

The value of foresight

Philip Ernst

Department of Statistics, Rice University

Support from NSF-DMS-1811936 (co-PI F. Viens) and ONR-N00014-18-1-2192 gratefully acknowledged.

IMA Financial and Economic Applications

June 11, 2018

Ernst, P.A., Rogers, L.C.G. and Zhou, Q. (2017)

The value of foresight

Stochastic Processes and their Applications **129**: 3913-3927.

Motivation I

Value of foresight

- Larry Shepp first asked the following problem: Consider a discrete-time Markov process for a stock. What is the optimal stopping strategy of an insider trader that can see m steps into the future?
- For the continuous-time model, the analogous question is: what would the trader gain who can see a units of time into the future?
- In the real world, market participants receive market information after a delay, and thus those who receive the information without delay have the same kind of advantage as insider traders.
 - Typical example: high frequency trading (500-microsecond advantage)


Motivation II

Look-back option

- We will see the value of foresight can be equivalently interpreted as the value of an American fixed-window look-back option: at a time of your choice, you can sell the stock for the best price in the past a units of time.
- The larger the value of a , the more you need to pay for the option
- Together with Rogers^a, we are the first to discuss this option.

^aRogers, L. C. G. (2015) Bermudan options by simulation. *Technical Report, University of Cambridge*.

The discretized path of stock price. The maximum stock price in $[0.6, 0.8]$ is 2.148.



Mathematical challenges

- Like many other option pricing problems, there is no closed-form solution to this problem.
 - For the insider trader, the decision whether or not to stop depends on the entire path of the stock price in the next a units of time – a path-valued state variable (infinite-dimensional optimization).
 - Even after discretizing time, the problem is high-dimensional.
 - The optimization is with respect to a non-stopping time since one makes decisions using information about the future.
 - The current literature on filtration enlargement is not applicable.
- Our goal is to find explicit expressions for
 - the bound of the maximum expected reward
 - “close-to-optimal” exercise rules
- We take a completely different approach from the existing methods.
 - Core: Brownian excursion theory

Mathematical challenges

Existing enlargement of filtration theory is not applicable

- The advantage of insider trading is usually modelled using the theory of enlargement of filtration (see Aksamit and Jeanblanc, 2017). We use (\mathcal{F}_t) to denote the usual filtration and (\mathcal{F}'_t) to denote the enlarged filtration.
- However, results are only established for particular classes of enlargement such that any (\mathcal{F}_t) -local martingale is a (\mathcal{F}'_t) -semimartingale.
- For our problem, the information available to the insider trader is modelled by $(\mathcal{F}'_t) = (\mathcal{F}_{t+a})$.
- One may show that a (\mathcal{F}_t) -Wiener process is not a (\mathcal{F}'_t) -semimartingale.

How so?

Proposition

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

PROOF. Consider the simple integrands

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

where

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

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

Proof (continued)

Now consider the (elementary) stochastic integral

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

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 Bichteler-Dellacherie theorem 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

And so...

The problem addressed is *concrete, challenging, and not amenable to general theory* — which is why it appealed to Larry Shepp!

Model of the stock price

- The stock price S_t is modelled by geometric Brownian motion,

$$S_t = \exp(\sigma W_t + (r - \sigma^2/2)t) \stackrel{\text{def}}{=} \exp(X_t), \quad S_0 = 1,$$

where W_t is the Wiener process, r is the fixed riskless interest rate and σ is the volatility. X_t is Brownian motion with drift.

- Without loss of generality, let $r = 0$ and $\sigma = 1$.
 - We are only interested in the discounted stock price $\tilde{S}_t = e^{-rt} S_t$ and thus r is always cancelled out.
 - Changing σ is equivalent to rescaling all the time parameters.
- Our model thus simplifies to

$$S_t = \exp(X_t), \quad X_t = W_t + ct \quad (c = -1/2).$$

Value of 'look-ahead'

- Consider a trader that can look into the future by a units of time. How much can the trader profit from the foresight?
- The insider trader can stop at any non-negative random time τ' s.t.

$$\tau' + a \text{ is a stopping time} \iff \{\tau' \leq t\} \in \mathcal{F}_{t+a}.$$

- The insider trader can stop at τ' , although τ' is not a stopping time
- Since $\mathcal{F}_t \subseteq \mathcal{F}_{t+a}$, the insider trader gains advantage over non-insiders.
- Let $\tau = \tau' + a$. Our primary goal is to find

$$v(a) \stackrel{\text{def}}{=} \sup_{0 \leq \tau' \leq T} E[S_{\tau'}] = \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}]$$

where T denotes an expiration time.

- This is the "look-ahead" interpretation of the problem.

Value of 'look-back'

- Another equivalent formulation of $v(a)$ (by the proposition to come) is

$$v(a) = \sup_{0 \leq \tau \leq T} E[Z_\tau], \quad Z_t \stackrel{\text{def}}{=} \sup\{S_u : t - a \leq u \leq t\}$$

where τ is a stopping time. So $v(a)$ is the value of an American fixed-window look-back option.

- This formulation also shows why the problem will never be solved explicitly. Plainly, it is because Z_t is *not Markov*, and *cannot be made Markov*. One would have to take as the state the path fragment $(S_u)_{t-a \leq u \leq t}$.

Proposition

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

Proposition

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

$$v(a) \equiv \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}] = \sup_{0 \leq \tau \leq T} E[Z_\tau], \quad (4)$$

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

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

$$v(a) \equiv \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}] \leq \sup_{0 \leq \tau \leq T} E[Z_\tau]. \quad (5)$$

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

$$\tilde{\tau} = \inf\{u \geq \tau \vee a : S_{u-a} = Z_\tau\}. \quad (6)$$

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

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

Proof (Continued)

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

$$Z_\tau = S_{\tilde{\tau}-a}, \quad (7)$$

and therefore

$$E[Z_\tau] = E[S_{\tilde{\tau}-a}] \leq \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}], \quad (8)$$


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

$$\sup_{0 \leq \tau \leq T} E[Z_\tau] \leq \sup_{a \leq \tau \leq T+a} E[S_{\tau-a}], \quad (9)$$

and the proof is complete.


□

'Look-ahead'



For the 'look-ahead' trader, at time 0.15, the trader can foresee that the highest price in the next 0.2 units of time will occur at time 0.259. Therefore the trader will continue till at least time 0.259 (.109 time units ahead).

'Look-back'



This is exactly the same situation for a look-back option holder at time 0.35. If the holder sells the stock, the reward would be the stock price at time 0.259. The holder will continue since at time 0.459 (also .109 time units ahead) the stock can still be sold for that price.

Interpretation


- If $r \neq 0$, these two interpretations may be slightly different.
- The insider trader gets the look-back option for free.
- We will work with the model of the look-back option.
- Although 'look-ahead' and 'look-back' are mathematically equivalent, we will work with the look-back option since it is a bit more intuitive and it is more amenable for simulation study.

Stopping time τ_0

- Our theory for the value of look-back option starts with the observation:

the first time we might stop before T is $\tau_0 \stackrel{\text{def}}{=} \inf\{t : Z_t = S_{t-a}\}$.

- At τ_0 , we have $Z_{\tau_0} = S_{\tau_0-a} = \max_{0 \leq u \leq \tau_0} S_u$.

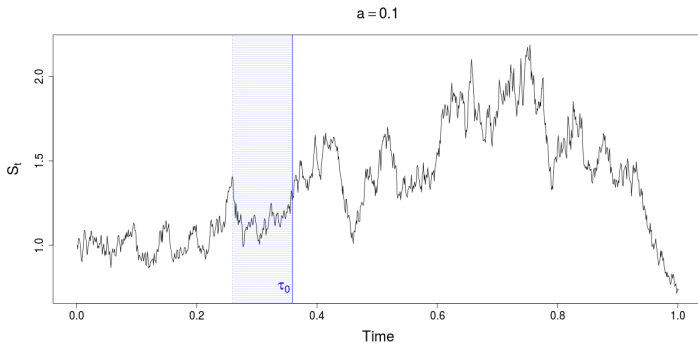


The value of foresight

Stopping time τ_0

τ_0 is not the optimal stopping time. We may want to continue because

- There is still a lot of remaining time.
- If we continue, the risk we need take may be very small.



An approximate model

The optimal stopping rule would be a function of the whole path in the fixed window, for which we do not have a closed-form expression. However, we can derive the closed-form expressions for both the optimal reward and the optimal exercise rule under the following (modified) model.

- The modified model makes two changes to the original model.
 - There is no fixed expiration time T . Instead, we assume the killing of the process comes at an exponential rate and denote it by $\alpha \sim Exp(\eta)$.
 - At a potential stopping time like τ_0 , if we decide to continue, we forget the previous path so that the whole process restarts.
- The exponential killing is memoryless, so the exercise rule at τ_0 should not depend on time.
- The second 'forget and continue' assumption says that the stopping rule at τ_0 should only depend on $S_{\tau_0-a}/S_{\tau_0} = \exp(X_{\tau_0-a} - X_{\tau_0})$, i.e., the endpoints of the past window.

Optimal exercise rule of the approximate model

Let K^* be the maximum expected reward of the approximate model with $S_0 = 1$. The optimal stopping rule must be

- Wait until $\tau_0 \wedge \alpha$.
- If $\alpha < \tau_0$, stop and receive Z_α since the process is killed.
- If $\tau_0 < \alpha$,
 - (a) if we stop, we receive $Z_{\tau_0} = S_{\tau_0-a}$;
 - (b) if we continue, we expect to receive $K^* S_{\tau_0}$.Hence, we should only stop if $S_{\tau_0-a} > K^* S_{\tau_0}$, which is equivalent to

$$X_{\tau_0} - X_{\tau_0-a} < q^* \stackrel{\text{def}}{=} -\log K^*.$$

- If $\tau_0 < \alpha$ and we have decided to continue at τ_0 , everything restarts.

Exercise rule under the exponential killing

To find q^* , consider the following class of exercise rules with parameter q .

- Wait until $\tau_0 \wedge \alpha$.
- If $\alpha < \tau_0$, stop and receive Z_α .
- If $\tau_0 < \alpha$,
 - (a) if $X_{\tau_0} - X_{\tau_0-a} < q$, stop and receive $Z_{\tau_0} = S_{\tau_0-a}$.
 - (b) if $X_{\tau_0} - X_{\tau_0-a} > q$, forget and continue.
- If $\tau_0 < \alpha$ and we have decided to continue at τ_0 , everything restarts and of course these rules still apply.

Optimal exercise rule

- Let $K(q)$ be the expected reward of the exercise rule using q .
- The optimal q^* must satisfy

$$\frac{\partial K}{\partial q} \Big|_{q^*} = 0, \quad \frac{\partial^2 K}{\partial q^2} \Big|_{q^*} \leq 0.$$

- q^* can be obtained by simple numerical optimization methods.
- We will use excursion theory to calculate $K(q)$.

Excursion theory

- Define running maximum $\bar{X}_t = \sup_{0 \leq u \leq t} X_u$ and consider the process

$$Y_t \stackrel{\text{def}}{=} X_t - \bar{X}_t, \quad 0 \leq t \leq \tau_0.$$

- By Levy's Theorem, (Y_t, \bar{X}_t) has the same distribution as $(-|X_t|, L_t(0))$ where $|X_t|$ is the reflected Brownian motion and $L_t(0)$ is the local time at 0 of X_t (also $|X_t|$).
- The path of Y_t is decomposed to a Poisson point process, on the local time $L_t(0)$, of excursions away from 0.
- We use n to denote the excursion (rate) measure on the excursion space

$$E = \{\text{continuous } f : \mathbb{R}^+ \rightarrow \mathbb{R} \text{ s.t. } f^{-1}(\mathbb{R} \setminus \{0\}) = (0, \zeta), \zeta > 0\},$$

where ζ is called the (real) lifetime of the excursion.

- τ_0 happens when Y_t makes first excursion of which $\zeta > a$.

Rules for the original problem

Let $q^*(\eta)$ be the optimal value for q given the exponential killing rate η . How to use $q^*(\eta)$ to define exercise rules for the American fixed-window look-back option?

Rule

- Wait until τ_0 .
- If $X_{\tau_0} - X_{\tau_0-a} > q'$, continue till next τ_0 .

Choice of q'

- Rule 1: use $q' = q^* \left(\frac{1}{T - \tau_0} \right)$.
- Rule 2: use $q' = q^* \left(\frac{1}{T - \tau_0} \right) - q^* \left(\frac{1}{a} \right) - \log \lambda(a)$ (to be explained).
- Both rules are explicit and deterministic.

Rule 2

- If at τ_0 we happen to have $\tau_0 = T - a$, the optimal rule by assuming “forget and continue” would simply be:

continue if $X_{\tau_0} - X_{\tau_0-a} > -\log \lambda(a)$

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

- However, by Rule 1, we would:

continue if $X_{\tau_0} - X_{\tau_0-a} > q^*(1/a)$.

- This motivates us to shift the q of Rule 1 by a constant and at any τ_0 ,

$$\text{continue if } X_{\tau_0} - X_{\tau_0-a} > q^* \left(\frac{1}{T - \tau_0} \right) - q^* \left(\frac{1}{a} \right) - \log \lambda(a).$$

Discretization of time

- In simulation, we have to discretize time. Consequently, we are actually looking for the solutions to the discrete-time version of the problem. (h : the length of each time step)

$$v_h(a) \stackrel{\text{def}}{=} \sup_{0 \leq \tau_h \leq T} E[Z_{\tau_h}^{(h)}], \quad \tau_h \in h\mathbb{N}$$

$$Z_t^{(h)} \stackrel{\text{def}}{=} \max\{S_{kh} : t - a \leq kh \leq t\}$$

- Let $m = a/h$. The $(m + 1)$ -dimensional process

$$U_t^{(h)} \stackrel{\text{def}}{=} (S_{t-mh}, \dots, S_t), \quad t \in h\mathbb{N}$$

is Markovian. Moreover, $(U_t^{(h)}, Z_t^{(h)})$ is a Markov process too.

- In simulation, we only need to keep record of $m + 1$ stock prices.
- The superscript h will be dropped.

Bounds of $v(a)$ by simulation

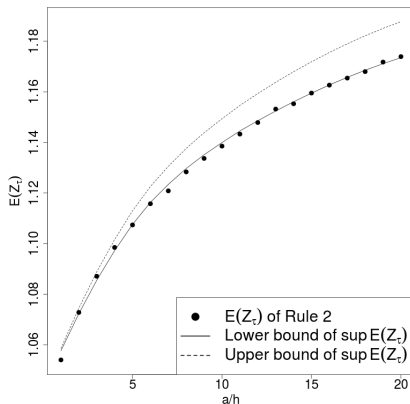

- We use the simulation method proposed by Rogers^a to calculate the bounds of $v(a)$.
- In the simulation, we discretize time with step length $h = 1/2500$ and use termination time $T = 0.1$. Hence there are 250 steps.

a/h	Lower	SE($\times 10^{-4}$)	Upper	SE ($\times 10^{-4}$)
1	1.054	0.43	1.055	0.53
2	1.074	0.62	1.076	0.77
3	1.088	0.73	1.092	0.90
4	1.100	0.83	1.104	1.0
5	1.109	0.95	1.114	1.1
6	1.117	1.0	1.123	1.2
7	1.123	1.1	1.131	1.2
8	1.129	1.2	1.137	1.3
9	1.135	1.3	1.144	1.4
10	1.140	1.4	1.149	1.4
15	1.159	1.7	1.172	1.8
20	1.174	2.0	1.188	1.9

^aRogers, L. C. G. (2015) Bermuda options by simulation. *Technical Report, University of Cambridge*.

Simulation of Rule 1 and Rule 2

- For each choice of a , we average the reward over 50,000 sample paths.
- The lines are the bounds for $v(a)$ computed by simulation.



Simulation of Rule 1 and Rule 2

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

The standard error for the expected reward of our rules is approximately 0.001 in all cases.

Concluding remarks

- Rule 2 is slightly better than Rule 1 as expected.
- The dots of Rule 2 coincide with the lower bound of $v(a)$, which implies it could be very useful in practice and replace the lower bound.
- We never know the true value for $v(a)$, which could be just slightly greater than the lower bound.
- Our rules are clearly motivated, precisely specified and easy to calculate, being different from randomly-generated rules.
- Using $q^*(\eta)$, one can search for other rules, deterministic or stochastic, that have better performance. But that goes beyond the scope of this study.