBLEM

Let $(\Omega, \mathbb{P}, \mathcal{F})$ be a complete probability space supporting a Brownian motion W and a random variable B with distribution μ such that W and B are independent. Define

$$X_t = Bt + W_t.$$

Writing $\mathbb{F}^X = \{\mathcal{F}_t^X\}_{t \geq 0}$ for the filtration generated by the process X and the null sets in \mathcal{F} , our goal is to find a pair (τ, d) consisting of an \mathbb{F}^X -stopping time τ and an \mathcal{F}_τ^X -measurable decision rule $d : \Omega \rightarrow \{0, 1\}$, indicating which of the hypotheses $H_0 : B < 0$ or $H_1 : B \geq 0$ to accept, in order to minimise the Bayes risk

$$R(\tau, d) := \mathbb{E}[\mathbb{1}_{\{d=1, B < 0\}}] + \mathbb{E}[\mathbb{1}_{\{d=0, B \geq 0\}}] + c\mathbb{E}[\tau].$$

Since d is \mathcal{F}_τ^X -measurable, we have

$$R(\tau, d) = \mathbb{E} \left[\mathbb{E} \left[\mathbb{1}_{\{B < 0\}} | \mathcal{F}_\tau^X \right] \mathbb{1}_{\{d=1\}} + \mathbb{E} \left[\mathbb{1}_{\{B \geq 0\}} | \mathcal{F}_\tau^X \right] \mathbb{1}_{\{d=0\}} + c\tau \right], \quad (1)$$

which shows that, at a given stopping time τ , the decision rule

$$d = \begin{cases} 1 & \text{if } \mathbb{P}(B \geq 0 | \mathcal{F}_\tau^X) \geq \mathbb{P}(B < 0 | \mathcal{F}_\tau^X), \\ 0 & \text{otherwise} \end{cases}$$

is optimal. Consequently, writing

$$\Pi_t := \mathbb{P}(B \geq 0 | \mathcal{F}_t^X),$$

the sequential testing problem (1) reduces to an optimal stopping problem

$$V = \inf_{\tau \in \mathcal{T}} \mathbb{E}[g(\Pi_\tau) + c\tau], \quad (2)$$

where $g(\pi) = \pi \wedge (1 - \pi)$ and \mathcal{T} denotes the set of \mathbb{F}^X -stopping times. We also consider the same sequential testing problem but with a finite horizon $T < \infty$. The corresponding optimal stopping problem is then

$$V^T = \inf_{\tau \in \mathcal{T}^T} \mathbb{E}[g(\Pi_\tau) + c\tau],$$

where $\mathcal{T}^T = \{\tau \in \mathcal{T} : \tau \leq T\}$.

Remark By translation, our study readily extends to testing the hypotheses $H_0 : B < \theta$ and $H_1 : B \geq \theta$ for any given $\theta \in \mathbb{R}$. The methods also extend to the case when the two types of possible errors are associated with different costs, i.e. when

$$R(\tau, d) = a\mathbb{E}[\mathbb{1}_{\{d=1, B < 0\}}] + b\mathbb{E}[\mathbb{1}_{\{d=0, B \geq 0\}}] + c\mathbb{E}[\tau]$$

for constants $a > 0$ and $b > 0$ with $a \neq b$. For simplicity of the presentation, however, we assume throughout the article that $\theta = 0$ and $a = b = 1$.

Note that in the cases when $\mu((-\infty, 0)) = 0$ or $\mu([0, \infty)) = 0$, the sequential testing problem becomes trivial as we can make the correct statement about the sign of the drift at time zero. Hence, from now onwards, we always assume that

$$0 < \mu([0, \infty)) < 1. \quad (3)$$

3. CONDITIONAL PROBABILITY OF NON-NEGATIVE DRIFT

In this section, we derive a filtering equation for the distribution of B conditional on the observations of X , which is then applied to prove some elementary results concerning the conditional distribution of the sign of B . We also show that there is an explicit one-to-one correspondence between Π and the observation process X at each fixed time, and we determine the dynamics of X and Π in terms of the innovation process.

3.1. Filtering of the unknown drift.

Proposition 3.1. *Assume that $q : \mathbb{R} \rightarrow \mathbb{R}$ satisfies $\int_{\mathbb{R}} |q(b)|\mu(db) < \infty$. Then*

$$\mathbb{E} [q(B)|\mathcal{F}_t^X] = \frac{\int_{\mathbb{R}} q(b)e^{bX_t - b^2t/2}\mu(db)}{\int_{\mathbb{R}} e^{bX_t - b^2t/2}\mu(db)} \quad (4)$$

for any $t > 0$.

Proof. The proof is based on standard methods in filtering theory, see e.g. [1, Section 3.3], yet we include it for completeness. First define an enlarged filtration $\mathbb{G} = \{\mathcal{G}_t\}_{0 \leq t < \infty}$ as the completion of $\{\sigma(B, W_s : 0 \leq s \leq t)\}_{0 \leq t < \infty}$. Clearly, $\mathcal{F}_t^X \subseteq \mathcal{G}_t$ for any $t \geq 0$, so \mathbb{G} is an enlargement of \mathbb{F}^X . Observing that $Z_t := e^{-BW_t - B^2t/2}$ is a \mathbb{G} -martingale, we define a new probability measure \mathbb{P}_* on the restriction (Ω, \mathcal{G}_T) for some large enough T by

$$\frac{d\mathbb{P}_*}{d\mathbb{P}}|_{\mathcal{G}_T} := Z_T.$$

It can be shown that under \mathbb{P}_* , X_t is a Brownian motion independent of \mathcal{G}_0 and therefore also of B and that the law of B is μ (see [1, Proposition 3.13]). Thus Bayes' rule (cf., for example, [17]) gives that

$$\mathbb{E} [q(B)|\mathcal{F}_t^X] = \frac{\mathbb{E}_* [q(B)/Z_t|\mathcal{F}_t^X]}{\mathbb{E}_* [1/Z_t|\mathcal{F}_t^X]} = \frac{\int_{\mathbb{R}} q(b)e^{bX_t - b^2t/2}\mu(db)}{\int_{\mathbb{R}} e^{bX_t - b^2t/2}\mu(db)}$$

for $t > 0$ and for any function $q : \mathbb{R} \rightarrow \mathbb{R}$ with $\int_{\mathbb{R}} |q(b)|\mu(db) < \infty$. \square

3.2. Conditional probability of non-negative drift. According to Proposition 3.1,

$$\Pi_t = \mathbb{E} [\mathbb{1}_{[0, \infty)}(B)|\mathcal{F}_t^X] = \pi(t, X_t),$$

where the function $\pi(t, x) : (0, \infty) \times \mathbb{R} \rightarrow (0, 1)$ is given by

$$\pi(t, x) := \frac{\int_{[0, \infty)} e^{bx - b^2t/2}\mu(db)}{\int_{\mathbb{R}} e^{bx - b^2t/2}\mu(db)}. \quad (5)$$

Denoting by

$$\mu_{t,x}(db) := \frac{e^{bx - \frac{b^2}{2}t}\mu(db)}{\int_{\mathbb{R}} e^{bx - \frac{b^2}{2}t}\mu(db)} \quad (6)$$

the distribution of B at time t conditional on $X_t = x$, we thus have

$$\pi(t, x) = \int_{[0, \infty)} \mu_{t,x}(db).$$

Proposition 3.2. *Assume that $q : \mathbb{R} \rightarrow \mathbb{R}$ is non-decreasing and satisfies $\int_{\mathbb{R}} |q(b)|\mu(db) < \infty$. Then the function*

$$u(t, x) := \frac{\int_{\mathbb{R}} q(b)e^{bx - b^2t/2}\mu(db)}{\int_{\mathbb{R}} e^{bx - b^2t/2}\mu(db)} = \int_{\mathbb{R}} q(b)\mu_{t,x}(db) \quad (7)$$

is non-decreasing in x for any fixed $t > 0$.

Proof. We will prove the claim by showing that $u(t, \cdot)$ is differentiable with a non-negative derivative on \mathbb{R} . By a standard differentiation lemma (see, for example, [14, Theorem 6.28]), both the numerator and the denominator in (7) are differentiable with respect to x with their derivatives obtained by differentiating under the integral sign. Thus the derivative of u with respect to the second argument x is given by

$$\begin{aligned} \partial_2 u(t, x) &= \int_{\mathbb{R}} q(b) b \mu_{t,x}(db) - \int_{\mathbb{R}} q(b) \mu_{t,x}(db) \int_{\mathbb{R}} b \mu_{t,x}(db) \\ &= \mathbb{E}_{t,x} [q(B)B] - \mathbb{E}_{t,x} [q(B)] \mathbb{E}_{t,x} [B], \end{aligned} \quad (8)$$

where $\mathbb{E}_{t,x}$ is the expectation operator under the probability measure $\mathbb{P}_{t,x}(\cdot) := \mathbb{P}(\cdot | X_t = x)$. Since

$$\begin{aligned} &\mathbb{E}_{t,x} [q(B)B] - \mathbb{E}_{t,x} [q(B)] \mathbb{E}_{t,x} [B] \\ &= \mathbb{E}_{t,x} [(B - \mathbb{E}_{t,x}[B])(q(B) - q(\mathbb{E}_{t,x}[B]))] \geq 0, \end{aligned} \quad (9)$$

this finishes the proof. \square

Corollary 3.3. *Let $a \in \mathbb{R}$ and $t > 0$. Then*

- (1) $\mathbb{P}(B > a | X_t = x)$ is non-decreasing in x ,
- (2) $\mathbb{P}(B < a | X_t = x)$ is non-increasing in x .

Proof. The first claim follows by applying Proposition 3.2 to the function $q(b) = 1_{(a, \infty)}(b)$. The second claim follows from $\mathbb{P}(B < a | X_t = x) = 1 - \mathbb{P}(B \geq a | X_t = x)$ and by applying Proposition 3.2 to the function $q(b) = 1_{[a, \infty)}(b)$. \square

Proposition 3.4. *For any given $t > 0$, the function $\pi(t, \cdot) : \mathbb{R} \rightarrow (0, 1)$ defined in (5) is a strictly increasing continuous bijection.*

Proof. First note that

$$(B - \mathbb{E}_{t,x}[B])(\mathbb{1}_{[0, \infty)}(B) - \mathbb{1}_{[0, \infty)}(\mathbb{E}_{t,x}[B])) \geq 0$$

and that

$$\mathbb{P}_{t,x}((B - \mathbb{E}_{t,x}[B])(\mathbb{1}_{[0, \infty)}(B) - \mathbb{1}_{[0, \infty)}(\mathbb{E}_{t,x}[B])) > 0) > 0$$

since (3) implies $\mu_{t,x}((-\infty, 0)) > 0$ and $\mu_{t,x}([0, \infty)) > 0$. Consequently, the inequality in (9) is strict, so $x \mapsto \pi(t, x)$ is strictly increasing.

Next, note that

$$\pi(t, x) = \frac{1}{1 + A(t, x)},$$

where

$$A(t, x) = \frac{\int_{(-\infty, 0)} e^{bx - \frac{b^2}{2}t} \mu(db)}{\int_{[0, \infty)} e^{bx - \frac{b^2}{2}t} \mu(db)}. \quad (10)$$

By monotone convergence, we find that $A(t, x) \rightarrow 0$ as $x \rightarrow \infty$ and $A(t, x) \rightarrow \infty$ as $x \rightarrow -\infty$. Consequently, $\lim_{x \rightarrow \infty} \pi(t, x) = 1$ and $\lim_{x \rightarrow -\infty} \pi(t, x) = 0$, which finishes the proof. \square

An immediate consequence of Proposition 3.4 is that for any fixed $t > 0$, the spatial inverse $\pi(t, \cdot)^{-1} : (0, 1) \rightarrow \mathbb{R}$ exists. To facilitate intuition, we denote the inverse by $x(t, \cdot)$.

We end this subsection with a result that describes the long-term behaviour of the process Π .

Proposition 3.5. $\Pi_t \rightarrow \Pi_\infty$ a.s. as $t \rightarrow \infty$, where Π_∞ is a Bernoulli random variable with $\mathbb{P}(\Pi_\infty = 0) = \mu((-\infty, 0))$ and $\mathbb{P}(\Pi_\infty = 1) = \mu([0, \infty))$.

Proof. Firstly, since $\Pi_t = \mathbb{E}[\mathbb{1}_{[0, \infty)}(B) | \mathcal{F}_t^X]$ is a bounded martingale, by the martingale convergence theorem, the pointwise limit $\Pi_\infty := \lim_{t \rightarrow \infty} \Pi_t$ is a well-defined random variable closing the martingale Π . By the law of large numbers for Brownian motion and Proposition 6.1 below (the proof of which is independent of the current result), for any $b < 0$ in the support of μ we have

$$\mathbb{P}(\Pi_\infty = 0 | B = b) = 1,$$

so

$$\mathbb{P}(\Pi_\infty = 0) = \int_{(-\infty, 0)} \mathbb{P}(\Pi_\infty = 0 | B = b) \mu(db) = \mu((-\infty, 0)).$$

Hence as Π_∞ can only take values in $[0, 1]$, the fact that $\mathbb{E}[\Pi_\infty] = \mathbb{E}[\Pi_0] = \mu([0, \infty))$ implies $\mathbb{P}(\Pi_\infty = 1) = \mu([0, \infty))$. \square

3.3. SDE for the conditional probability of non-negative drift. Assuming that B has a first moment, the conditional expectation of B exists and is given by

$$\mathbb{E}[B | \mathcal{F}_t^X] = \frac{\int_{\mathbb{R}} b e^{bX_t - b^2 t/2} \mu(db)}{\int_{\mathbb{R}} e^{bX_t - b^2 t/2} \mu(db)}, \quad (11)$$

compare (4). Moreover, the observation process X is represented in terms of the innovation process

$$\hat{W}_t := X_t - \int_0^t \mathbb{E}[B | \mathcal{F}_s^X] ds$$

as

$$dX_t = \mathbb{E}[B | \mathcal{F}_t^X] dt + d\hat{W}_t.$$

Here \hat{W} is a standard \mathbb{F}^X -Brownian motion (see [1, Proposition 2.30 on p. 33]). Moreover, writing $\mathbb{F}^{\hat{W}} = \{\mathcal{F}_t^{\hat{W}}\}_{t \geq 0}$ for the completion of the filtration $\{\sigma(\hat{W}_s : 0 \leq s \leq t)\}_{t \geq 0}$, we have $\mathbb{F}^X = \mathbb{F}^{\hat{W}}$ (see the remark on p. 35 in [1]).

From now onwards, the following integrability condition on μ will be imposed throughout the article.

Assumption. $\int_{\mathbb{R}} e^{\epsilon b^2} \mu(db) < \infty$ for some $\epsilon > 0$. (12)

Note that this assumption is a minor restriction on our hypothesis testing problem since, given any probability distribution μ , the distributions $\mu_{t,x}$ all satisfy (12) for $t > 0$. In other words, no matter what prior distribution μ one starts with, the condition (12) will be satisfied after any infinitesimal time of observation. Also, note that the assumption allows us to extend the definition of $\mu_{t,x}$ in (6) to $t = 0$. Moreover, if we have a prior distribution ξ on \mathbb{R} given by

$$\xi(db) := \frac{e^{\epsilon b^2} \mu(db)}{\int_{\mathbb{R}} e^{\epsilon b^2} \mu(db)},$$

then

$$\mu_{0,x}(db) = \xi_{2\epsilon,x}(db) := \frac{e^{bx-b^2(2\epsilon)/2} \xi(db)}{\int_{\mathbb{R}} e^{bx-b^2(2\epsilon)/2} \xi(db)}. \quad (13)$$

Consequently, the distribution $\mu_{0,x}$ can be identified with a conditional distribution at time 0 given that the prior distribution at time -2ϵ was ξ and the current value of the observation process is x . This gives us a generalisation of the notion of the starting point of the observation process X to allow $X_0 = x \neq 0$, and we may regard time 0 as an interior point of the time interval.

A closer look at the condition (12) and the expression (5) assures that the standard differentiability lemma can be applied to differentiate $\pi(t, x)$ with respect to both variables multiple times inside $(-2\epsilon, \infty) \times \mathbb{R}$. Applying Ito's formula to $\Pi_t = \pi(t, X_t)$, we find that

$$\begin{aligned} d\Pi_t &= \left(\partial_1 \pi(t, X_t) + \mathbb{E}[B | \mathcal{F}_t^X] \partial_2 \pi(t, X_t) + \frac{1}{2} \partial_2^2 \pi(t, X_t) \right) dt + \partial_2 \pi(t, X_t) d\hat{W}_t \\ &= \partial_2 \pi(t, x(t, \Pi_t)) d\hat{W}_t, \end{aligned}$$

where the second equality is verified using the expression (11) and

$$\begin{aligned} \partial_1 \pi(t, x) &= - \int_{[0, \infty)} \frac{b^2}{2} \mu_{t,x}(db) + \int_{[0, \infty)} \mu_{t,x}(db) \int_{\mathbb{R}} \frac{b^2}{2} \mu_{t,x}(db), \\ \partial_2 \pi(t, x) &= \int_{[0, \infty)} b \mu_{t,x}(db) - \int_{\mathbb{R}} b \mu_{t,x}(db) \int_{[0, \infty)} \mu_{t,x}(db), \end{aligned} \quad (14)$$

$$\begin{aligned} \partial_2^2 \pi(t, x) &= \int_{[0, \infty)} b^2 \mu_{t,x}(db) - 2 \int_{\mathbb{R}} b \mu_{t,x}(db) \int_{[0, \infty)} b \mu_{t,x}(db) \\ &\quad - \int_{\mathbb{R}} b^2 \mu_{t,x}(db) \int_{[0, \infty)} \mu_{t,x}(db) + 2 \int_{[0, \infty)} \mu_{t,x}(db) \left(\int_{\mathbb{R}} b \mu_{t,x}(db) \right)^2. \end{aligned}$$

Thus the dynamics of Π_t are specified by a zero drift and the volatility

$$\sigma(t, \Pi_t) = \partial_2 \pi(t, x(t, \Pi_t)), \quad (15)$$

being a positive function of the current time and the current value of Π . Using (8), the volatility function can be expressed as

$$\begin{aligned} \sigma(t, \pi) &= \mathbb{E}_{t,x(t,\pi)}[B \mathbb{1}_{\{B \geq 0\}}] - \mathbb{P}_{t,x(t,\pi)}(B \geq 0) \mathbb{E}_{t,x(t,\pi)}[B] \\ &= (1 - \pi) \mathbb{E}_{t,x(t,\pi)}[B \mathbb{1}_{\{B \geq 0\}}] - \pi \mathbb{E}_{t,x(t,\pi)}[B \mathbb{1}_{\{B < 0\}}]. \end{aligned} \quad (16)$$

Example (The two-point distribution). Assume that $\mathbb{P}(B = a_1) = 1 - p$ and $\mathbb{P}(B = a_2) = p$ for some constants $a_1 < 0 \leq a_2$ and $p \in (0, 1)$. Then

$$\mathbb{P}(B = a_2 | X_t = x) = \pi(t, x) = \frac{pe^{a_2x - a_2^2t/2}}{(1-p)e^{a_1x - a_1^2t/2} + pe^{a_2x - a_2^2t/2}},$$

and

$$\sigma(t, \pi) = (a_2 - a_1)\pi(1 - \pi).$$

This example with a two-point prior distribution is a special case of the Wonham filter.

Example (The normal distribution). Assume that μ is the normal distribution with mean m and variance γ^2 . Then the conditional distribution $\mathbb{P}(\cdot | X_t = x) = \mu_{t,x}$ is also normal but with mean $\frac{m + \gamma^2 x}{1 + t\gamma^2}$ and variance $\frac{\gamma^2}{1 + t\gamma^2}$. Consequently,

$$\pi(t, x) = \Phi\left(\frac{m + \gamma^2 x}{\gamma\sqrt{1 + t\gamma^2}}\right) \quad (17)$$

and

$$\sigma(t, \pi) = \varphi(\Phi^{-1}(\pi)) \frac{\gamma}{\sqrt{1 + t\gamma^2}},$$

where

$$\varphi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz$$

and

$$\Phi(y) = \int_{-\infty}^y \varphi(z) dz$$

are the density and the cumulative distribution of the standard normal random variable, respectively. Note that this instance of a normal prior distribution is a special case of the well-known Kalman-Bucy filter, see for example [17, Chapter 6].

3.4. Volatility of the conditional probability process. In this section we study the volatility function σ . The main result, Corollary 3.8, states that the volatility is non-increasing as a function of time.

Let $\pi \in (0, 1)$ be a fixed number and consider the map $x(\cdot, \pi) : [0, \infty) \rightarrow \mathbb{R}$ sending $t \mapsto x(t, \pi)$. Note that the graph of this function is the trajectory that the process X_t has to follow in order for the conditional probability process Π_t to stay constant at the value π . Thus we call $x(\cdot, \pi)$ *the π -level curve*. Some handy regularity of $x(\cdot, \cdot)$ and $\sigma(\cdot, \cdot)$ is brought to light in the following.

Proposition 3.6. *The functions $x(\cdot, \cdot) : [0, \infty) \times (0, 1) \rightarrow \mathbb{R}$ and $\sigma(\cdot, \cdot) : [0, \infty) \times (0, 1) \rightarrow \mathbb{R}$ are both C^1 .*

Proof. Define $F : (-\epsilon, \infty) \times \mathbb{R} \rightarrow (-\epsilon, \infty) \times (0, 1)$ by $F(t, x) = (t, \pi(t, x))$, where ϵ is as in the assumption (12). The function F is C^1 , which is evident by applying the standard differentiation lemma as in the proof of Proposition 3.2. The Jacobian matrix of F is

$$J_F(t, x) = \begin{pmatrix} 1 & 0 \\ \partial_1 \pi(t, x) & \partial_2 \pi(t, x) \end{pmatrix}.$$

Since F is invertible and $\det(J_F(t, x)) = \partial_2 \pi(t, x) > 0$ for all $t > -\epsilon$ and all $x \in \mathbb{R}$, the inverse function theorem tells us that the inverse of F is also C^1 , with

$$J_{F^{-1}}(F(t, x)) = (J_F(t, x))^{-1}.$$

Consequently, $x(\cdot, \cdot)$ is C^1 on $(-\epsilon, \infty) \times (0, 1)$ with the derivatives $\partial_1 x(t, \pi) = -\partial_1 \pi(t, x(t, \pi)) / \partial_2 \pi(t, x(t, \pi))$ and $\partial_2 x(t, \pi) = 1 / \partial_2 \pi(t, x(t, \pi))$. Finally, since a product of continuous functions is continuous, by the chain rule, $\sigma(\cdot, \cdot)$ is continuously differentiable on $(-\epsilon, \infty) \times (0, 1)$ and so on $[0, \infty) \times (0, 1)$. \square

Next, denoting the initial value Π_0 by $\pi_0 \in (0, 1)$, we show that the tails of the conditional distribution $\mu_{t,x}$ are decreasing along the level curve $x(\cdot, \pi_0)$.

Proposition 3.7.

(1) *If $a \geq 0$, then for any $t > 0$,*

$$\mathbb{P}(B > a | X_t = x(t, \pi_0)) - \mathbb{P}(B > a) \leq 0. \quad (18)$$

Supposing $\mu((a, \infty)) > 0$, the inequality above is strict if and only if $\mu([0, a]) > 0$.

(2) *Likewise, if $a < 0$, then for any $t > 0$,*

$$\mathbb{P}(B < a | X_t = x(t, \pi_0)) - \mathbb{P}(B < a) \leq 0. \quad (19)$$

Supposing $\mu((-\infty, a)) > 0$, the inequality above is strict if and only if $\mu([a, 0)) > 0$.

Proof. We prove only the first of the two claims as the proof of the second one follows the same argument with straightforward modifications.

In the case $\mu((a, \infty)) = 0$, the claim holds trivially with equality in (18). Thus we assume that $\mu((a, \infty)) > 0$ in what follows. Writing $x(t)$ instead of $x(t, \pi_0)$ for brevity, we note that using $\mu_{t,x(t)}([0, \infty)) = \mu([0, \infty))$, the inequality (18) is easily seen to be equivalent to

$$\frac{\int_{(a, \infty)} e^{bx(t) - b^2 \frac{t}{2}} \mu(db)}{\int_{[0, \infty)} e^{bx(t) - b^2 \frac{t}{2}} \mu(db)} \leq \frac{\mu((a, \infty))}{\mu([0, \infty))}. \quad (20)$$

Now, we will split the proof into consideration of two separate cases.

Case 1: $a \geq 2x(t)/t$. Here

$$\begin{aligned} \frac{\int_{[0, a]} e^{bx(t) - b^2 \frac{t}{2}} \mu(db)}{\int_{(a, \infty)} e^{bx(t) - b^2 \frac{t}{2}} \mu(db)} &\geq \frac{e^{ax(t) - a^2 \frac{t}{2}} \mu([0, a])}{e^{ax(t) - a^2 \frac{t}{2}} \mu((a, \infty))} \\ &= \frac{\mu([0, a])}{\mu((a, \infty))}, \end{aligned}$$

which is equivalent to (20). Since $\mu((a, \infty)) > 0$, the inequality above is strict if and only if $\mu([0, a]) > 0$.

Case 2: $x(t) > 0$ and $0 < a < 2x(t)/t$. Using that $\mu_{t,x(t)}([0, \infty)) = \mu([0, \infty))$, we get

$$\begin{aligned} \int_{[0, \infty)} e^{bx(t) - b^2 t/2} \mu(db) &= \mu([0, \infty)) \frac{\int_{(-\infty, 0)} e^{bx(t) - b^2 t/2} \mu(db)}{\mu((-\infty, 0))} \\ &< \mu([0, \infty)), \end{aligned}$$

where the inequality holds since $x(t) > 0$. Hence rewriting

$$\frac{\int_{(a, \infty)} e^{bx(t) - b^2 t/2} \mu(db)}{\mu((a, \infty))} = \frac{\int_{[0, \infty)} e^{bx(t) - b^2 t/2} \mu(db) - \int_{[0, a]} e^{bx(t) - b^2 t/2} \mu(db)}{\mu([0, \infty)) - \mu([0, a])}$$

and keeping in mind that $0 < a < 2x(t)/t$, one clearly sees that

$$\frac{\int_{(a, \infty)} e^{bx(t) - b^2 t/2} \mu(db)}{\mu((a, \infty))} \leq \frac{\int_{[0, \infty)} e^{bx(t) - b^2 t/2} \mu(db)}{\mu([0, \infty))} \quad (21)$$

with the strict inequality if and only if $\mu([0, a]) > 0$. As (21) is equivalent to (20), the proof is complete. \square

Corollary 3.8. *For any $\pi \in (0, 1)$ fixed, the volatility function $\sigma(\cdot, \pi) : [0, \infty) \rightarrow \mathbb{R}$ defined in (15) is non-increasing in time. Moreover, it is strictly decreasing for any initial prior μ except a two-point distribution in which case $t \mapsto \sigma(t, \pi)$ is a constant function.*

Proof. A key to the proof is a realisation that it is sufficient to prove that $\sigma(0, \pi_0) \geq \sigma(s, \pi_0)$ for any $s > 0$; the rest will immediately follow by a ‘moving-frame’ argument. More precisely, by ‘moving-frame’ we mean that for any $\pi \in (0, 1)$, $t \geq 0$, one can think of $\mu_{t,x(t,\pi)}$ as the initial prior distribution at time zero and so immediately obtain that $\sigma(t, \pi) \geq \sigma(t+s, \pi)$ for any $s > 0$.

Using a shorthand $x(t)$ for $x(t, \pi_0)$ as before, recall from (16) that

$$\sigma(t, \pi_0) = (1 - \pi_0)\mathbb{E}_{t,x(t)}[B\mathbb{1}_{\{B \geq 0\}}] - \pi_0\mathbb{E}_{t,x(t)}[B\mathbb{1}_{\{B < 0\}}].$$

Consequently,

$$\begin{aligned} \sigma(0, \pi_0) - \sigma(t, \pi_0) &= (1 - \pi_0) (\mathbb{E}[B\mathbb{1}_{\{B \geq 0\}}] - \mathbb{E}_{t,x(t)}[B\mathbb{1}_{\{B \geq 0\}}]) \\ &\quad + \pi_0 (\mathbb{E}_{t,x(t)}[B\mathbb{1}_{\{B < 0\}}] - \mathbb{E}[B\mathbb{1}_{\{B < 0\}}]) \\ &\geq 0 \end{aligned}$$

by Proposition 3.7. Moreover, by the same proposition, the inequality reduces to an equality if and only if μ is a two-point distribution. \square

Remark It seems difficult to find an easy intuitive argument for the monotonicity of the volatility function. As an example, consider a symmetric prior distribution, and a strictly positive time-point t at which the observation process satisfies $X_t = 0$. Then the conditional distribution $\mu_{t,0}$ is also symmetric, so $\Pi_t = \Pi_0 = 1/2$. One certainly expects that $\mu_{t,0}$ is obtained from the prior distribution μ by pushing mass towards zero (this is also verified in Proposition 3.7 above). One could expect that the Π -process of a distribution with a lot of mass close to zero is sensitive to small changes in the observation process since the mass easily may ‘spill over’ to the other side of zero, and thus such a distribution gives rise to a comparatively large volatility. On the other hand, a concentrated distribution makes it difficult to distinguish possible drifts from each other, and changes in the observation process would to a higher degree be attributed to the Brownian fluctuations. This implies a slow learning process, which indicates a small volatility. Corollary 3.8 shows that the latter effect outweighs the former one.

4. ANALYSIS OF THE OPTIMAL STOPPING PROBLEM

In this section, we study the perpetual optimal stopping problem (36) and its finite-horizon counterpart under the integrability condition (12). Most of the time the emphasis is on the perpetual case, though the corresponding results also hold for the finite horizon case by the same arguments. If the analogy is straightforward, we only comment on it, otherwise, more details are provided.

4.1. The value function with arbitrary starting points. Recall that

$$d\Pi_t = \sigma(t, \Pi_t) d\hat{W}_t,$$

where

$$\sigma(t, \pi) = \partial_2 \pi(t, x(t, \pi)) > 0$$

for all $(t, \pi) \in [0, \infty) \times (0, 1)$ (beware that $\pi(\cdot, \cdot)$ is a function, while π is a real number). We embed the optimal stopping problem (36), in which the

starting point of the process Π is given by $\Pi_0 := \mathbb{P}(B \geq 0)$, into the optimal stopping problem

$$v(t, \pi) = \inf_{\tau \in \mathcal{T}} \mathbb{E} \left[g(\Pi_{t+\tau}^{t, \pi}) + c\tau \right], \quad (t, \pi) \in [0, \infty) \times (0, 1), \quad (22)$$

for the process $\Pi^{t, \pi} = \Pi$ given by

$$\begin{cases} d\Pi_{t+s} = \sigma(t+s, \Pi_{t+s}) d\hat{W}_{t+s}, & (s > 0) \\ \Pi_t = \pi, \end{cases} \quad (23)$$

where \mathcal{T} denotes the set of stopping times with respect to the completed filtration of $\{\Pi_{t+s}^{t, \pi}\}_{s \geq 0}$. The SDE (23) possesses a unique solution since $\sigma(\cdot, \cdot)$ is locally Lipschitz by Proposition 3.6. Furthermore, the embedding has a consistent interpretation also at time $t = 0$, which is given by (13) and the remark following it. Note that choosing $\tau = 0$ gives $v(t, \pi) \leq g(\pi)$.

Proposition 4.1. *The value function $v(t, \pi)$ is concave in π for any fixed $t \geq 0$.*

Proof. This follows by a standard approximation argument using optimal stopping problems where stopping is only allowed at a discrete set of time-points, compare [8].

To outline this, denote by $\mathcal{T}_{t,n}$, where $t \leq n$, the set of stopping times in \mathcal{T} taking values in $\{k2^{-n}, k = 0, 1, \dots, n2^n\} \cap [0, n-t]$, $n = 1, 2, \dots$, and let

$$v_n(t, \pi) = \inf_{\tau \in \mathcal{T}_{t,n}} \mathbb{E} \left[g(\Pi_{t+\tau}^{t, \pi}) + c\tau \right].$$

Then $v_n(n, \pi) = g(\pi)$ is concave in π . By preservation of concavity for martingale diffusions, see [11] (the results of [11] extend to the current setting with both an upper and a lower bound on the state space), $\pi \mapsto v_n(t, \pi)$ is concave also for $t \in (n - 2^{-n}, n)$. Next, at time $t = n - 2^{-n}$ the value is given by dynamic programming as

$$v_n(t, \pi) = \min \left\{ g(\pi), \mathbb{E} \left[v_n(n, \Pi_n^{t, \pi}) + c2^{-n} \right] \right\},$$

which is concave (being the minimum of two concave functions). Proceeding recursively shows that v_n is concave in π at all times $t \in [0, n]$. Since v_n converges pointwise to v as $n \rightarrow \infty$, this implies that also v is concave in π . \square

Proposition 4.2. *The value function $v(t, \pi)$ is non-decreasing in t for every fixed $\pi \in (0, 1)$.*

Proof. This can be proven using approximation by Bermudan options as in the proof of Proposition 4.1 above. Indeed, for a fixed time $t \geq 0$ one may approximate $v(t, \pi)$ by the optimal value in the case when stopping times are restricted to take values in the set $\{k2^{-n} : n \in \mathbb{N}, k \in \{0, 1, \dots, n2^n\}\}$. Since the expected value of a concave function of a martingale diffusion is non-increasing in the volatility, see [11], the approximation is non-decreasing in t by Corollary 3.8. Letting $n \rightarrow \infty$ finishes the proof. \square

Remark It is straightforward to check that the monotonicity of v in the time-variable also holds in cases when the rate c of the observation cost is increasing in time (instead of a constant as in our set-up). As the non-decreasing value function implies the monotonicity of the optimal stopping boundaries (see Proposition 4.5), the same monotonicity of the optimal stopping boundaries would be present also in the case of the observation rate c being increasing in time. Accordingly, we expect the subsequent results to extend to that case as well.

Proposition 4.3. *The value function v is continuous on $[0, \infty) \times [0, 1]$.*

Proof. By concavity of v in the second variable together with the bounds $0 \leq v \leq g$, we have that v is Lipschitz continuous in π for any fixed t , with Lipschitz coefficient 1. Thus it suffices to check that v is continuous in time. To do this, let $t_2 > t_1 \geq 0$ and note that

$$v(t_1, \pi) \geq \mathbb{E}[v(t_2, \Pi_{t_2}^{t_1, \pi})] \geq v(t_2, \pi) - \mathbb{E}[|\Pi_{t_2}^{t_1, \pi} - \pi|],$$

where the first inequality holds since $\mathbb{E}[v(t_2, \Pi_{t_2}^{t_1, \pi})]$ represents the value of a sequential testing problem, started at t_1 , with the running cost of observation not started until time t_2 , the second inequality holds by the concavity of v in the second variable and the bounds $0 \leq v \leq g$. Thus

$$0 \leq v(t_2, \pi) - v(t_1, \pi) \leq \mathbb{E}[|\Pi_{t_2}^{t_1, \pi} - \pi|].$$

Since the expected value of a convex function of a martingale diffusion is non-decreasing in the volatility (again by [11]) and $\sigma(0, \cdot) \geq \sigma(\cdot, \cdot)$ on $[0, \infty) \times (0, 1)$, we deduce that $\mathbb{E}[|\Pi_{t_2}^{t_1, \pi} - \pi|] \leq \mathbb{E}[|\Pi_{t_2-t_1}^{0, \pi} - \pi|] \rightarrow 0$ as $t_2 - t_1 \searrow 0$. This finishes the proof. \square

Lemma 4.4. *We have $v(t, 1/2) < g(1/2)$ for all times $t \geq 0$.*

Proof. Let $t \geq 0$, and define $A_\epsilon := [t, t + \epsilon] \times [1/2 - (c + 1)\epsilon, 1/2 + (c + 1)\epsilon]$ for ϵ small enough so that $A_\epsilon \subseteq [t, \infty) \times (0, 1)$. Let

$$\tau_\epsilon := \inf\{s \geq 0 : (t + s, \Pi_{t+s}^{t, 1/2}) \notin A_\epsilon\}$$

be the first exit time from A_ϵ . By Proposition 3.6 and Corollary 3.8, $\sigma(\cdot, \cdot)$ is continuous and strictly positive on $[0, \infty) \times (0, 1)$. Thus $\sigma_\epsilon := \inf_{(s, \pi) \in A_\epsilon} \sigma(s, \pi)$ is strictly positive and non-increasing as a function of ϵ , so σ_ϵ is bounded away from 0 as $\epsilon \rightarrow 0$. Now,

$$\begin{aligned} g(1/2) - v(t, 1/2) &\geq 1/2 - \mathbb{E}\left[g(\Pi_{t+\tau_\epsilon}^{t, 1/2}) + c\tau_\epsilon\right] \\ &\geq 1/2 - (1/2 - (c + 1)\epsilon + c\epsilon)\mathbb{P}(\tau_\epsilon < \epsilon) - (1/2 + c\epsilon)\mathbb{P}(\tau_\epsilon = \epsilon) \\ &= \epsilon - (c + 1)\epsilon\mathbb{P}(\tau_\epsilon = \epsilon). \end{aligned} \tag{24}$$

Here

$$\begin{aligned}
\mathbb{P}(\tau_\epsilon = \epsilon) &= \mathbb{P}\left(\sup_{0 \leq s \leq \epsilon} \left| \int_0^s \sigma(t+u, \Pi_{t+u}^{t, 1/2}) d\hat{W}_{t+u} \right| \leq (c+1)\epsilon\right) \\
&\leq \mathbb{P}\left(\sup_{0 \leq s \leq \epsilon} \left| \sigma_\epsilon \hat{W}_s \right| \leq (c+1)\epsilon\right) \\
&\leq \mathbb{P}\left(\sup_{0 \leq s \leq \epsilon} \hat{W}_s \leq (c+1)\epsilon/\sigma_\epsilon\right) \rightarrow 0
\end{aligned} \tag{25}$$

as $\epsilon \rightarrow 0$, where the first inequality follows from [11, Lemma 10]. Consequently, (24) and (25) yield that $g(1/2) - v(t, 1/2) > 0$, which finishes the proof of the claim. \square

4.2. The structure of an optimal strategy. Recalling that $0 \leq v(t, \pi) \leq g(\pi)$, we denote by

$$\mathcal{C} := \{(t, \pi) \in [0, \infty) \times (0, 1) : v(t, \pi) < g(\pi)\}$$

the continuation region, and by

$$\mathcal{D} := \{(t, \pi) \in [0, \infty) \times (0, 1) : v(t, \pi) = g(\pi)\}$$

the stopping region. Since v is continuous, \mathcal{C} is open and \mathcal{D} is closed. Resorting to intuition from optimal stopping theory, we expect that the stopping time

$$\tau^* := \inf\{s \geq 0 : (t+s, \Pi_{t+s}^{t, \pi}) \in \mathcal{D}\} \tag{26}$$

is an optimal stopping time in (22). (Note that standard optimal stopping theory does not apply since the pay-off process is not uniformly integrable.) The optimality of τ^* is verified below, see Theorem 4.6.

Proposition 4.5. *There exist two functions $b_1 : [0, \infty) \rightarrow [0, 1/2)$ and $b_2 : [0, \infty) \rightarrow (1/2, 1]$ such that*

$$\mathcal{C} = \{(t, \pi) : b_1(t) < \pi < b_2(t)\}.$$

The function b_1 is non-decreasing and right-continuous with left limits. Similarly, b_2 is non-increasing and right-continuous with left limits.

Proof. The existence of b_1 and b_2 follows from the concavity of v and Lemma 4.4. The monotonicity properties are immediate consequences of Proposition 4.2. Moreover, by the continuity of v , the function b_1 is upper semi-continuous and b_2 is lower semi-continuous. Hence, they are right-continuous with left limits. \square

Let us also consider the same optimal stopping problem with a finite horizon $T > 0$. It is written as

$$v^T(t, \pi) = \inf_{\tau \in \mathcal{T}_{T-t}} \mathbb{E} \left[g(\Pi_{t+\tau}^{t, \pi}) + c\tau \right], \tag{27}$$

where \mathcal{T}_{T-t} denotes the set of stopping times less or equal to $T-t$ with respect to the completed filtration of $\{\Pi_{t+s}^{t, \pi}\}_{s \geq 0}$. Note that all results for the perpetual problem (22) described above in this section also hold for the

finite horizon problem (35), with the obvious modifications regarding the time horizon, by the same proofs. Moreover, the pay-off process in (35) is continuous and bounded, so standard optimal stopping theory (see, for example, [20, Corollary 2.9 on p. 46]) yields that

$$\tau^T := \inf\{s \geq 0 : \Pi_{t+s}^{t,\pi} \notin (b_1^T(t+s), b_2^T(t+s))\}$$

is an optimal stopping time in (35), where b_1^T and b_2^T are the corresponding boundaries enclosing the finite-horizon continuation region

$$\mathcal{C}^T := \{(t, \pi) \in [0, T) \times (0, 1) : v^T(t, \pi) < g(\pi)\}.$$

The infinite-horizon problem can be approximated by finite-horizon problems in the following sense.

Theorem 4.6. *The functions $v^T \searrow v$, $b_1^T \searrow b_1$, and $b_2^T \nearrow b_2$ pointwise as $T \nearrow \infty$. The stopping times $\tau^T \nearrow \tau^*$ a.s. as $T \nearrow \infty$, where τ^* is defined in (26). Moreover, τ^* is optimal in (22).*

Proof. Since $v^T \geq v$, we have that $b_1 \leq b_1^T < b_2^T \leq b_2$ and $\tau^T \leq \tau^*$. By bounded and monotone convergence, $v^T(t, \pi) \searrow v(t, \pi)$ pointwise as $T \rightarrow \infty$, so $b_1^T \searrow b_1$ and $b_2^T \nearrow b_2$ pointwise as $T \rightarrow \infty$. Consequently, by the monotonicity of b_i and b_i^T , it follows that $\tau^T \nearrow \tau^*$ a.s. Thus

$$v^T(t, \pi) = \mathbb{E} \left[g(\Pi_{t+\tau^T}^{t,\pi}) + c\tau^T \right] \rightarrow \mathbb{E} \left[g(\Pi_{t+\tau^*}^{t,\pi}) + c\tau^* \right]$$

by bounded and monotone convergence. By uniqueness of limits,

$$v(t, \pi) = \mathbb{E} \left[g(\Pi_{t+\tau^*}^{t,\pi}) + c\tau^* \right],$$

so τ^* is optimal. \square

4.3. Optimal stopping boundaries and the free-boundary problem.

Proposition 4.7. *The boundaries b_1 and b_2 satisfy $0 < b_1(t) < 1/2 < b_2(t) < 1$ for all times $t \geq 0$.*

Proof. The two middle inequalities are granted by Lemma 4.4. To see that $b_1 > 0$ on $[0, \infty)$, without loss of generality, it is sufficient to show that $b_1 > 0$ on $(0, \infty)$; this is due to the possibility provided by (13) to start the process Π slightly earlier. Let us assume, to reach a contradiction, that $b_1(t) = 0$ for some $t > 0$. Then, by monotonicity, $b_1 \equiv 0$ on $[0, t]$. By the martingale inequality,

$$\mathbb{P}(\tau^* \leq t) \leq \mathbb{P} \left(\sup_{0 \leq s \leq t} \Pi_s^{0,\pi} \geq 1/2 \right) \leq 2\pi.$$

Consequently,

$$\mathbb{E}[\tau^*] \geq t(1 - 2\pi),$$

and so $0 = \lim_{\pi \searrow 0} g(\pi) \geq \lim_{\pi \searrow 0} v(0, \pi) \geq c\mathbb{E}[\tau^*] \geq ct > 0$, which is a clear contradiction. Therefore $b_1 > 0$ at all times. The proof that $b_2 < 1$ is analogous. \square

Proposition 4.8. *The triplet (v, b_1, b_2) satisfies the free boundary problem*

$$\begin{cases} \partial_1 v(t, \pi) + \frac{\sigma(t, \pi)^2}{2} \partial_2^2 v(t, \pi) + c = 0 & b_1(t) < \pi < b_2(t) \\ v(t, \pi) = \pi & \pi \leq b_1(t) \\ v(t, \pi) = 1 - \pi & \pi \geq b_2(t). \end{cases} \quad (28)$$

Moreover, the smooth-fit condition holds in the sense that the function $\pi \mapsto v(t, \pi)$ is C^1 for all $t \geq 0$.

Proof. The proof that the differential equation in (28) holds is based on the strong Markov property and the continuity of v . However, the procedure is standard and we therefore omit the argument, referring to the proof of [13, Theorem 7.7] for the details instead. The value of v for $\pi \notin (b_1(t), b_2(t))$ follows from concavity and the definition of b_1 and b_2 .

For the smooth-fit condition, note that the value function $\pi \mapsto v(t, \pi)$ is continuous on $(0, 1)$ and C^1 for $\pi \in (b_1(t), b_2(t))$ as well as for $\pi \in (0, b_1(t)) \cup (b_2(t), 1)$. Thus it remains to check the C^1 property at $b_1(t)$ and $b_2(t)$. To prove the C^1 property at $b_1(t)$ (the C^1 property at $b_2(t)$ being completely analogous), note that since v is concave in π , it suffices to show that

$$\liminf_{\epsilon \downarrow 0} \frac{v(t, b(t) + \epsilon) - v(t, b(t))}{\epsilon} \geq 1. \quad (29)$$

Without loss of generality, we do this for $t = 0$, letting $\pi = b_1(0)$.

Let $\epsilon \in (0, 1/2 - \pi)$ and denote by τ^ϵ the first hitting time of the stopping region for $\Pi^{0, \pi + \epsilon}$. Then

$$\begin{aligned} v(0, \pi + \epsilon) - v(0, \pi) &\geq \mathbb{E} \left[g(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon}) - g(\Pi_{\tau^\epsilon}^{0, \pi}) \right] \\ &\geq \epsilon - 2\mathbb{E} \left[(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} - \Pi_{\tau^\epsilon}^{0, \pi}) \mathbb{1}_{\{\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} > 1/2\}} \right]. \end{aligned}$$

Thus, to prove (29) it suffices (by the Cauchy-Schwartz inequality) to show that

$$\mathbb{E} \left[(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} - \Pi_{\tau^\epsilon}^{0, \pi})^2 \right] \mathbb{P} \left(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} > 1/2 \right) = o(\epsilon^2) \quad (30)$$

as $\epsilon \rightarrow 0$. To do this, first assume that σ is Lipschitz continuous in π on any compact time interval, and define

$$h(t) := \mathbb{E} \left[(\Pi_{t \wedge \tau^\epsilon}^{0, \pi + \epsilon} - \Pi_{t \wedge \tau^\epsilon}^{0, \pi})^2 \right].$$

Fixing $T > 0$, for $t \in [0, T]$ we have

$$\begin{aligned} h(t) &= \mathbb{E} \left[\left(\epsilon + \int_0^{t \wedge \tau^\epsilon} \sigma(s, \Pi_s^{0, \pi + \epsilon}) - \sigma(s, \Pi_s^{0, \pi}) d\hat{W}_s \right)^2 \right] \\ &\leq \epsilon^2 + \int_0^t \mathbb{E} \left[D(T)^2 (\Pi_s^{0, \pi + \epsilon} - \Pi_s^{0, \pi})^2 \mathbb{1}_{\{s \leq \tau^\epsilon\}} \right] ds \\ &\leq \epsilon^2 + D(T)^2 \int_0^t h(s) ds, \end{aligned}$$

where $D(T)$ is a Lipschitz constant for σ on $[0, T] \times (0, 1)$. Consequently, Gronwall's inequality yields

$$h(T) \leq \epsilon^2 e^{D(T)^2 T}. \quad (31)$$

Next, denote by $f(y) := \pi - \frac{\pi}{1-\pi}(y - \pi)$ the affine function satisfying $f(\pi) = \pi$ and $f(1) = 0$, and note that $f \leq g$ on $[\pi, 1]$. Therefore,

$$\begin{aligned} c\mathbb{E}[\tau^\epsilon] &= v(0, \pi + \epsilon) - \mathbb{E}[g(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon})] \\ &\leq g(\pi + \epsilon) - \mathbb{E}[f(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon})] \\ &= \pi + \epsilon - \left(\pi - \frac{\pi}{1-\pi}\mathbb{E}[\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} - \pi]\right) \\ &= \epsilon/(1-\pi), \end{aligned}$$

where the inequality follows from the monotonicity of b_1 and the last equality by optional sampling. Thus, writing $D = 1/(1-\pi)$, we have

$$\mathbb{P}(\tau_\epsilon > T) \leq D\epsilon/(cT). \quad (32)$$

Moreover, writing

$$\tau_{\pi, 1/2} := \inf\{s \geq 0 : \Pi_s^{0, \pi + \epsilon} \notin (\pi, 1/2)\},$$

we have

$$\mathbb{P}\left(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} > 1/2\right) \leq \mathbb{P}\left(\Pi_{\tau_{\pi, 1/2}}^{0, \pi + \epsilon} = 1/2\right) \leq \frac{\epsilon}{\frac{1}{2} - \pi} = C\epsilon \quad (33)$$

for $C = 2/(1-2\pi)$, where we used the martingality of Π to obtain the second inequality. Putting together (31), (32) and (33) yields

$$\begin{aligned} &\mathbb{E}\left[(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} - \Pi_{\tau^\epsilon}^{0, \pi})^2\right] \mathbb{P}\left(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} > 1/2\right) \\ &\leq \left(\mathbb{E}\left[(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} - \Pi_{\tau^\epsilon}^{0, \pi})^2 \mathbb{1}_{\{\tau_\epsilon \leq T\}}\right] + \mathbb{P}(\tau_\epsilon > T)\right) \mathbb{P}\left(\Pi_{\tau^\epsilon}^{0, \pi + \epsilon} > 1/2\right) \\ &\leq \epsilon^2 C(\epsilon e^{D(T)^2 T} + D/(cT)). \end{aligned}$$

Given $\delta > 0$, it is possible to choose T large enough so that $CD/(cT) \leq \delta/2$, and then to choose $\epsilon > 0$ small enough so that $C\epsilon e^{D^2(T)T} \leq \delta/2$. This proves (30) and thus finishes the proof of the smooth-fit property if σ is Lipschitz in π , locally uniformly in t .

For a general σ , due to the C^1 regularity of σ on $[0, \infty) \times (0, 1)$, one can find another volatility function $\hat{\sigma}$ that is Lipschitz continuous in π on any given compact interval in time, and that satisfies $0 \leq \hat{\sigma} \leq \sigma$ everywhere and $\hat{\sigma} = \sigma$ on $[0, \infty) \times [b_1(0), b_2(0)]$. By monotonicity in the volatility, the corresponding value function \hat{v} satisfies $\hat{v} \geq v$. On the other hand, since $\hat{\sigma} = \sigma$ on $[0, \infty) \times [b_1(0), b_2(0)]$ and since τ^* is optimal for the volatility σ , we also have $\hat{v} \leq v$, so $\hat{v} = v$. By the above argument, \hat{v} is C^1 , which finishes the proof. \square

Theorem 4.9. *The boundaries b_1 and b_2 are both continuous.*

Proof. Let us prove continuity of b_1 (the proof for b_2 is analogous). We know that b_1 is right-continuous, so it suffices to assume for a contradiction that b_1 is not continuous at some time $t_0 > 0$. By monotonicity, $b_1(t_0) > b_1(t_0-)$. In the continuation region, $\partial_1 v \geq 0$, so (28) yields

$$\frac{\sigma^2}{2} \partial_2^2 v \leq -c.$$

Since σ is locally bounded away from zero, this means that on each compact set we can find some constant $d > 0$ such that $\partial_2^2 v \leq -d$. By Proposition 4.8, the map $\pi \mapsto v(t, \pi)$ is C^1 on $[b_1(t), b_2(t)]$ for any $t \geq 0$, so for $t < t_0$ and $b_1(t) < \pi < b_1(t_0)$, we have

$$\begin{aligned} v(t, \pi) - g(\pi) &= \int_{b_1(t)}^{\pi} \int_{b_1(t)}^w \partial_2^2 (v - g)(t, u) \, du \, dw \\ &\leq -d(\pi - b_1(t))^2 / 2. \end{aligned}$$

Choosing $\pi = \frac{b_1(t_0-) + b_1(t_0)}{2}$ and letting $t \rightarrow t_0$ gives

$$v(t_0, \frac{b_1(t_0-) + b_1(t_0)}{2}) - g(\frac{b_1(t_0-) + b_1(t_0)}{2}) \leq -d(b_1(t_0) - b_1(t_0-))^2 / 2 < 0.$$

This contradicts the assumption that $(t_0, \frac{b_1(t_0-) + b_1(t_0)}{2})$ belongs to the stopping region, so b_1 has to be continuous. \square

Remark Even though, in this section, all the results are formulated for the perpetual problem (22), it is straightforward to check that the corresponding results for the finite-horizon problem (35) also hold. In that case, the boundaries $b_1^T : [0, 1] \rightarrow (0, 1)$ and $b_2^T : [0, 1] \rightarrow (0, 1)$ are continuous and monotone, with $0 < b_1^T < 1/2 < b_2^T < 1$ on $[0, T)$ and $b_1^T(T) = b_2^T(T) = 1/2$. Also, the assertions of Proposition 4.8 hold for (v^T, b_1^T, b_2^T) on the time interval $[0, T)$ in place of (v, b_1, b_2) .

5. INTEGRAL EQUATIONS FOR THE BOUNDARIES

It is well-known that optimal stopping boundaries, under some conditions, can be characterized by certain integral equations, compare [10] and [19]. In this section, we study the integral equations for the optimal stopping boundaries arising in our sequential testing problem. For the problem (35) with finite horizon, a pair of integral equations is shown to completely characterise the optimal stopping boundaries within the class of continuous solutions. The situation in the perpetual case is more delicate, and uniqueness of solutions remains an open question.

5.1. A pair of integral equations for the finite-horizon boundaries.

Theorem 5.1. *Assume that $T < \infty$. Then the pair (b_1^T, b_2^T) is the unique continuous solution of*

$$\begin{cases} c_1(t) = \mathbb{E} \left[g(\Pi_T^{t, c_1(t)}) \right] + c \int_0^{T-t} \mathbb{P}(c_1(t+u) < \Pi_{t+u}^{t, c_1(t)} < c_2(t+u)) \, du \\ 1 - c_2(t) = \mathbb{E} \left[g(\Pi_T^{t, c_2(t)}) \right] + c \int_0^{T-t} \mathbb{P}(c_1(t+u) < \Pi_{t+u}^{t, c_2(t)} < c_2(t+u)) \, du \end{cases} \quad (34)$$

such that $0 < c_1(t) \leq 1/2 \leq c_2(t) < 1$ for all $t \in [0, T]$.

Proof. For $(t, \pi) \in [0, T] \times (0, 1)$ and $\epsilon > 0$ small enough, let

$$\tau_\epsilon := \inf\{s \geq 0 : \Pi_{t+s}^{t, \pi} \notin (\epsilon, 1 - \epsilon)\} \wedge (T - t - \epsilon).$$

Applying Ito's formula (more precisely, an extension of Ito's formula, see [18, Theorem 3.1 and Remark 3.2], which can be applied thanks to the monotonicity of b_1^T and b_2^T) to the process $v^T(t + s \wedge \tau_\epsilon, \Pi_{t+s \wedge \tau_\epsilon}^{t, \pi})$ and then taking expectations yields

$$\mathbb{E} \left[v(t + \tau_\epsilon, \Pi_{t+\tau_\epsilon}^{t, \pi}) \right] = v^T(t, \pi) - c \mathbb{E} \left[\int_0^{\tau_\epsilon} \mathbb{1}_{(b_1^T(t+u), b_2^T(t+u))}(\Pi_{t+u}^{t, \pi}) du \right].$$

Since $\tau_\epsilon \rightarrow T - t$ as $\epsilon \rightarrow 0$, it follows from Proposition 4.3 and bounded convergence that

$$\mathbb{E} \left[g(\Pi_T^{t, \pi}) \right] = v^T(t, \pi) - c \int_0^{T-t} \mathbb{P}(b_1^T(t+u) < \Pi_{t+u}^{t, \pi} < b_2^T(t+u)) du. \quad (35)$$

Plugging in $\pi = b_1^T(t)$ and $\pi = b_2^T(t)$ shows that (b_1^T, b_2^T) solves (34).

For uniqueness, assume that (c_1, c_2) is another continuous solution to (34) with $0 < c_1(t) \leq 1/2 \leq c_2(t) < 1$, and define

$$V(t, \pi) := \mathbb{E} \left[g(\Pi_T^{t, \pi}) \right] + c \int_0^{T-t} \mathbb{P}(c_1(t+u) < \Pi_{t+u}^{t, \pi} < c_2(t+u)) du. \quad (36)$$

Then $V(t, c_1(t)) = c_1(t)$ and $V(t, c_2(t)) = 1 - c_2(t)$ by (34), and $V(T, \pi) = g(\pi)$. Moreover, by the Markov property, the process

$$M_s := V(t + s, \Pi_{t+s}^{t, \pi}) + c \int_0^s \mathbb{1}_{(c_1(t+u), c_2(t+u))}(\Pi_{t+u}^{t, \pi}) du$$

is a martingale for any (t, π) . In particular, the process

$$\tilde{M}_s := v^T(t + s, \Pi_{t+s}^{t, \pi}) + c \int_0^s \mathbb{1}_{(b_1^T(t+u), b_2^T(t+u))}(\Pi_{t+u}^{t, \pi}) du$$

is also a martingale.

Claim 1: $V(t, \pi) = g(\pi)$ for $\pi \notin (c_1(t), c_2(t))$.

Assume that $\pi \leq c_1(t)$ (the case $\pi \geq c_2(t)$ is similar), and let

$$\gamma_c := \inf\{s \geq 0 : \Pi_{t+s}^{t, \pi} \geq c_1(t + s)\} \wedge (T - t).$$

Then

$$V(t, \pi) = \mathbb{E} \left[V(t + \gamma_c, \Pi_{t+\gamma_c}^{t, \pi}) \right] = \mathbb{E} \left[\Pi_{t+\gamma_c}^{t, \pi} \right] = \pi = g(\pi),$$

whith the first equality being justified by optional sampling and the martingale property of M , the second by (34), and the third by optional sampling and the martingale property of Π .

Claim 2: $V \geq v^T$.

Take (t, π) such that $c_1(t) < \pi < c_2(t)$, and let

$$\tau_c := \inf\{s \geq 0 : \Pi_{t+s}^{t, \pi} \notin (c_1(t + s), c_2(t + s))\} \wedge (T - t).$$

Then

$$\begin{aligned} V(t, \pi) &= \mathbb{E} \left[V(t + \tau_c, \Pi_{t+\tau_c}^{t, \pi}) \right] + c \mathbb{E} \left[\int_0^{\tau_c} \mathbb{1}_{(c_1(t+u), c_2(t+u))}(\Pi_{t+u}^{t, \pi}) \, du \right] \\ &= \mathbb{E} \left[g(\Pi_{t+\tau_c}^{t, \pi}) \right] + c \mathbb{E} [\tau_c] \geq v^T(t, \pi). \end{aligned}$$

From this and Claim 1, Claim 2 follows.

Claim 3: $b_1^T \leq c_1$ and $c_2 \leq b_2^T$.

Assume that $b_1^T(t) > c_1(t)$ for some t . Choose $\pi = c_1(t)$, and let

$$\gamma_b := \inf\{s \geq 0 : \Pi_{t+s}^{t, \pi} \geq b_1^T(t+s)\} \wedge (T-t).$$

Then, by right-continuity of b_1^T and c_1 ,

$$\mathbb{E} \left[\int_0^{\gamma_b} \mathbb{1}_{(c_1(t+u), c_2(t+u))}(\Pi_{t+u}^{t, \pi}) \, du \right] > 0. \quad (37)$$

On the other hand, by optional sampling and martingality of M and \tilde{M} we have

$$\begin{aligned} 0 &= V(t, \pi) - v^T(t, \pi) \\ &= \mathbb{E} \left[V(t + \gamma_b, \Pi_{t+\gamma_b}^{t, \pi}) - v^T(t + \gamma_b, \Pi_{t+\gamma_b}^{t, \pi}) \right] \\ &\quad + c \mathbb{E} \left[\int_0^{\gamma_b} \mathbb{1}_{(c_1(t+u), c_2(t+u))}(\Pi_{t+u}^{t, \pi}) \, du \right]. \end{aligned}$$

Since $V \geq v$ by Claim 2, this contradicts (37) and thus $b_1^T \leq c_1$. The claim $c_2 \leq b_2^T$ is proved similarly.

Now, it follows from (35), (36) and Claim 3 that $V = v^T$. Therefore, since $v^T < g$ for $\pi \in (b_1^T(t), b_2^T(t))$ it follows from Claims 1 and 3 that $b_1^T = c_1$ and $b_2^T = c_2$, which finishes the proof. \square

Remark A closer inspection of the proof above shows that uniqueness holds in the larger class of pairs (c_1, c_2) such that c_1 is right-continuous with no negative jumps and c_2 is right-continuous with no positive jumps.

5.2. A pair of integral equations for the infinite-horizon boundaries.

Theorem 5.2. *The pair (b_1, b_2) is a solution of*

$$\begin{cases} b_1(t) = c \int_0^\infty \mathbb{P}(b_1(t+u) < \Pi_{t+u}^{t, b_1(t)} < b_2(t+u)) \, du \\ 1 - b_2(t) = c \int_0^\infty \mathbb{P}(b_1(t+u) < \Pi_{t+u}^{t, b_2(t)} < b_2(t+u)) \, du. \end{cases} \quad (38)$$

Proof. For fixed $T > 0$, an application of Itô's formula as in the preceding proof gives

$$\mathbb{E} \left[v(T, \Pi_T^{t, \pi}) \right] = v(t, \pi) - c \int_0^{T-t} \mathbb{P}(b_1(t+u) < \Pi_{t+u}^{t, \pi} < b_2(t+u)) \, du.$$

Since v is bounded and $\Pi_T^{t,\pi}$ converges to either 0 or 1 as $T \rightarrow \infty$ by Proposition 3.5, we find that

$$v(t, \pi) = c \int_0^\infty \mathbb{P}(b_1(t+u) < \Pi_{t+u}^{t,\pi} < b_2(t+u)) du.$$

Plugging in $\pi = b_1(t)$ and $\pi = b_2(t)$ shows that (b_1, b_2) solves (38). \square

Remark The main technical difficulty when trying to apply the uniqueness proof of Theorem 5.1 to the perpetual problem lies in the lack of a straightforward extension of the optional sampling theorem to unbounded, possibly infinite stopping times.

5.3. The case of a symmetric volatility function. Now assume that the volatility function is symmetric about $\pi = 1/2$, i.e. $\sigma(t, \pi) = \sigma(t, 1-\pi)$. This is the case, for example, if the prior distribution μ is symmetric about zero in the sense that $\mu([0, a]) = \mu((-a, 0))$ for all $a > 0$. Then, by symmetry, $b_1^T = 1 - b_2^T$, and we set $b^T := b_1^T$. The following result is a straightforward consequence of Theorem 5.1.

Theorem 5.3. *Assume that σ is symmetric about $\pi = 1/2$. Then the boundary b^T is the unique continuous solution of*

$$c(t) = \mathbb{E} \left[g(\Pi_T^{t,c(t)}) \right] + c \int_0^{T-t} \mathbb{P}(c(t+u) < \Pi_{t+u}^{t,c(t)} < 1 - c(t+u)) du \quad (39)$$

such that $0 < c(t) \leq 1/2$ for all $t \in [0, T]$.

Remark Although not necessarily symmetric, all normal prior distributions as well as all two-point priors give rise to symmetric volatilities, compare Section 3.

6. LONG-TERM ASYMPTOTICS OF THE VOLATILITY AND THE BOUNDARIES

Since the boundaries b_1 and b_2 are monotone, the limits $b_i(\infty) := \lim_{t \rightarrow \infty} b_i(t)$, $i = 1, 2$, exist with $b_1(\infty) \leq 1/2$ and $b_2(\infty) \geq 1/2$. In this section we determine these limits. To do that, we first derive a few limiting properties of level curves as well as study the limit $\sigma(\infty, \pi) := \lim_{t \rightarrow \infty} \sigma(t, \pi)$ of the volatility.

Let us define

$$r = \inf\{s \geq 0 : \mu([s, s + \epsilon]) > 0 \text{ for all } \epsilon > 0\} \quad (40)$$

and

$$l = \sup\{s < 0 : \mu((s - \epsilon, s]) > 0 \text{ for all } \epsilon > 0\}. \quad (41)$$

We write $m = (l + r)/2$ for the midpoint between l and r .

The following proposition will serve as a useful device for understanding long-term volatility.

Proposition 6.1.

1. *If $\alpha > m$, then $\pi(t, \alpha t) \rightarrow 1$ as $t \rightarrow \infty$.*

2. If $\alpha < m$, then $\pi(t, \alpha t) \rightarrow 0$ as $t \rightarrow \infty$.

Proof. Given $\alpha \in \mathbb{R}$, define

$$h(t) := \frac{\int_{(-\infty, 0)} \exp\left(-\frac{(b-\alpha)^2 t}{2}\right) \mu(db)}{\int_{[0, \infty)} \exp\left(-\frac{(b-\alpha)^2 t}{2}\right) \mu(db)},$$

so that $\pi(t, \alpha t) = 1/(1 + h(t))$. We will prove the claims in two different cases separately.

(i) First case: $l < r$.

1. First note that, in view of Proposition 3.4, it suffices to treat the case $\alpha \in (m, r)$. For such α , fix $\gamma > r$ such that $\gamma - \alpha < \alpha - l$. Then

$$h(t) \leq \frac{\exp\left(-\frac{(\alpha-l)^2 t}{2}\right) \int_{(-\infty, 0)} \mu(db)}{\exp\left(-\frac{(\gamma-\alpha)^2 t}{2}\right) \int_{[0, \gamma]} \mu(db)} \rightarrow 0$$

as $t \rightarrow \infty$. Hence $\pi(t, \alpha t) \rightarrow 1$ as $t \rightarrow \infty$.

2. For the second result, suppose that $\alpha < m$, and note that it suffices to treat the case $\alpha \in (l, m)$. Let $\gamma < l$ be such that $\alpha - \gamma < r - \alpha$. Then

$$h(t) \geq \frac{\exp\left(-\frac{(\alpha-\gamma)^2 t}{2}\right) \int_{(\gamma, 0)} \mu(db)}{\exp\left(-\frac{(r-\alpha)^2 t}{2}\right) \int_{[0, \infty)} \mu(db)} \rightarrow \infty$$

as $t \rightarrow \infty$. Hence $\pi(t, \alpha t) \rightarrow 0$ as $t \rightarrow \infty$.

(ii) Second case: $l = r = 0$.

1. Assume that $\alpha > 0$, and let $\epsilon > 0$. Then

$$\begin{aligned} h(t) &\leq \frac{\int_{(-\infty, -\epsilon)} \exp\left(-\frac{(\alpha+\epsilon)^2 t}{2}\right) \mu(db) + \int_{[-\epsilon, 0)} \exp\left(-\frac{\alpha^2 t}{2}\right) \mu(db)}{\int_{[0, \alpha]} \exp\left(-\frac{\alpha^2 t}{2}\right) \mu(db)} \\ &\rightarrow \frac{\mu([- \epsilon, 0))}{\mu([0, \alpha])} \end{aligned}$$

as $t \rightarrow \infty$. Thus, since $\epsilon > 0$ is arbitrary and $\mu([- \epsilon, 0)) \rightarrow 0$ as $\epsilon \rightarrow 0$, we conclude that $h(t) \rightarrow 0$ as $t \rightarrow \infty$. Consequently, $\pi(t, \alpha t) \rightarrow 1$.

2. Next, assume that $\alpha < 0$. Choosing $\gamma \in (\alpha, 0)$ with $\mu((\alpha, \gamma)) > 0$, we find that

$$h(t) \geq \frac{\exp\left(-\frac{(\alpha-\gamma)^2 t}{2}\right) \int_{(\alpha, \gamma)} \mu(db)}{\exp\left(-\frac{\alpha^2 t}{2}\right) \int_{[0, \infty)} \mu(db)} \rightarrow \infty$$

as $t \rightarrow \infty$. Consequently, $\pi(t, \alpha t) \rightarrow 0$ as $t \rightarrow \infty$, which finishes the proof. \square

Remark Notice that Proposition 6.1 implies that for any fixed value π , the corresponding level curve $x(\cdot, \pi)$ satisfies $\lim_{t \rightarrow \infty} (ct - x(t, \pi)) = \infty$ if $\alpha > m$, and $\lim_{t \rightarrow \infty} (\alpha t - x(t, \pi)) = -\infty$ if $\alpha < m$.

6.1. Long-term behaviour of the volatility. Now, we are in a position to determine the limit $\sigma(\infty, \pi) := \lim_{t \rightarrow \infty} \sigma(t, \pi)$ of the volatility as time increases.

Proposition 6.2. *The long-term limit of volatility satisfies $\sigma(\infty, \pi) = (r - l)\pi(1 - \pi)$.*

Remark Note that if $l = r = 0$, then the volatility converges to zero as time tends to infinity. Also, note that if $l < r$, then the volatility tends to the volatility from the case of a two-point prior distribution.

Proof of Proposition 6.2. We first claim that

$$\mathbb{E}_{t,x(t,\pi)} [B \mathbb{1}_{[0,\infty)}(B)] \rightarrow \pi r$$

as $t \rightarrow \infty$. To see this, suppose that $a > r$ and take $\gamma \in (r, a)$ such that $\gamma - r < a - \gamma$. By Corollary 3.3 and Proposition 6.1, for all large enough t ,

$$\begin{aligned} \mathbb{E}_{t,x(t,\pi)} [B \mathbb{1}_{(a,\infty)}(B)] &\leq \mathbb{E}_{t,\gamma t} [B \mathbb{1}_{(a,\infty)}(B)] \\ &= \frac{\int_{(a,\infty)} b e^{-(b-\gamma)^2 \frac{t}{2}} \mu(db)}{\int_{\mathbb{R}} e^{-(b-\gamma)^2 \frac{t}{2}} \mu(db)} \\ &\leq \frac{\exp\left(-\frac{(a-\gamma)^2 t}{2}\right) \int_{(a,\infty)} b \mu(db)}{\exp\left(-\frac{(\gamma-r)^2 t}{2}\right) \mu([r, \gamma])}, \end{aligned}$$

which tends to 0 as $t \rightarrow \infty$. Now, the fact that $\mathbb{P}_{t,x(t,\pi)}(B \in [0, r]) = 0$ for all $t \geq 0$ finishes the claim.

Next, straightforward modifications of the arguments above show that

$$\mathbb{E}_{t,x(t,\pi)} [B \mathbb{1}_{(-\infty, 0)}(B)] \rightarrow (1 - \pi)l$$

as $t \rightarrow \infty$. Since

$$\sigma(t, \pi) = (1 - \pi)\mathbb{E}_{t,x(t,\pi)} [B \mathbb{1}_{[0,\infty)}(B)] - \pi\mathbb{E}_{t,x(t,\pi)} [B \mathbb{1}_{(-\infty, 0)}(B)],$$

this finishes the proof. \square

Remark Similar arguments as in the proof above show that $\mu_{t,x(t,\pi)} \Rightarrow (1 - \pi)\delta_l + \pi\delta_r$ as $t \rightarrow \infty$. Thus, along a level curve $x(\cdot, \pi)$ the conditional distribution of B converges weakly to the two-point distribution with mass π at r and mass $1 - \pi$ at l .

6.2. Long-term behaviour of the boundaries.

Theorem 6.3.

- If $l = r = 0$, then $b_1(\infty) = b_2(\infty) = 1/2$.
- If $l < r$, then $b_1(\infty) = b_1^{r-l}$ and $b_2(\infty) = b_2^{r-l}$, where $b_1^{r-l} < 1/2 < b_2^{r-l}$ are the optimal boundaries for a two-point prior distribution with mass at points separated by 0 and at a distance $r - l$ from each other.

Proof. Since the volatility $\sigma(\cdot, \cdot)$ is non-increasing in time, Proposition 6.2 and Dini's theorem yield that $\sigma(t, \cdot)$ converges to $\sigma(\infty, \pi) = (r-l)\pi(1-\pi)$ uniformly on the compact interval $[b_1(0), b_2(0)]$ as $t \rightarrow \infty$. Therefore, given $\epsilon > 0$ we can find t_0 large enough so that $\sigma(t_0, \pi) \leq (\epsilon + r - l)\pi(1 - \pi)$ for $\pi \in [b_1(0), b_2(0)]$. Define

$$\hat{\sigma}(t, \pi) := \sigma(t, \pi)\mathbb{1}_{[b_1(0), b_2(0)]}(\pi),$$

and denote by \hat{v} the corresponding value function. Since the optimal stopping problem (22) is monotone in the volatility (compare e.g. [11, Lemma 10]), we have that $\hat{v} \geq v$. On the other hand, since $\hat{\sigma} = \sigma$ on the continuation region $\{(t, \pi) : b_1(t) < \pi < b_2(t)\}$, we also have $\hat{v} \leq v$, so $\hat{v} = v$. Moreover, by monotonicity in the volatility,

$$v^{\epsilon+r-l} \leq \hat{v} = v \leq v^{r-l}$$

on $[t_0, \infty) \times (0, 1)$, where v^a denotes the value function corresponding to a volatility function $a\pi(1-\pi)$. Since the value function v is squeezed in between the value functions $v^{\epsilon+r-l}$ and v^{r-l} from time t_0 , the optimal stopping boundaries b_1 and b_2 are squeezed in between the corresponding optimal stopping boundaries for $v^{\epsilon+r-l}$ and v^{r-l} . By inspection of the explicit formulas in the two-point distribution case, see [20, Theorem 21.1], the gaps $b_1^{r-l} - b_1^{\epsilon+r-l}$ and $b_2^{\epsilon+r-l} - b_2^{r-l}$ between the boundaries vanish as $\epsilon \rightarrow 0$, which finishes the proof. \square

Remark It is also of interest to determine $b_i(0)$ for $i = 1, 2$ in order to find the best bounds for the continuation region. It seems difficult to determine these quantities in general, but an upper bound for the continuation region initially (and thus at all times) can be established by solving the free-boundary problem for the time-homogeneous volatility $\sigma(0, \pi)$. However, we expect these bounds to be rather crude, and therefore do not provide any details.

7. THE NORMAL PRIOR DISTRIBUTION

In this final section, we study the case of a normal prior distribution in further detail. In particular, we show that the kernel in the integral equations determined in Section 5 can be calculated explicitly for normal priors. For similar results in the case of the two-point distribution, see [9].

First, recall from Section 3 that a normal prior distribution with mean m and variance γ^2 leads to a volatility surface $\sigma(\cdot, \cdot)$ that is symmetric around the line $\pi = 1/2$. As a result, the stopping boundaries b_1 and b_2 are also symmetric around $\pi = 1/2$ with $b_2(t) = 1 - b_1(t)$, so it suffices to solve a single integral equation to determine both boundaries. Next, recall that the conditional distribution $\mu_{t,x}$ is normal with standard deviation $\gamma(t) := \gamma/\sqrt{1+t\gamma^2}$. Consequently, the x -value that gives $\pi(t, x) = b(t)$ is

such that the conditional drift equals

$$m(t) := \Phi^{-1}(b(t))\gamma/\sqrt{1+t\gamma^2}.$$

Now, given $s > 0$, let Y denote a $N(m(t)s, s+s^2\gamma^2(t))$ -distributed random variable. Then using (17), we calculate

$$\begin{aligned} K(t, s, b(t), b(t+s)) &:= \mathbb{P}\left(b(t+s) < \Pi_{t+s}^{t, b(t)} < 1 - b(t+s)\right) \\ &= \mathbb{P}\left(b(t+s) < \Phi\left(\frac{m(t) + \gamma^2(t)Y}{\gamma(t)\sqrt{1+s\gamma^2(t)}}\right) < 1 - b(t+s)\right) \\ &= \Phi(d_2) - \Phi(d_1), \end{aligned}$$

where

$$d_1 := \frac{\Phi^{-1}(b(t+s))\gamma(t)\sqrt{1+s\gamma^2(t)} - m(t)(1+s\gamma^2(t))}{\gamma^2(t)\sqrt{s+s^2\gamma^2(t)}}$$

and

$$d_2 := \frac{-\Phi^{-1}(b(t+s))\gamma(t)\sqrt{1+s\gamma^2(t)} - m(t)(1+s\gamma^2(t))}{\gamma^2(t)\sqrt{s+s^2\gamma^2(t)}}.$$

Thus the kernel K appearing in the integral equation (39) and in the corresponding equation for the infinite-horizon formulation is explicit.




FIGURE 1. The boundaries b_1 and b_2 calculated numerically for the case of $N(m, 1)$ -prior (note that the boundaries do not depend on $m \in \mathbb{R}$) and the cost of observation $c = 0.5$ per unit time.

REFERENCES

- [1] Bain, A., Crisan, D. *Fundamentals of stochastic filtering*. Stochastic Modelling and Applied Probability, 60. Springer, New York, 2009.
- [2] Bather, J. A. Bayes procedures for deciding the sign of a normal mean. *Proc. Cambridge Philos. Soc.*, vol. 58, 1962, 599-620.
- [3] Bickel P. J., Yahav, Y. A. On the Wiener process approximation to Bayesian sequential testing problems. *Proc. Sixth Berkeley Symp. Math. Statist. Probab.*, vol. 1, 1972, 57-84.
- [4] Breakwell, J., Chernoff, H. Sequential tests for the mean of a normal distribution. II. (Large t). *Ann. Math. Statist.*, vol. 35, 1964, 162-173.
- [5] Chernoff, H. Sequential tests for the mean of a normal distribution. *Proc. 4th Berkeley Sympos. Math. Statist. and Prob.*, vol. 1, 1961, 79-91.
- [6] Chernoff, H. Sequential tests for the mean of a normal distribution III (small t). *Ann. Math. Statist.*, vol. 36, 1965, 28-54.
- [7] Chernoff, H. Sequential tests for the mean of a normal distribution IV (discrete case). *Ann. Math. Statist.*, vol. 36, 1965, 55-68.
- [8] Ekström, E. Properties of American option prices. *Stochastic Process. Appl.*, vol. 114, no. 2, 2004, 265-278.
- [9] Gapeev, P., Peskir, G. The Wiener sequential testing problem with finite horizon. *Stoch. Stoch. Rep.* vol 76, no. 1, 2004, 59-75.
- [10] Jacka, S. Optimal stopping and the American put. *Math. Finance*, vol. 1, 1991, 1-14.
- [11] Janson, S. and Tysk, J. Volatility time and properties of option prices. *Ann. Appl. Probab.*, vol. 13, no. 3, 2003, 890-913.
- [12] Karatzas, I. Shreve, S. *Brownian motion and stochastic calculus. Second edition*. Graduate Texts in Mathematics, 113. Springer-Verlag, New York, 1991.
- [13] Karatzas, I., Shreve, S. *Methods of mathematical finance*. Applications of Mathematics, 39. Springer-Verlag, New York, 1998.
- [14] A. Klenke, *Probability Theory: A Comprehensive Course* Universitext. Springer-Verlag, London, 2008.
- [15] Lai, T. L. Nearly optimal sequential tests of composite hypotheses. *Ann. Statist.*, vol. 16, 1988, 856-886.
- [16] Lai, T. L. On optimal stopping problems in sequential hypothesis testing. *Statistica Sinica*, vol. 7, 1997, 33-51.
- [17] Øksendal, B. *Stochastic differential equations. An introduction with applications. Sixth edition*. Universitext. Springer-Verlag, Berlin, 2003.
- [18] Peskir, G. A change-of-variable formula with local time on curves. *J. Theoret. Probab.*, vol. 18, 2005, 499-535.
- [19] Peskir, G. On the American option problem. *Math. Finance*, vol. 15, no. 1, 2005, 169-181.
- [20] Peskir, G., Shiryaev, A. *Optimal stopping and free-boundary problems*. Lectures in Mathematics, ETH Zürich. Birkhäuser Verlag, Basel, 2006.

- [21] Shiryaev, A. N. Two problems of sequential analysis. *Cybernetics*, vol. 3, no. 2, 1967, 63-69.
- [22] Zhitlukhin, M., Muravlev, A. On Chernoff's Hypotheses Testing Problem for the Drift of a Brownian Motion, *Theory of Probability & Its Applications*, vol. 57, no. 4, 2013, 708-717.

UPPSALA UNIVERSITY, BOX 480, 75106 UPPSALA, SWEDEN.