



The value of foresight

Philip A. Ernst^{a,*}, L.C.G. Rogers^b, Quan Zhou^a

^a *Department of Statistics, Rice University, Houston, TX 77005, USA*

^b *Statistical Laboratory, University of Cambridge, Cambridge CB3 0WB, UK*

Received 20 July 2016; received in revised form 9 December 2016; accepted 22 March 2017

Available online 31 March 2017

We dedicate this work to our colleague, mentor, and friend, Professor Larry Shepp (1936–2013)

Abstract

Suppose you have one unit of stock, currently worth 1, which you must sell before time T . The Optional Sampling Theorem tells us that whatever stopping time we choose to sell, the expected discounted value we get when we sell will be 1. Suppose however that we are able to see a units of time into the future, and base our stopping rule on that; we should be able to do better than expected value 1. But how much better can we do? And how would we exploit the additional information? The optimal solution to this problem will never be found, but in this paper we establish remarkably close bounds on the value of the problem, and we derive a fairly simple exercise rule that manages to extract most of the value of foresight.

© 2017 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

MSC: primary 60G40; secondary 60H30

Keywords: Explicit stopping rules; Excursion processes; Insider trading; Bermudan fixed-window lookback option; Brownian motion

1. Introduction

What is the value of foresight in a financial market? This is a question that intrigued Larry Shepp and seems an interesting question in the context of insider trading; if we could know one

* Corresponding author.

E-mail addresses: philip.ernst@rice.edu (P.A. Ernst), chris@statslab.cam.ac.uk (L.C.G. Rogers), qz9@rice.edu (Q. Zhou).

minute in advance what the price of a stock was going to do, what would we be prepared to pay for that information? Of course, it is rather fanciful to imagine that we could possibly be told the price of the stock at some time in the future, but we might imagine a situation where some market participants received information only after a delay, which would confer the same kind of advantage on those who got the information earlier. In a modern financial market, any such differences would be measured in microseconds, a timescale on which conventional models of stock prices could not be trusted, but one of the first observations of this paper is that the value of foresight can be equivalently interpreted in terms of the value of a fixed-window lookback option; at a time of your choosing, you may sell the stock for the best price which the stock achieved in the previous a units of time. This transforms the question into an American option pricing problem, but not one that is possible to solve in closed form, since the state variable at time t is the entire history of the stock from time $t - a$ to time t —a path-valued state. Moreover, even if we were to discretize time, the state vector will be high dimensional, so existing numerical methods will struggle to cope. Nevertheless, recent developments allow good progress to be made on the question, as we shall see.

To begin with, we set some notation. We shall take the sample space to be the path space $C(\mathbb{R}^+, \mathbb{R})$ with canonical process W , Wiener measure P and the usual P -augmentation (\mathcal{F}_t) of the canonical filtration of W . We denote by \mathcal{T} the class of all (\mathcal{F}_t) -stopping times. We fix some $a > 0$ which represents the foresight available to the insider. The insider may choose to stop at any stopping time of the larger filtration $(\mathcal{F}'_t) \equiv (\mathcal{F}_{t+a})$; we denote the set of all (\mathcal{F}'_t) -stopping times by \mathcal{T}' .

At an abstract level, this set-up can be considered an example of *grossissement*—the enlargement of filtrations. This theory was developed from the late 1970s on, starting with the works of Barlow [6], Jeulin & Yor [25–27], and further developed by others including Yoeurp [35], and by Itô’s extension of the stochastic integral (see [23]). The topic has since continued to flourish, most recently in the works of [1,2,11], and [16], primarily because of its natural connection with insider trading. Varied formulations of the insider trader’s advantage over other agents in a financial market are addressed in [3,5,8,12,17,19–21,24,28,29], and [30], among others. It has to be understood that the theory of enlargement of filtrations is not a *universal* theory; results are only established for particular classes of enlargement, such as filtrations enlarged with an honest time, or filtrations with an initial enlargement. The results proved say that if the enlargement has one of these particular structural forms, then any (\mathcal{F}_t) -local martingale is a (\mathcal{F}'_t) -semimartingale, and the semimartingale decomposition is then identified. The following proposition shows that none of these general results can be applied to the problem we consider here.

Proposition 1.1. *The process W is not a semimartingale in the filtration (\mathcal{F}'_t) .*

Proof. Consider the simple integrands

$$H_t^n \equiv n^{-1/2} \sum_{j=1}^n \text{sgn}(\Delta_j^n) I_{\{(j-1)a < nt \leq ja\}}, \tag{1}$$

where

$$\Delta_j^n \equiv W(ja/n) - W((j - 1)a/n). \tag{2}$$

The processes H^n are left-continuous, bounded, and (\mathcal{F}_t') -previsible; indeed, H_t^n is measurable on $\mathcal{F}'_0 \equiv \mathcal{F}_a$. Now consider the (elementary) stochastic integral

$$H^n \cdot W = n^{-1/2} \sum_{j=1}^n |\Delta_j^n| = n^{-1} \sum_{j=1}^n \sqrt{n} |\Delta_j^n|. \tag{3}$$

The random variables $\sqrt{n} \Delta_j^n$ are independent zero-mean Gaussians with common variance a . By the Weak Law of Large Numbers, $H^n \cdot W$ converges in probability to $E|W_a|$ as $n \rightarrow \infty$. But the Dellacherie–Bichteler theorem (see, for example, Theorem IV.16.4 in [34]) says that W is a semimartingale if and only if whenever a sequence H^n of bounded previsible simple processes tends uniformly to zero, then the simple stochastic integrals $H^n \cdot W$ tend to zero in probability. We conclude that W is not an (\mathcal{F}_t') -semimartingale. \square

The message from Proposition 1.1 is that none of the results from enlargement of filtrations will help us here – the problem addressed is *concrete, challenging, and not amenable to general theory* – which is why it appealed to Larry Shepp.

So what *are* we able to do? To begin with, applying the methods of [33], we obtain remarkably tight upper and lower bounds on the value of foresight. This methodology is based on simulations, so there is no simple interpretation of the exercise rule which arises. However, in Section 3 we develop simple and transparent rules based on heuristic arguments which we are then able to compare with the bounds from Section 2; the resulting (explicit) rules essentially achieve the lower bound which comes from the simulation approach of [33] applied in Section 2.

Preliminaries.

To start with, we set out some notation and make various standardizations of the question. There is a fixed time horizon $T > a > 0$ by which time the investor must have sold the stock. The stock price process S will be the solution to an SDE

$$dS_t = S_t(\sigma dW_t + \mu_t dt) I_{\{t \leq T\}}, \quad S_0 = 1, \tag{4}$$

where μ is a bounded (\mathcal{F}_t) -previsible process, and $\sigma > 0$ is a constant. Notice that $S_t = S_T$ for all $t \geq T$, and where necessary we make the convention that $S_t = 1$ for $t < 0$. It is immediate from the definitions that $\tau' \in T'$ if and only if $\tau' + a \in T$. We let Q denote the pricing measure, given in terms of the Cameron–Martin–Girsanov martingale

$$d\Lambda_t = \frac{1}{\sigma} \Lambda_t(r - \mu_t) dW_t, \quad \Lambda_0 = 1, \tag{5}$$

by

$$\left. \frac{dQ}{dP} \right|_{\mathcal{F}_t} = \Lambda_t. \tag{6}$$

What is the time-0 value of stopping at $\tau' \in T'$? The obvious immediate answer to this question is just $E^Q[\exp(-r\tau') S_{\tau'}]$, but is it clear that there is no issue arising from the fact that $S_{\tau'}$ is not $\mathcal{F}_{\tau'}$ -measurable? We can see that this is in fact correct by the following argument. At time τ' , the investor receives $S_{\tau'}$ which he can then place in the bank account for a units of time, so that at (\mathcal{F}_t) -stopping time $\tau = \tau' + a$ he has $e^{ra} S_{\tau'} = e^{ra} S_{\tau-a}$. This random variable is \mathcal{F}_{τ} -measurable, so the time-0 value of it is just given by the usual expression,

$$E^Q[e^{-r\tau} e^{ra} S_{\tau-a}] = E^Q[e^{-r\tau'} S_{\tau'}],$$

as expected. So what we find is that the time-0 value of stopping at $\tau' \in \mathcal{T}'$ is the Q -expectation of the discounted value of the stock at the time of exercise. We may as well therefore work with the discounted stock, and work in the pricing measure Q . Equivalently, we may (and shall) assume for simplicity that

$$r = \mu = 0. \tag{7}$$

Since the choice of σ amounts to the choice of a time unit, we may and shall assume that

$$\sigma = 1. \tag{8}$$

Our model for the asset dynamics is therefore

$$S_t = \exp\left(W_t - \frac{1}{2}t\right) \equiv \exp(X_t), \tag{9}$$

where X is the drifting Brownian motion

$$X_t = W_t + ct, \tag{10}$$

with the special value¹ $c = -\frac{1}{2}$. Notice that we have by convention fixed $S_0 = 1$.

What we want to understand then is

$$v(a) \equiv \sup_{\tau' \in \mathcal{T}', 0 \leq \tau' \leq T} E[S_{\tau'}] \equiv \sup_{\tau \in \mathcal{T}, a \leq \tau \leq T+a} E[S_{\tau-a}]. \tag{11}$$

It is clear that v will be increasing, and $v(0) = 1$, but our aim is to determine as accurately as possible what $v(a)$ is, and to identify a good approximation to the optimal stopping time. The first item of business, dealt with in Section 2, is to show that the value v can be alternatively expressed as

$$v(a) = \sup_{\tau \in \mathcal{T}, 0 \leq \tau \leq T} E[Z_\tau], \tag{12}$$

where²

$$Z_t \equiv \sup\{S_u : t - a \leq u \leq t\}. \tag{13}$$

In other words, the value $v(a)$ is the value of an *American fixed-window lookback option*.

2. Foresight as lookback

Recalling the convention that $S_u = 1$ for $u < 0$, and $S_u = S_T$ for $u \geq T$, we have a simple proposition.

Proposition 2.1. *With τ denoting a generic (\mathcal{F}_t) -stopping time,*

$$v(a) \equiv \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}] = \sup_{0 \leq \tau \leq T} E[Z_\tau], \tag{14}$$

where $Z_t \equiv \sup\{S_u : t - a \leq u \leq t\}$.

¹ Later on, we shall derive expressions for various probabilities and expectations associated with X , and it turns out to be notationally cleaner to work with a general drift, which is why we write X as (10).

² ... using the convention that $S_u = 1$ for $u < 0$, and $S_u = S_T$ for $u \geq T$...

Proof. Because $S_u = S_T$ for all $u \geq T$, it is clear that $Z_t \leq Z_{t \wedge T}$. Therefore, for any stopping time τ such that $a \leq \tau \leq T + a$, we have

$$S_{\tau-a} \leq Z_\tau \leq Z_{\tau \wedge T}.$$

Therefore

$$v(a) \equiv \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}] \leq \sup_{0 \leq \tau \leq T} E[Z_\tau]. \tag{15}$$

For the reverse inequality, suppose that τ is a stopping time, $0 \leq \tau \leq T$, and define a new random time $\tilde{\tau}$ by

$$\tilde{\tau} = \inf\{u \geq \tau \vee a : S_{u-a} = Z_\tau\}. \tag{16}$$

Clearly $\tau \vee a \leq \tilde{\tau} \leq \tau + a$. We claim that $\tilde{\tau}$ is a stopping time, as follows:

$$\begin{aligned} \{\tilde{\tau} \leq v\} &= \{\text{for some } u \in [\tau \vee a, v], S_{u-a} = Z_\tau\} \\ &= \{\tau \vee a \leq v\} \cap \{\text{for some } u \in [(\tau \vee a) \wedge v, v], S_{u-a} = Z_{\tau \wedge v}\} \\ &\in \mathcal{F}_v, \end{aligned}$$

since the event $\{\exists u \in [(\tau \vee a) \wedge v, v], S_{u-a} = Z_{\tau \wedge v}\}$ is \mathcal{F}_v -measurable, as is $(\tau \vee a) \wedge v$. Now we see that

$$Z_\tau = S_{\tilde{\tau}-a}, \tag{17}$$

and therefore

$$E[Z_\tau] = E[S_{\tilde{\tau}-a}] \leq \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}], \tag{18}$$

since $a \leq \tilde{\tau} \leq T + a$. Since $0 \leq \tau \leq T$ was any stopping time, we deduce that

$$\sup_{0 \leq \tau \leq T} E[Z_\tau] \leq \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}], \tag{19}$$

and the proof is complete. \square

The importance of Proposition 2.1 is that it turns the problem of calculating the value of foresight into the calculation of an American fixed-window lookback option. By discretizing the time, we shall instead calculate numerically the value of a Bermudan fixed-window lookback option. Of course, we need to account for the difference between American and Bermudan prices, but this is in essence a solved problem; see Broadie, Glasserman & Kou [10].

So we shall take the standardized asset dynamics (9), fix some time horizon $T > 0$ which is subdivided into N_T steps of length $h = T/N_T$, and consider the problem of bounding

$$v_h(a) \equiv \sup_{0 \leq \tau_h \leq T} E[Z_{\tau_h}^{(h)}], \tag{20}$$

where τ_h is a stopping time taking values in the set $h\mathbb{Z}^+$, and

$$Z_t^{(h)} = \max\{S_{kh} : t - a \leq kh \leq t\}. \tag{21}$$

Since this discretization is now fixed for the rest of the section, we shall drop the appearance of h in the notation and refer to Z for $Z^{(h)}$, v for v_h , τ for τ_h . We are now exclusively considering the

optimal stopping of a functional of a discrete-time Markov process. A moment’s thought shows that the process Z is not Markov, but the process

$$x_t = (S_{t-mh}, \dots, S_t) \quad (t \in h\mathbb{Z}^+) \tag{22}$$

is Markovian (where $mh = a$), and the payoff process Z is simply a function of the Markov process x :

$$Z_t = g(x_t), \tag{23}$$

where $g(x) \equiv \max\{x_0, \dots, x_m\}$ is the largest component of the $(m + 1)$ -vector x . We use the approach of [33], which is a combination of several techniques developed over the last twenty or so years, and in summary consists of the four steps:

- pretend that the stopping reward process Z is itself Markovian, and by discretizing Z onto a suitably-chosen finite set of values estimate the transition probabilities of this finite state Markov chain by simulation (this is the approach of Barraquand & Martineau [7]);
- solve the optimal stopping problem for this finite state Markov chain by dynamic programming;
- use the solution to generate a stopping rule whose performance is evaluated by simulation;
- use the dual characterization of the value of the problem (see [4,18,32]) to find a hedging martingale.

The method is fully explained in [33], and illustrated with examples, so there is no need to discuss it further here, except to highlight one point, which is used in various places in [33] and is needed here. Suppose that $(M_t)_{t \geq 0}$ is any strictly positive martingale, $M_0 = 1$. Then we may equivalently express the value

$$\begin{aligned} v(a) &\equiv \sup_{0 \leq \tau \leq T} E[Z_\tau] \\ &= \sup_{0 \leq \tau \leq T} E[M_\tau (Z_\tau / M_\tau)] \\ &= \sup_{0 \leq \tau \leq T} \tilde{E}[Z_\tau / M_\tau], \end{aligned} \tag{24}$$

where the probability \tilde{P} equivalent to P is defined by using the likelihood-ratio martingale M :

$$\left. \frac{d\tilde{P}}{dP} \right|_{\mathcal{F}_t} = M_t. \tag{25}$$

In the present application, it is natural to use $M_t = S_t$ as the change-of-measure martingale, and the effect of this is to change the Brownian motion W into a Brownian motion with drift 1. Thus when we do simulations, we simulate a Brownian motion with drift 1 in place of W , and the stopping reward process is changed to Z_t/S_t . This is a good thing to do, because the value of Z_t will be close to S_t , so the binning procedure of the Barraquand–Martineau approach should achieve a lot better accuracy than we would get without this measure transformation. Moreover, after the change of measure the stopping rule depends on $(Z_t/S_t, t)$ instead of (S_t, t) , which is more sensible and resembles the stopping rule we will propose below. We do indeed find that the accuracy is substantially improved by doing this change of measure; the results are reported in [Table 1](#).

Table 1

Upper and lower bounds of $\sup_{\tau < T} E[Z_\tau]$ from simulation using method of [33] with $h = 1/2500$, $N_T = 250$. ‘SE’ is the standard error, in basis points ($1e-4$). Simulation parameters: number of bins = 200, number of samples per bin = 200, lower-bound simulation = 50,000, upper-bound simulation = 10,000, sub-simulation per step = 50.

a/h	Lower	SE (low)	Upper	SE (up)	Gap (%)
1	1.054	0.43	1.055	0.53	0.09
2	1.074	0.62	1.076	0.77	0.18
3	1.088	0.73	1.092	0.90	0.37
4	1.100	0.83	1.104	1.0	0.37
5	1.109	0.95	1.114	1.1	0.45
6	1.117	1.0	1.123	1.2	0.53
7	1.123	1.1	1.131	1.2	0.71
8	1.129	1.2	1.137	1.3	0.70
9	1.135	1.3	1.144	1.4	0.79
10	1.140	1.4	1.149	1.4	0.78
11	1.144	1.4	1.154	1.5	0.87
12	1.149	1.5	1.159	1.6	0.86
13	1.152	1.6	1.164	1.7	1.03
14	1.156	1.6	1.168	1.7	1.03
15	1.159	1.7	1.172	1.8	1.11
16	1.163	1.8	1.175	1.8	1.02
17	1.166	1.8	1.179	1.9	1.10
18	1.168	1.9	1.182	1.8	1.18
19	1.171	1.9	1.185	1.9	1.18
20	1.174	2.0	1.188	1.9	1.18

3. Explicit stopping rules

As was stated earlier, the stopping rules which are derived in Section 2 are the output of a simulation; they have no particular structure or interpretation, and a different simulation run will generate a different stopping rule. The methodology of [33] is generic, and works just the same for essentially any Markov process, and any stopping function of the process, but we may hope to improve in individual applications by exploiting the specific structure of that application, which is what we shall do here.

We work in terms of the log price $X_t = W_t + ct$, and use the notations

$$\bar{X}_t \equiv \sup_{0 \leq u \leq t} X_u, \quad \bar{X}_{[t-a,t]} \equiv \sup_{t-a \leq u \leq t} X_u \equiv \log Z_t. \tag{26}$$

The first time we may need to consider stopping is the stopping time

$$\tau_0 \equiv \inf\{ t : \bar{X}_{[t-a,t]} = X_{t-a} \}, \tag{27}$$

because up until that time we have $X_{t-a} < \bar{X}_{[t-a,t]}$, so it will be suboptimal to exercise—waiting a little longer may improve, and will not make the reward less. But at time τ_0 , continuation means we have to let go of the good value X_{τ_0-a} in the hope of doing better in the future, and we may in fact do worse. Whether we should optimally continue will depend on the entire path of X from $\tau_0 - a$ to τ_0 , but we will simplify to consider only rules where continuation is decided by the value of $X_{\tau_0} - X_{\tau_0-a}$; if this is higher than some threshold $q < 0$ we shall continue, otherwise we stop.

An important observation is the fact that

$$\bar{X}_{[\tau_0-a, \tau_0]} = \bar{X}_{\tau_0}. \tag{28}$$

We prove this by contradiction. Suppose $\bar{X}_{[\tau_0-a, \tau_0]} \neq \bar{X}_{\tau_0}$, which implies $\bar{X}_{\tau_0} = X_m$ for some $m \in [0, \tau_0 - a)$. Since \bar{X}_m is the maximum till τ_0 , it must be greater than X_t for $t \in (m, m + a]$. However, this means $m = \tau_0$, which gives the contradiction. Therefore we have

$$\tau_0 = \inf\{ t : \bar{X}_t = X_{t-a} \}. \tag{29}$$

Now it is clear that the choice of threshold value q will depend on the time to go when we have to make the stop/continue decision, and this makes complete solution of the problem much more complicated. So what we shall do is to propose a *modified* problem, where the time by which we must sell is not a fixed time T but an independent random time $\alpha \sim \exp(\eta)$ for some $\eta > 0$. This gives us a renewal property that allows us to make progress, and obtain explicit expressions. When it comes to converting the solution to the modified problem into an exercise rule for the original problem, what we do is set the threshold according to the value of η which makes the expectation of α equal to the time to go. The point is that the problem considered with an exponential lifetime can be solved completely by excursion theory, whereas the problem with a fixed lifetime cannot.

For this modified problem, let us define a stopping rule $R(q)$ depending on the chosen threshold as follows:

- (i) Wait until $\tau_0 \wedge \alpha$;
- (ii) If $\alpha < \tau_0$, stop and receive $Z_\alpha = \exp(\bar{X}_\alpha)$;
- (iii.a) If $\tau_0 < \alpha$ and $X_{\tau_0} - \bar{X}_{\tau_0} < q$, stop and receive $Z_{\tau_0} = \exp(\bar{X}_{\tau_0})$;
- (iii.b) If $\tau_0 < \alpha$ and $X_{\tau_0} - \bar{X}_{\tau_0} > q$, forget and continue.

In the final eventuality (iii.b), ‘forget and continue’ means that we wipe away the whole path of X in the time interval $[0, \tau_0)$, keeping only the value X_{τ_0} , and restart the rule from that point. If the value of this strategy is K , then we have the identity

$$K = E[\exp(\bar{X}_\alpha) : \alpha < \tau_0] + E[\exp(\bar{X}_{\tau_0}) : \tau_0 < \alpha, X_{\tau_0} - \bar{X}_{\tau_0} < q] + K E[\exp(X_{\tau_0}) : \tau_0 < \alpha, X_{\tau_0} - \bar{X}_{\tau_0} > q]. \tag{30}$$

It is therefore apparent that we can evaluate this particular stopping rule provided we can find explicit expressions for

$$A_0 = E[\exp(\bar{X}_\alpha) : \alpha < \tau_0], \tag{31}$$

$$A_-(q) = E[\exp(\bar{X}_{\tau_0}) : \tau_0 < \alpha, X_{\tau_0} - \bar{X}_{\tau_0} < q], \tag{32}$$

$$A_+(q) = E[\exp(X_{\tau_0}) : \tau_0 < \alpha, X_{\tau_0} - \bar{X}_{\tau_0} > q]. \tag{33}$$

If φ denotes the density of the standard $N(0, 1)$ distribution, and Φ denotes its distribution function, we have the following result.

Proposition 3.1. For $q < 0$, denote

$$\beta = \sqrt{c^2 + 2\eta}, \tag{34}$$

$$\nu_a = \frac{2}{\sqrt{a}} \varphi(c\sqrt{a}) - 2c \Phi(-c\sqrt{a}), \tag{35}$$

$$\begin{aligned}
 \nu_\alpha &= 2\beta \left(\Phi(\beta\sqrt{a}) - \frac{1}{2} \right) - 2c \left(\Phi(c\sqrt{a}) - \frac{1}{2} \right) + \frac{2}{\sqrt{a}} \{ \varphi(\beta\sqrt{a}) - \varphi(c\sqrt{a}) \}, \\
 \nu &\equiv \nu_a + \nu_\alpha = \frac{2}{\sqrt{a}} \varphi(\beta\sqrt{a}) + 2\beta\Phi(\beta\sqrt{a}) - c - \beta,
 \end{aligned}
 \tag{36}$$

$$\psi_0(q) = e^{-\eta a} \left\{ \frac{2}{\sqrt{a}} \varphi((q - ca)/\sqrt{a}) - 2c \Phi((q - ca)/\sqrt{a}) \right\},
 \tag{37}$$

$$\begin{aligned}
 \psi_1(q) &= e^{(c+\frac{1}{2}-\eta)a} \left[\frac{2}{\sqrt{a}} \{ \varphi(\bar{c}\sqrt{a}) - \varphi((q - \bar{c}a)/\sqrt{a}) \} \right. \\
 &\quad \left. + 2\bar{c} \{ \Phi(\bar{c}\sqrt{a}) - \Phi(-(q - \bar{c}a)/\sqrt{a}) \} \right],
 \end{aligned}
 \tag{38}$$

where $\bar{c} \equiv 1 + c$. Then:

$$A_0 = \frac{\nu_\alpha + (1 - e^{-\eta a})\nu_a}{\nu - 1},
 \tag{39}$$

$$A_-(q) = \frac{\psi_0(q)}{\nu - 1},
 \tag{40}$$

$$A_+(q) = \frac{\psi_1(q)}{\nu - 1}.
 \tag{41}$$

Proof. The process $Y_t \equiv X_t - \bar{X}_t$ is a diffusion taking values in the negative half-line, and reflecting from zero. The process \bar{X}_t is its local time at zero, and the path of Y can be decomposed into a Poisson process of excursions, as Itô [22] explained. For a more extensive discussion of excursion theory, in particular, the notion of marked excursion processes, see [31] or Chapter VI.8 of [34].

The process X runs until $\tau_0 \wedge \alpha$, which happens in the first excursion of Y which either lasts at least a , or contains an η -mark. Let n denote the excursion measure, ζ denote the lifetime of an excursion, P^c denote the law of $X_t = W_t + ct$, and H_b denote the first time X hits b . The excursion law can be characterized as the limit as $\varepsilon \downarrow 0$ of the (rescaled) law of X started at $-\varepsilon < 0$ run until it hits zero; see VI.50.20 in [34]. The rescaling required is to multiply by $n(f : \inf f < -\varepsilon)$, and it is known that $n(f : \inf f < -\varepsilon) \sim \varepsilon^{-1}$ (see VI.51.2 in [34]). We therefore have

$$\begin{aligned}
 n(\zeta > a) &= \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} P^c[H_0 > a \mid X_0 = -\varepsilon] \\
 &= \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} P^c[\bar{X}_a < \varepsilon \mid X_0 = 0] \\
 &= \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} \left\{ \Phi\left(\frac{\varepsilon - ca}{\sqrt{a}}\right) - e^{2c\varepsilon} \Phi\left(\frac{-\varepsilon - ca}{\sqrt{a}}\right) \right\}
 \end{aligned}
 \tag{42}$$

$$\begin{aligned}
 &= \frac{2}{\sqrt{a}} \varphi(c\sqrt{a}) - 2c \Phi(-c\sqrt{a}), \\
 &\equiv \nu_a,
 \end{aligned}
 \tag{43}$$

where (42) is a standard result on Brownian motion; see, for example, [9]. By similar reasoning,

$$n(\alpha < \zeta < a) = \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} \int_0^a P^c[H_0 \in ds \mid X_0 = -\varepsilon] (1 - e^{-\eta s})$$

$$\begin{aligned}
 &= \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} \int_0^a \frac{\varepsilon e^{-\varepsilon^2/2s}}{\sqrt{2\pi s^3}} e^{c\varepsilon - c^2s/2} (1 - e^{-\eta s}) ds \\
 &= \int_0^a \frac{e^{-c^2s/2}}{\sqrt{2\pi s^3}} (1 - e^{-\eta s}) ds \\
 &= 2\beta \left(\Phi(\beta\sqrt{a}) - \frac{1}{2} \right) - 2c \left(\Phi(c\sqrt{a}) - \frac{1}{2} \right) \\
 &\quad + \frac{2}{\sqrt{a}} \{ \varphi(\beta\sqrt{a}) - \varphi(c\sqrt{a}) \} \\
 &\equiv \nu_\alpha
 \end{aligned} \tag{44}$$

where we recall that $\beta \equiv \sqrt{c^2 + 2\eta}$. The rate of excursions which either last at least a , or contain an η -mark is

$$n(\alpha \wedge a < \zeta) = n(\alpha < \zeta < a) + n(\zeta > a) \tag{45}$$

and this is simply the sum of the two expressions (43) and (44), which is therefore known explicitly. Simple calculations confirm that it is the expression ν defined at (36). Let B_a be the set of excursions for which $a < \alpha \wedge \zeta$, and let B_α be the set of excursions for which $\alpha < a \wedge \zeta$. Clearly, $n(B_a) = \nu_a e^{-\eta a}$, and $n(B_a \cup B_\alpha) = n(a \wedge \alpha < \zeta) = \nu$. The stopping time of interest to us, $\tau_0 \wedge \alpha$, occurs in the first excursion in $B_a \cup B_\alpha$. By Itô excursion theory, the distribution of the local time \bar{X} at that moment will be exponential with rate $\nu = n(B_a \cup B_\alpha)$; and the event that the excursion falls into B_a will be independent of the local time, with probability $n(B_a)/\{n(B_a) + n(B_\alpha)\}$. To summarize:

- $\bar{X}_{\tau_0 \wedge \alpha} \sim \exp(\nu)$;
- $P[\alpha < \tau_0] = \{ \nu_\alpha + (1 - e^{-\eta a})\nu_a \} / \nu$;
- $P[\tau_0 < \alpha] = \nu_a e^{-\eta a} / \nu$;

and from this the expression (39) for A_0 follows.

To deal with A_\pm , we need to find the measure of excursions which get to time a without killing, and which are in dy at time a . We use the reflection principle and the Cameron–Martin–Girsanov theorem to derive

$$\begin{aligned}
 g(y)dy &\equiv n(a < \zeta \wedge \alpha, Y_a \in dy) \\
 &= \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} P^c[H_0 \wedge \alpha > a, Y_a \in dy \mid Y_0 = -\varepsilon] \\
 &= \lim_{\varepsilon \downarrow 0} \varepsilon^{-1} e^{c(y+\varepsilon) - c^2a/2 - \eta a} \{ \varphi((y + \varepsilon)/\sqrt{a}) - \varphi((y - \varepsilon)/\sqrt{a}) \} dy / \sqrt{a} \\
 &= \frac{-2y}{a^{3/2}} \varphi((y - ca)/\sqrt{a}) e^{-\eta a} dy.
 \end{aligned} \tag{47}$$

Straightforward calculations lead us to

$$\begin{aligned}
 \int_{-\infty}^q g(y) dy &= e^{-\eta a} \left\{ \frac{2}{\sqrt{a}} \varphi((q - ca)/\sqrt{a}) - 2c \Phi((q - ca)/\sqrt{a}) \right\}, \\
 &\equiv \psi_0(q), \\
 \int_q^0 e^y g(y) dy &= e^{(c+\frac{1}{2}-\eta)a} \left[\frac{2}{\sqrt{a}} \{ \varphi(\bar{c}\sqrt{a}) - \varphi((q - \bar{c}a)/\sqrt{a}) \} \right]
 \end{aligned}$$

$$\begin{aligned} & \times 2\bar{c} \left\{ \Phi(\bar{c}\sqrt{a}) - \Phi(-(q - \bar{c}a)/\sqrt{a}) \right\} \Big], \\ & \equiv \psi_1(q). \end{aligned}$$

Arguments similar to those used for A_0 give us

$$\begin{aligned} A_-(q) &= E[\exp(\bar{X}_{\tau_0}) : \tau_0 < \alpha, X_{\tau_0} - \bar{X}_{\tau_0} < q] = \frac{\psi_0(q)}{\nu-1}; \\ A_+(q) &= E[\exp(X_{\tau_0}) : \tau_0 < \alpha, X_{\tau_0} - \bar{X}_{\tau_0} > q] = \frac{\psi_1(q)}{\nu-1}. \end{aligned}$$

These are the forms claimed in (40) and in (41). \square

Now we can draw all the pieces together. Proposition 3.1 was a step on the way to evaluating K , the value of the proposed strategy; we found K expressed in terms of A_0 , $A_-(q)$ and $A_+(q)$ at (30), and it remains just to write the answer cleanly.

Theorem 3.2. *The value of the rule $R(q)$ is given by*

$$K = \frac{A_0 + A_-(q)}{1 - A_+(q)} = \frac{\nu_\alpha + (1 - e^{-\eta a})\nu_a + \psi_0(q)}{\nu - 1 - \psi_1(q)}. \tag{48}$$

The denominator $\nu - 1 - \psi_1(q)$ is always positive.

Proof. The expression (48) is a trivial arrangement of (30), so all that remains is to deal with the final assertion. Using the fact that $c = -\frac{1}{2}$, we see that

$$\psi_1(-\infty) = \uparrow \lim_{q \downarrow -\infty} \psi_1(q) = e^{-\eta a} \left\{ \frac{2}{\sqrt{a}} \varphi(\sqrt{a}/2) - \Phi(-\sqrt{a}/2) \right\},$$

is positive. Therefore

$$\begin{aligned} \nu - 1 - \psi_1(q) &> \nu - 1 - \psi_1(-\infty) \\ &= \nu_a + \nu_\alpha - 1 - \psi_1(-\infty) \\ &> \nu_a - 1 - \psi_1(-\infty) \\ &= \frac{2}{\sqrt{a}} \varphi(\sqrt{a}/2) + \Phi(\sqrt{a}/2) - 1 - \psi_1(-\infty) \\ &= (1 - e^{-\eta a}) \left\{ \frac{2}{\sqrt{a}} \varphi(\sqrt{a}/2) - \Phi(-\sqrt{a}/2) \right\} \\ &> 0. \quad \square \end{aligned}$$

Using the stopping rule $R(q)$.

Rule 1. How is the stopping rule $R(q)$ analyzed in the preceding section relevant to the original problem? Holding η fixed, there will be an optimal $q^* = q^*(\eta)$ which maximizes the value K given by (48). Notice that as $\eta \uparrow \infty$ we shall have $q^*(\eta) \uparrow 0$; the larger the η , the less the expected time remaining after τ_0 and therefore the less likely one is to continue. If we fix η and let $a \downarrow 0$, we shall also have $q^*(\eta) \uparrow 0$. This is so because when $a = 0$ the reward process $Z_{\tau \wedge T}$ is a martingale for any valid exercise rule τ , by the optimal sampling theorem. Fig. 1 characterizes these relationships. Note that for small values of a we find the value of lookback behaves roughly like \sqrt{a} . This is consistent with the findings of [8], and is entirely to be expected; when we stop at time τ , we receive the current value (mean 1) plus the maximum of $S_t - S_\tau$ over $t \in [\tau - a, \tau]$, and by the Brownian modulus of continuity this maximum varies as \sqrt{a} .

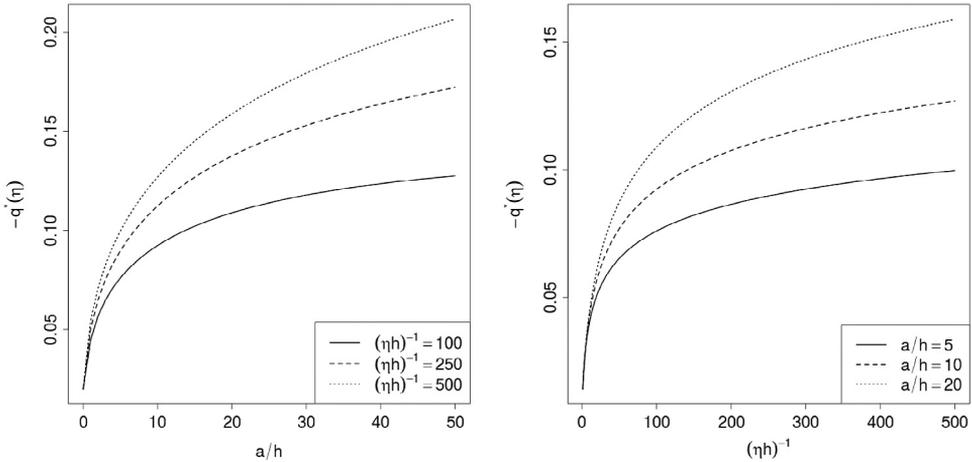


Fig. 1. $h = 1/2500$ is the time interval. a/h is the number of steps that we can look back. $(\eta h)^{-1}$ is the expected number of remaining steps. $\exp(-q^*(\eta))$ can be interpreted as an estimate for the maximum expected reward, which must be larger either if we can “look back” for more steps or if we have more time remaining.

Now recall how the stopping rule works; we let the process run until the stopping time

$$\tau_0 = \inf\{ t : \bar{X}_{[t-a,t]} = X_{t-a} \} = \inf\{ t : \bar{X}_t = X_{t-a} \},$$

and at that moment we stop if and only if $Y_{\tau_0} \equiv X_{\tau_0} - \bar{X}_{\tau_0} < q$, else we forget and continue. What value of q do we use? A natural choice is to take

$$q = q^* ((T - \tau_0)^{-1}). \tag{49}$$

This is because at time τ_0 there is time $(T - \tau_0)$ still to go, and an exponential random variable with rate $(T - \tau_0)^{-1}$ has this as its mean. The result of using this rule is shown in the left-hand panel of Fig. 2. In this figure, $h = 1/2500$ can be thought of as one trading day; by using $T = 0.1$ we assume the option contract will expire in 250 trading days. We have assumed $\sigma = 1$, which after rescaling amounts to an annual volatility of approximately 0.3. The dots (evaluated by simulating 50,000 runs of the rule) are visibly close to (but below) the lower bound we obtained by the Barraquand–Martineau technique of [33].

Rule 2. Can we do better than this? Indeed we can. Firstly, we observe that

$$K(q^*(\eta)) = \exp(-q^*(\eta)). \tag{50}$$

This is because when we arrive at time τ_0 , we have to choose between stopping and receiving S_{τ_0-a} , or continuing and receiving $K(q^*(\eta)) \cdot S_{\tau_0}$ in expectation. Optimal behavior requires us to continue if and only if

$$K(q^*(\eta)) \cdot S_{\tau_0} > S_{\tau_0-a}.$$

On the other hand, the rule $R(q^*(\eta))$ requires us to continue if and only if

$$S_{\tau_0} > S_{\tau_0-a} \cdot \exp(q^*(\eta)).$$

Since $q^*(\eta)$ was chosen optimally, (50) must therefore hold.

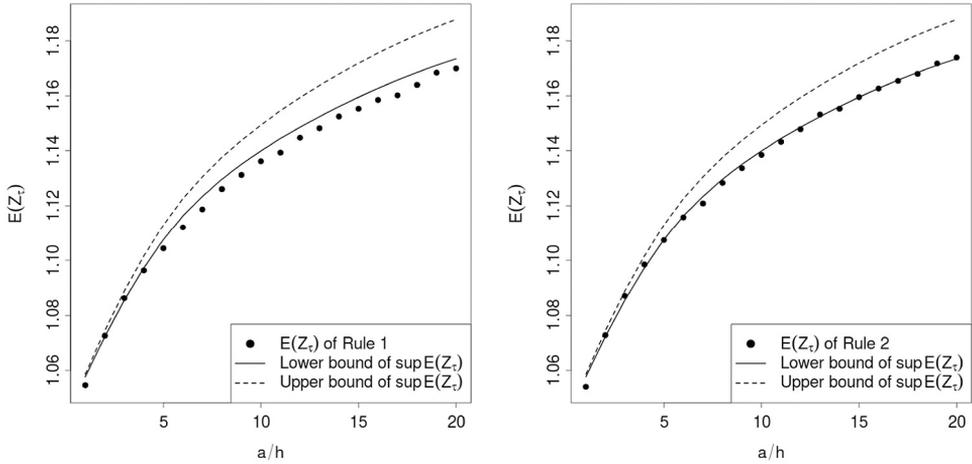


Fig. 2. The simulation results with $h = 1/2500$ and $N_T = 250$.

Table 2

$E[Z_T]$ estimates from simulation of Rules 1 and 2 with $h = 1/2500$ and $N_T = 250$, averaged over 50,000 sample paths. Standard errors are 0.001 (to one significant figure) in all cases.

a/h	Rule 1	Rule 2	Lower bound
1	1.055	1.054	1.054
2	1.073	1.073	1.074
3	1.086	1.087	1.088
4	1.096	1.098	1.100
5	1.104	1.107	1.109
6	1.112	1.116	1.117
7	1.119	1.121	1.123
8	1.126	1.128	1.129
9	1.131	1.134	1.135
10	1.136	1.139	1.140
11	1.139	1.143	1.144
12	1.145	1.148	1.149
13	1.148	1.153	1.152
14	1.152	1.155	1.156
15	1.155	1.160	1.159
16	1.158	1.163	1.163
17	1.160	1.165	1.166
18	1.164	1.168	1.168
19	1.168	1.172	1.171
20	1.170	1.174	1.174

Secondly, if $a = T$, then the value will just be the expectation of the overall maximum, $E \exp(\bar{X}_T)$. Routine but tedious calculations give

$$E[\exp(\bar{X}_a)] = \lambda(a) \equiv \left(2 + \frac{a}{2}\right) \Phi\left(\frac{\sqrt{a}}{2}\right) + \sqrt{a}\varphi\left(\frac{a}{4}\right). \tag{51}$$

If we had arrived at time τ_0 and it turns out that $\tau_0 = T - a$, then by forgetting and continuing we will actually receive expected reward $\lambda(a)S_{\tau_0}$, whereas if we used Rule 1 we would think we were going to receive $K(q^*(a^{-1})) \cdot S_{\tau_0} = e^{-q^*(a^{-1})} S_{\tau_0}$. This suggests that we modify Rule 1, replacing (49) by

$$q = q^*((T - \tau_0)^{-1}) - q^*(a^{-1}) - \log \lambda(a). \quad (52)$$

By making this modification, the stop/continue decision we make in the event that $T - \tau_0 = a$ will be exactly correct.

Of course, the argument just presented is only a rough heuristic, but if we look at the right-hand panel in Fig. 2 we see the results of using Rule 2 (see Table 2 for numerical values). The dots are now essentially coincident with the lower bounds, which is very encouraging, and argues for the use of Rule 2 rather than Rule 1. This rule is something which can be clearly motivated and precisely specified, in contrast to the randomly-generated rules which come from the Barraquand–Martineau technique in Section 2. One could continue to search for other explicit rules which do even better, but such a study is beyond the scope of this paper.

4. Supplemental materials

The code for all simulations can be found on <http://stat.rice.edu/~pe6/blb.cpp>. The pseudocode is available on http://stat.rice.edu/~pe6/Bermuda_code.pdf.

Acknowledgements

We are grateful for the efforts of two anonymous referees, whose very helpful suggestions greatly improved the quality of this work. The first author is grateful to Professor Larry Shepp's mentorship in mathematical finance (see [13–15]).

References

- [1] B. Acciaio, C. Fontana, C. Kardaras, Arbitrage of the first kind and filtration enlargements in semimartingale financial models, *Stochastic Process. Appl.* 126 (2016) 1761–1784.
- [2] A. Aksamit, T. Choulli, M. Jeanblanc, On an optional semimartingale decomposition and the existence of a deflator in an enlarged filtration, in: *In Memoriam Marc Yor–Séminaire de Probabilités XLVII*, Springer, 2015, pp. 187–218.
- [3] J. Amendinger, P. Imkeller, M. Schweizer, Additional logarithmic utility of an insider, *Stochastic Process. Appl.* 75 (1998) 263–286.
- [4] L. Andersen, M. Broadie, A primal–dual simulation algorithm for pricing multidimensional American options, *Manage. Sci.* 50 (2004) 1222–1234.
- [5] K. Back, Insider trading in continuous time, *Rev. Financ. Stud.* 5 (1992) 387–409.
- [6] M. Barlow, Study of a filtration expanded to include an honest time, *Probab. Theory Related Fields* 44 (1978) 307–324.
- [7] J. Barraquand, D. Martineau, Numerical valuation of high dimensional multivariate American securities, *J. Finan. Quant. Anal.* 30 (1995) 383–405.
- [8] E. Bayraktar, Z. Zhou, On an optimal stopping problem of an insider, *Theory Probab. Appl.* 61 (2016) 181–186.
- [9] A.N. Borodin, P. Salminen, *Handbook of Brownian Motion - Facts and Formulae*, Birkhäuser, 2012.
- [10] M. Broadie, P. Glasserman, S. Kou, A continuity correction for discrete barrier options, *Math. Finance* 7 (1997) 325–349.
- [11] E. Ekström, H. Wanntorp, Optimal stopping of a brownian bridge, *J. Appl. Probab.* 46 (2009) 170–180.
- [12] R. Elliott, H. Geman, B. Korkie, Portfolio optimization and contingent claim pricing with different information, *Stoch. Stoch. Rep.* 60 (1997) 185–203.
- [13] P.A. Ernst, On the arbitrage price of European call options, *Stoch. Models* 33 (2017) 48–58.
- [14] P.A. Ernst, D.P. Foster, L.A. Shepp, On optimal retirement, *J. Appl. Probab.* 51 (2014) 333–345.

- [15] P.A. Ernst, L.A. Shepp, Revisiting a theorem of L.A. Shepp on optimal stopping, *Commun. Stoch. Anal.* 9 (2015) 419–423.
- [16] C. Fontana, M. Jeanblanc, S. Song, On arbitrages arising with honest times, *Finance Stoch.* 18 (2014) 515–543.
- [17] A. Grorud, M. Pontier, Insider trading in a continuous time market model, *Int. J. Theor. Appl. Finance* 1 (1998) 331–347.
- [18] M.B. Haugh, L. Kogan, Pricing American options: a duality approach, *Oper. Res.* 52 (2004) 258–270.
- [19] P. Imkeller, Enlargement of the Wiener filtration by an absolutely continuous random variable via Malliavian’s calculus, *Probab. Theory Related Fields* 106 (1996) 105–135.
- [20] P. Imkeller, Random times at which insiders can have free lunches, *Stoch. Stoch. Rep.* 74 (2002) 465–487.
- [21] P. Imkeller, M. Pontier, F. Weisz, Free lunch and arbitrage possibilities in a financial market model with an insider, *Stochastic Process. Appl.* 92 (2001) 103–130.
- [22] K. Itô, Poisson point processes attached to Markov processes, in: *Proc. 6th Berk. Symp. Math. Stat. Prob.*, Vol. 3, 1971, pp. 225–240.
- [23] K. Itô, Extension of stochastic integral, in: *Proceedings of the International Symposium on Stochastic Differential Equations*, 1978, pp. 95–109.
- [24] J. Jacod, Grossissement initial, hypothèse (H’), et théorème de Girsanov, in: *Grossissements de Filtrations: Exemples et Applications*, Vol. 1118, Springer-Verlag, 1985.
- [25] T. Jeulin, M. Yor, Grossissement d’une filtration et semi-martingales: formules explicites, in: *Séminaire de Probabilités XII*, in: *Lecture Notes in Mathematics*, Springer-Verlag, 1978, pp. 78–97.
- [26] T. Jeulin, M. Yor, Inégalité de Hardy, semimartingales, et faux-amis, in: *Séminaire de Probabilités XIII*, in: *Lecture Notes in Mathematics*, Springer-Verlag, 1979, pp. 332–359.
- [27] T. Jeulin, M. Yor, Grossissement de filtrations: exemples et applications, in: *Lecture Notes in Mathematics*, Vol. 1118, Springer-Verlag, 1985.
- [28] I. Karatzas, I. Pikovsky, Anticipative portfolio optimization, *Adv. Appl. Probab.* 28 (1996) 1095–1122.
- [29] R. Mansuy, M. Yor, *Random Times and Enlargements of Filtrations in a Brownian Setting*, Springer, 2006.
- [30] P. Protter, *Stochastic Integration and Stochastic Differential Equations*, Springer, 2005.
- [31] L.C.G. Rogers, A guided tour through excursions, *Bull. Lond. Math. Soc.* 21 (1989) 305–341.
- [32] L.C.G. Rogers, Monte Carlo valuation of American options, *Math. Finance* 12 (2002) 271–286.
- [33] L.C.G. Rogers, Bermudan options by simulation. Technical Report, University of Cambridge, 2015.
- [34] L.C.G. Rogers, D. Williams, *Diffusions, Markov Processes and Martingales: Volume 2, Itô Calculus*, Vol. 2, Cambridge University Press, 2000.
- [35] C. Yoeurp, Grossissement progressif d’une filtration, grossissement de filtrations: exemples et applications, in: *Lecture Notes in Mathematics*, Vol. 1118, Springer-Verlag, 1985, pp. 197–315.