UPPSALA
UNIVERSITET

U.U.D.M. Project Report 2018:33

Optimal Stopping with Discrete Costly Observations

Yuqiong Wang

Examensarbete i matematik, 30 hp
Handledare: Erik Ekström
Examinator: Kaj Nyström
Juli 2018

Department of Mathematics
Uppsala University

Contents

Abstract

We study the optimal stopping problem where one can only observe the underlying process at discrete random time points, and each of these observations comes with a constant cost. One should decide how to distribute the future observation times and when to stop the process.

To solve this problem, we define an associated operator and prove that its unique fixed point characterizes the value function. We then provide an optimal strategy where one chooses how to distribute the future observations and when to stop in terms of the fixed point. We then use an iterative procedure to reach the fixed point, and provide a specific example of this procedure which has an exponential convergence rate.

We conclude that when we can make at most finitely many observations, the value functions can be characterized by the elements in the sequence constructed by the iterative procedure. We further prove that when one can stop the process at any time, it reduces to the previous problem.

Finally, we discuss the quickest detection problem under discrete costly observations, and we see it can also be solved by our approach.

Acknowledgements

I would like to give my deepest thanks to my supervisor Professor Erik Ekström for introducing me to this exciting topic, and for his invaluable guidance, stimulating discussions, constant help and encouragement whenever I needed it.

I gratefully acknowledge the support and generosity of my sponsor Svenska institutet.

Chapter 1

Introduction

In the general optimal stopping problems we often consider choosing a stopping time in order to maximize the expected gain or minimize the expected loss. In this study we consider such a case where we can only get information from making discrete observations with costs, and we decide two things: when to observe and when to stop. The question then arises: How do we maximize our expected gain in this case?

There are some discussions on similar topics. In their book [1] Peskir and Shiryaev discuss a similar case where the observation times are deterministic, and one pays for the information in the future. In our case we pay for the information today and we observe immediately, and the observation times are random. In their article [3], Bayraktar and Kravitz discuss a quickest detection problem under discrete observations with cost. This can be regarded as an example in our problem. In their article [2], Dyrssen and Ekström solve the sequential testing problem under discrete observations with cost. They show that the value function equals the unique fixed point of an associated operator, which can be constructed using an iterative procedure. We follow their approach when solving our problem but prove the uniqueness of the fixed point in a different way.

In Chapter 2 we carefully answer the previous question. Firstly, in Section 2.1 we give a precise formulation of the problem, and define what information we can use and when we can observe and stop. Our main result is obtained in Section 2.2, where we define an operator which characterizes our choices at each observation. We assume that there exists a fixed point of this operator. The main result of this project is Theorem 2.2.3, where we conclude this fixed point equals our value function, and provide an optimal strategy, using Lemma 2.2.1 and Lemma 2.2.2. The uniqueness of the fixed point thus follows, assuming that it exists. The next step is to prove that such a fixed point indeed exists, and provide a method to find it. In Section 2.3 we talk about this: we use an iterative procedure to construct a sequence and prove that the point-wise limit of this sequence is the fixed point. We then discuss a specific example in detail, where the gain function equals the payoff of a put option. We discuss the structure of the continuation set in terms of the choice of parameters in this example, and prove that starting iterating from the lower bound, this sequence has the contraction property and an exponential rate of convergence. We provide some numerical examples which agree with our conclusions.

In Section 2.4 we follow the approach in [2], and conclude that in the sequence constructed by iterating from the lower bound, every element can be interpreted as the value function when one can make at most finitely many observations. It is natural to ask what happens if if we were allowed to stop the process continuously when observing the underlying process discretely, in Section 2.5 we show that it reduces to the problem we solve in Section 2.2. We do this by using the strong Markov property and constructing two optimal stopping times.

In Chapter 3 we discuss the quickest detection problem under discrete costly observations and the property of its underlying process, and show that it can also be solved by our approach.

1.1 An Overview of Optimal Stopping Theory

Suppose that we have a filtered probability space $(\Omega, \mathcal{F}_t, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$, and an arbitrary stochastic process $G = (G_t)_{t\geq 0}$ defined on it, where G_t is interpreted as the gain if observation is stopped at time t. We interpret $\mathcal{F}_t$ as the information available up to time t, and we ask G_t is $\mathcal{F}_t$-measurable. We can define the value function as:

$$V = \sup_{\tau \in \mathfrak{m}} \mathbb{E} G_\tau$$

where $\mathfrak{m}$ is the collection of the random variables τ such that $\mathbb{P}(\tau(\omega) < \infty) = 1$. The fact that G_t is adapted to $\mathcal{F}_t$ makes each τ a stopping time.

Here we would also assume that $G_\infty = 0$, and that [1]:

$$\mathbb{E}(\sup_{t\geq 0} |G_t|) < \infty$$

The optimal stopping problem involves two goals:

(i) Find the value function V;

(ii) Find an associated stopping time τ where we should stop the underlying process.

For general optimal stopping problems, we have two main approaches: martingale approach and Markovian approach. When choosing which method to use, one should consider the probabilistic structure of the stochastic processes which underly the problem. The two structures can be considered as "unconditional" and "conditional", respectively [1].

When we determine the probabilisitic structure of a process by its unconditional finite dimensional distribution, we refer to the martingale approach. As the techniques and theories to the solution of such a problem usually based on the results from the theory of martingales. In the martingale approach we can use backwards induction for discrete time finite horizon cases, and essential supremum for discrete and continuous cases, with finite or infinite horizon. For details, see Chapter I in [1].

When we determine the structure of the process by considering a (conditional) transition functions and its initial state, instead of directly from its distributions, we refer to the Markovian approach. As one can use the powerful tools from the theory of Markov processes. In the Markovian approach we assume that $G_t(\omega)$ has Markovian representation, which means there exists a Markov process X_t such that $G_t(\omega) = G(t, X_t(\omega))$.

One can choose one of the two methods depending on the features of the underlying process of the problem. We can think of a Markov process as a special case of the processes determined by its unconditional distribution, or, if we take the state space to be large enough, we can obtain a Markovian representation for any process, for some measurable function G. For details of the two methods, see [1]. In this project we use the idea of Markovian approach in discrete time, which we will give a brief introduction in Section 1.2.

1.2 Discrete Time: Markovian Approach

In this section let us recall the standard Markovian approach when the time is discrete. Let us consider a time-homogeneous Markov chain $X = (X_n)_{n\geq 0}$ defined on a filtered probability space $(\Omega, \mathcal{F}_t, (\mathcal{F}_n)_{n\geq 0}, \mathbb{P})$ and taking values in the measurable space $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. X starts from $x \in \mathbb{R}^d$ under $\mathbb{P}$, and that $x \mapsto \mathbb{E}_x(Z)$ is measurable for each Z.

Let us assume we have a measurable function $G : \mathbb{R}^d \to \mathbb{R}$ which satisfies:

$$\mathbb{E}_x(\sup_n |G(X_n)|) < \infty$$

Let us consider the finite horizon first, where we consider the value function:

$$V^N(x) = \sup_{0\leq\tau\leq N} \mathbb{E}_x G(X_\tau)$$

where τ is a stopping time with respect to the natural filtration $\mathcal{F}_n^X = \sigma(X_k : 0 \leq k \leq n)$. Let us introduce a sequence of random variables $(S_n^N)_{0\leq n\leq N}$ defined recursively as:

$$\begin{aligned} S_n^N &= G(X_N), \quad \text{when } n = N, \\ S_n^N &= \max(G(X_n), \mathbb{E}(S_{n+1}|\mathcal{F}_n)), \quad \text{for } n = N-1, \ldots, 0. \end{aligned}$$

and the key identity to this problem is:

$$S_n^N = V^{N-n}(X_n)$$

To prove this identity we first prove by induction that:

$$S_n^N = \mathbb{E}_x(G(X_{\tau_n^N})|\mathcal{F}_n)$$

where:

$$\tau_n^N = \inf\{n \leq k \leq N : S_k^N = G(X_k)\}$$

Then by the Markov property we have:

$$\begin{aligned}
S_n^N &= \mathbb{E}_x(G(X_{\tau_n^N})|\mathcal{F}_n) \\
&= \mathbb{E}_x(G(X_{n+\tau_0^{N-n}\circ\theta_n})|\mathcal{F}_n) \\
&= \mathbb{E}_x(G(X_{\tau_0^{N-n}})\circ\theta_n|\mathcal{F}_n) \\
&= \mathbb{E}_{X_n}G(X_{\tau_0^{N-n}}) \\
&= V^{N-n}(X_n)
\end{aligned}$$

Let us now let:

$$\begin{aligned}
C_n &= \{x : V^{N-n}(x) > G(x)\} \\
D_n &= \{x : V^{N-n}(x) = G(x)\}
\end{aligned}$$

Let us define:

$$\tau_D = \inf\{0 \leq n \leq N : X_n \in D_n\}$$

And the operator T:

$$TF(x) = \mathbb{E}_x[F(X_1)]$$

where F is measurable and $F(X_1)$ is integrable with respect to $\mathbb{P}$.

Then we have the following conclusions [1]:

(i) The value function V^n satisfies:

$$V^n(x) = \max(G(x), TV^{n-1}(x))$$

for $n = 1, \ldots, N$, where $V^0 = G$.

(ii) The stopping time τ_D is optimal.

(iii) If τ^* is an optimal stopping time, then $\tau_D \leq \tau^*$ $\mathbb{P}$-a.s. for all x.

(iv) The sequence $(V^{N-n}(X_n))_{\{0\leq n\leq N\}}$ is the smallest supermartingale which dominates $(G(X_n))_{\{0\leq n\leq N\}}$ under $\mathbb{P}$ for x fixed.

(v) The stopped sequence $(V^{N-n\wedge\tau_D}(X_{n\wedge\tau_D}))_{\{0\leq n\leq N\}}$ is a martingale under $\mathbb{P}$ for every x.

The proof of (i) follows from the Markov property:

$$\begin{aligned}
V^{N-n}(X_n) &= \max(G(X_n), \mathbb{E}_x(V^{N-n-1}(X_{n+1})|\mathcal{F}_n)) \\
&= \max(G(X_n), \mathbb{E}_x(V^{N-n-1}(X_1)\circ\theta_n|\mathcal{F}_n)) \\
&= \max(G(X_n), \mathbb{E}_{X_n}(V^{N-n-1}(X_1)|\mathcal{F}_n)) \\
&= \max(G(X_n), TV^{N-n-1}(X_n)
\end{aligned}$$

By letting $n = 0$ we get (i).

To prove (ii), first one needs to prove $S_n^N \geq \mathbb{E}_x(X_\tau|\mathcal{F}_n)$ for all stopping times $n \leq \tau \leq N$. Taking expectation and supremum on both sides, we have $\mathbb{E}_x(S_n^N) \geq V_n^N$, we also have $\mathbb{E}_x(S_n^N) = \mathbb{E}_x(G(X_{\tau_D})) \leq V_n^N$, which implies the equality of (ii).

To prove (iii), one needs to prove that τ^* being an optimal stopping time implies $S_{\tau^*}^N = G(X_{\tau^*})$ $\mathbb{P}$-a.s.. If we assume the contrary, then we get the strict equality $\mathbb{E}_x G(X_{\tau^*}) < V_n^N$ by optional sampling theorem, which contradicts the fact that τ^* is optimal.

In the proof of (iv), the supermartingality follows directly from $S_n^N = \max(G(X_n), \mathbb{E}(S_{n+1}|\mathcal{F}_n))$. Suppose there is another supermartingale $(\tilde{V}^{N-n}(X_n))_{\{0\leq n\leq N\}}$ which dominates $(G(X_n))_{\{0\leq n\leq N\}}$, one can prove by induction that $(\tilde{V}^{N-n}(X_n))_{\{0\leq n\leq N\}} \geq (V^{N-n}(X_n))_{\{0\leq n\leq N\}}$ almost surely.

To prove (v), fix k and use an indicator. To see the details of theses proofs, see [1].

If we introduce an operator Q defined by:

$$QF(x) = \max(G(x), TF(x))$$

then we can write the value function $V^n(x)$ as:

$$V^n(x) = Q^n G(x)$$

This recursive relation forms a constructive method for finding $V^N(x)$.

Now let us consider the infinite horizon where we write the value function as:

$$V(x) = \sup_\tau \mathbb{E}_x G(X_\tau)$$

Similarly, we define the continuation set and the stopping set:

$$C = \{x : V(x) > G(x)\}$$
$$D = \{x : V(x) = G(x)\}$$

Then we have the following conclusions [1]:

(i) The value function V satisfies:

$$V(x) = \max(G(x), TV(x))$$

where the operator T is defined as:

$$TF(x) = \mathbb{E}_x F(X_1)$$

(ii) Assume that $\mathbb{P}_x(\tau_D < \infty) = 1$, where:

$$\tau_D = \inf\{t \geq 0 : X_t \in D\}$$

then the stopping time τ_D is optimal.

(iii) The value function V is the smallest superharmonic function (Dynkin's Characterization) $(TV \leq V)$ which dominates the gain function G.

(iv) The stopped sequence $V(X_{n\wedge\tau_D})$ is a martingale.

(v) We have a constructive method for finding the value function:

$$V(x) = \lim_{n \to \infty} Q^n G(x)$$

The key identity here is $S_n = V(X_n)$ which can be proved by taking $N \to \infty$. The rest of the claims follow from the previous claims where N is finite. To see more detailed explanations and more examples, one can refer to [1].

Chapter 2

Stopping Problem with Discrete Costly Observations

2.1 Formulation of the Problem

Let us assume that we want to solve an optimal stopping problem where the underlying stochastic process $(X_t)_{t\geq 0}$ is a (strong) Markov process. We can choose to stop this process at some time point, and then get a payoff which depends on the value of the process at the time when we stop. However, instead of observing the underlying process continuously, we can only observe it at discrete time points, and only stop the process at those observation times. Let us assume that we can choose the sequence of time points at which we observe. However, the observations are not for free, but every observation comes with a cost, which is a constant in our case. Therefore, we cannot make infinitely many observations to get the most sufficient information, as we need to make as few observations as possible, so as to determine when to stop the process. Is there an optimal strategy based on the current value of our underlying process, such that we choose a certain sequence of observing times, as well as a certain time point that we stop the process and exercise the option, such that we could achieve the maximum discounted payoff at time 0? Our goal is to find such a strategy to optimize our value function accordingly.

Let us first define the observation times. Let us take an infinite sequence of random times: $\tau_1 \leq \tau_2 \leq \tau_3 \leq \dots$, such that: τ_i is $\sigma\{X_{\tau_1}, X_{\tau_2}, \dots, X_{\tau_{i-1}}, \tau_1, \tau_2, \dots, \tau_{i-1}\}$ measurable. We can see that τ_1 is deterministic, as it is measurable with respect to $\mathcal{F}_0^{\hat{\tau}}$.

Let us write the sequence $\hat{\tau} = \{\tau_k\}_{k\geq 1}$, and define the information associated with this infinite sequence $\hat{\tau}$ up to time t as:

$$(\mathcal{G}_t^{\hat{\tau}})_{t\geq 0} = \sigma\{X_{\tau_1}, X_{\tau_2}, \dots, X_{\tau_i}, \tau_1, \tau_2, \dots, \tau_i;\ \ i = \max\{j : \tau_j \leq t\}\}$$

We can see from the definition that this continuous time filtration $(\mathcal{G}_t^{\hat{\tau}})_{t\geq 0}$ satisfies that, $\forall j \geq 1$, if $\tau_j \leq t < \tau_{j+1}$, then $\mathcal{G}_t^{\hat{\tau}} = \mathcal{G}_{\tau_j}^{\hat{\tau}}$. In other words, the information up to time t comes from all the observations before it.

Let us now define two sets of stopping times admissible to the filtration $\mathcal{G}^{\hat{\tau}}$:

Definition 2.1.1. *Define:*

$$\mathcal{T}^{\hat{\tau}} = \{\mathcal{G}^{\hat{\tau}} \text{ - stopping times}\},$$
$$\text{and} \quad \mathcal{S}^{\hat{\tau}} = \{\tau \in \mathcal{T}^{\hat{\tau}} : \mathbb{P}(\exists i, \tau(\omega) = \tau_i(\omega)) = 1\}$$

Remark 2.1.1. *The set $\mathcal{T}^{\hat{\tau}}$ is the collection of all $\mathcal{G}^{\hat{\tau}}$-stopping times. The set $\mathcal{S}^{\hat{\tau}}$ is the collection of $\mathcal{G}^{\hat{\tau}}$-stopping times taking values in the sequence $\hat{\tau}$. Obviously $\mathcal{S}^{\hat{\tau}} \subset \mathcal{T}^{\hat{\tau}}$.*

Let us denote the payoff function by g. If we choose to stop the process at time t, the payoff we get at the moment we stop is $g(X_t)$. Let g be a real-valued non-negative function bounded by $M < \infty$: $g : \mathbb{R}_+ \to \mathbb{R}_+, g \leq M$. Let V be our value function, which optimizes our discounted payoff at time 0:

Definition 2.1.2. *Value Function V*

$$V : [0, +\infty) \to [0, +\infty)$$
$$V(x) = \sup_{\hat{\tau}} \sup_{\tau \in \mathcal{S}^{\hat{\tau}}} \mathbb{E}[e^{-r\tau} g(X_\tau) - \sum_{i=1}^{\infty} c e^{-r\tau_i} 1_{\{\tau_i \leq \tau\}}]$$

Where $r \geq 0$ is the discounting rate, $c > 0$ is the cost for each observation.

Remark 2.1.2. *Here we do not ask the stopping times to be finite, in other words, τ can take the value ∞ with positive probability. Let us assume that we have made n observations before τ, and the optimal strategy is to never stop the process, i.e. $\tau = \tau_{n+1} = \infty$. However, we will not make another observation at ∞. Therefore when looking at the value function, if we assume one must first pay the cost and then stop the process, it is sufficient to consider all the finite stopping times.*

Remark 2.1.3. *From Definition 2.1.2 we can see that, by taking $\tau = 0$, V is bounded by $g(x)$ from below. It is also bounded above by M from above.*

2.2 A Fixed Point Approach

In our formulation of the problem we only allow stopping the process at the observation times. i.e. we can either exercise the option immediately after making an observation, or continue observing for a deterministic time and make another observation. Let us assume that we have an optimal strategy where we can optimize the value function. Then standing on each observation point of this optimal sequence, our next step should always be optimal. Note that this optimal strategy might not be unique, but such a sequence would provide an optimal strategy. Therefore, it is natural to define an operator associated with this problem, so as to characterize our options at each step. We will now specify our choice, define the operator and solve the problem step by step.

Standing on each observation point, after making an observation, we take a look at the information we have up to now, and decide whether to continue observing the underlying process for once more or not. We have 2 choices:

(i) To decide to stop observing and stop the process immediately;

(ii) To decide to make another observation and face the same choices after making the next observation.

We now characterize these choices. Let us first define a family F of functions:

Definition 2.2.1. *Set F*

$$F := \{f \text{ Borel-measurable} : [0, +\infty) \to [0, +\infty), \ \ g \leq f \leq M\}$$

Let us now introduce an operator $\mathcal{J}$ to characterize our choices as described above:

Definition 2.2.2. *Operator $\mathcal{J}$*
Let $\mathcal{J}$ be an operator acting on F, defined by:

$$(\mathcal{J}f)(x) = \max(g(x), \ \sup_{t \geq 0} \mathbb{E}_x[e^{-rt}(f(X_t) - c)])$$

Let us assume that $\hat{V}$ is a fixed point of the operator $\mathcal{J}$ (We will prove the existence in Section 2.3.1), i.e. $\hat{V} = \mathcal{J}\hat{V}$. We can thus write:

$$\begin{aligned} &\hat{V} : [0, +\infty) \to [0, +\infty) \\ &\hat{V}(x) = \max(g(x), \ \sup_{t \geq 0} \mathbb{E}_x[e^{-rt}(\hat{V}(X_t) - c)]) \end{aligned}$$

Now we will prove that $\hat{V} = V$, where V is defined as in Definition 2.1.2, and then provide an optimal strategy. It thus follows that the fixed point, if it exists, is unique.

Let us define a discrete-time filtration $\mathcal{F}^{\hat{\tau}}$ indexed by k, by setting $\mathcal{F}^{\hat{\tau}}_k = \mathcal{G}^{\hat{\tau}}_{\tau_k}$.

Lemma 2.2.1. *Let $\hat{\tau}$ be any sequence of stopping times, then $k \to e^{-r\tau_k}\hat{V}(X_{\tau_k}) - c\sum_{1 \leq i \leq k} e^{-r\tau_i}$ is a supermartingale with respect to $\mathcal{F}^{\hat{\tau}}_k$.*

Proof. We have $\hat{V} = \mathcal{J}\hat{V}$, so:

$$\hat{V}(x) = \max(g(x), \ \sup_{t \geq 0} \mathbb{E}_x[e^{-rt}(\hat{V}(X_t) - c)])$$

Define $Y_k := e^{-r\tau_k}\hat{V}(X_{\tau_k}) - c\sum_{1 \leq i \leq k} e^{-r\tau_i}$, for $k = 0$:

$$\begin{aligned} \mathbb{E}_x[Y_1] &= \mathbb{E}_x[e^{-r\tau_1}\hat{V}(X_{\tau_1}) - ce^{-r\tau_1}] \\ &\leq \sup_t \mathbb{E}_x[e^{-rt}(\hat{V}(X_t) - c)] \\ &\leq \hat{V}(x) = Y_0 \end{aligned} \tag{2.1}$$

For $k = 1, 2, \ldots$:

$$\begin{aligned}
\mathbb{E}_x[Y_{k+1}|\mathcal{F}_k^{\hat{\tau}}] &= \mathbb{E}_x[e^{-r\tau_{k+1}}\hat{V}(X_{\tau_{k+1}}) - c\sum_{1\le i\le k+1} e^{-r\tau_i}|\mathcal{G}_{\tau_k}^{\hat{\tau}}] \\
&= \mathbb{E}_x[e^{-r\tau_{k+1}}(\hat{V}(X_{\tau_{k+1}}) - c)|\mathcal{F}_k^{\hat{\tau}}] - c\sum_{1\le i\le k} e^{-r\tau_i} \\
&= \mathbb{E}_x[e^{-r(\tau_k+(\tau_{k+1}-\tau_k))}(\hat{V}(X_{\tau_{k+1}}) - c)|\mathcal{F}_k^{\hat{\tau}}] - c\sum_{1\le i\le k} e^{-r\tau_i} \\
&= e^{-r\tau_k}\mathbb{E}_x[e^{-r(\tau_{k+1}-\tau_k)}(\hat{V}(X_{\tau_{k+1}}) - c)|\mathcal{F}_k^{\hat{\tau}}] - c\sum_{1\le i\le k} e^{-r\tau_i}, \quad \text{since } \tau_k \in \mathcal{G}_{\tau_k}^{\hat{\tau}} \\
&= e^{-r\tau_k}\mathbb{E}_{X_{\tau_k}}[e^{-r(\tau_{k+1}-\tau_k)}(\hat{V}(X_{\tau_1}) - c)] - c\sum_{1\le i\le k} e^{-r\tau_i}, \quad \text{by strong Markov property} \\
&\le e^{-r\tau_k}\hat{V}(X_{\tau_k}) - c\sum_{1\le i\le k} e^{-r\tau_i}, \quad \text{by (2.1)} \\
&= Y_k
\end{aligned}$$

Therefore we conclude that $\{Y_k\}_{k=0}^{\infty}$ is a supermartingale with respect to $\mathcal{F}_k^{\hat{\tau}}$.

To prove the following lemma, let us first define the continuation set C and stopping set D:

Definition 2.2.3. *Continuation and Stopping set*

$$\begin{aligned}
C &:= \{x : \hat{V}(x) > g(x)\} \\
D &:= \{x : x \notin C\}
\end{aligned}$$

Then define a specific sequence of observation time $\hat{\tau}^* = \tau_1^*, \tau_2^*, ...$, which we will prove to be an optimal strategy:

Definition 2.2.4. *Observation Times*

Define the observation times by the following recursive construction:

$$t^*(x) = \begin{cases} \inf\{t : \hat{V}(x) = \mathbb{E}_x[e^{-rt}(\hat{V}(X_t) - c)]\}, & x \in C \\ \infty, & x \in D \end{cases}$$

Define:

$$\begin{aligned}
\tau_0^* &:= 0 \\
\tau_{k+1}^* &:= \tau_k^* + t^*(X_{\tau_k^*}), \quad k \ge 0
\end{aligned}$$

Let $m \in 0, 1, ...$ be the index of the last observation time, i.e.:

$$\begin{aligned}
m &= \max\{k : \tau_k^* < \infty\} \\
\textit{i.e.} \quad \tau_m^* &< \infty, \quad \tau_{m+1}^* = \infty, ...
\end{aligned}$$

Lemma 2.2.2. *Let $\hat{\tau}^*, m$ be defined as in Definition 2.2.4, then $k \to e^{-r(\tau_k^* \wedge \tau_m^*)}\hat{V}(X_{\tau_{k\wedge m}}) - c\sum_{1\le i\le(k\wedge m)} e^{-r\tau_i^*}$ is a martingale with respect to $\mathcal{F}_k^{\hat{\tau}^*}$.*

Proof. Define $S_k := e^{-r(\tau_k^* \wedge \tau_m^*)} \hat{V}(X_{\tau_{k\wedge m}}) - c \sum_{1 \le i \le (k\wedge m)} e^{-r\tau_i^*}$, for $k = 0$:

$$\begin{aligned}
\mathbb{E}_x[S_1] =& \mathbb{E}_x[e^{-r(\tau_{k+1}^* \wedge \tau_m^*)} \hat{V}(X_{\tau_{k+1\wedge m}}) - c \sum_{1 \le i \le (k+1\wedge m)} e^{-r\tau_i^*}] \\
=& \mathbb{E}_x[e^{-r(\tau_1^* \wedge \tau_m^*)} \hat{V}(X_{\tau_{1\wedge m}}) - c \sum_{1 \le i \le (1\wedge m)} e^{-r\tau_i^*}] \\
=& 1_{m\ge 1} \mathbb{E}_x[e^{-r\tau_1^*}(\hat{V}(X_{\tau_1}) - c)] + 1_{m<1}\hat{V}(x) \\
=& 1_{m\ge 1}\hat{V}(x) + 1_{m=0}\hat{V}(x) \\
=& \hat{V}(x)
\end{aligned} \tag{2.2}$$

For $k = 1, 2, \dots$:

$$\begin{aligned}
& \mathbb{E}_x[S_{k+1} | \mathcal{F}_k^{\hat{\tau}^*}] \\
=& \mathbb{E}_x[e^{-r(\tau_{k+1}^* \wedge \tau_m^*)} \hat{V}(X_{\tau_{k+1\wedge m}}) - c \sum_{1 \le i \le (k+1\wedge m)} e^{-r\tau_i^*} | \mathcal{F}_k^{\hat{\tau}^*}] \\
=& 1_{m>k} \mathbb{E}_x[e^{-r\tau_{k+1}^*} \hat{V}(X_{\tau_{k+1\wedge m}}) - c \sum_{1 \le i \le k+1} e^{-r\tau_i^*} | \mathcal{F}_k^{\hat{\tau}^*}] + 1_{m\le k}(e^{-r\tau_m^*}\hat{V}(X_{\tau_m}) - c \sum_{1\le i \le m} e^{-r\tau_i^*}) \\
=& 1_{m>k}(e^{-r\tau_k^*} \mathbb{E}_x[e^{-r(\tau_{k+1}^* - \tau_k^*)}(\hat{V}(X_{\tau_{k+1\wedge m}}) - c) | \mathcal{F}_k^{\hat{\tau}^*}] - c \sum_{1\le i\le k} e^{-r\tau_i^*}) + 1_{m\le k}(e^{-r\tau_m^*}\hat{V}(X_{\tau_m}) - c \sum_{1\le i \le m} e^{-r\tau_i^*}) \\
& \text{since } \tau_k^* \in \mathcal{F}_k^{\hat{\tau}^*} \\
=& 1_{m>k}(e^{-r\tau_k^*} \mathbb{E}_{X_{\tau_k^*}}[e^{-r\tau_1^*}(\hat{V}(X_{\tau_1}) - c)] - c \sum_{1\le i\le k} e^{-r\tau_i^*}) + 1_{m\le k}(e^{-r\tau_m^*}\hat{V}(X_{\tau_m}) - c \sum_{1\le i \le m} e^{-r\tau_i^*}) \\
& \text{by strong Markov property} \\
=& 1_{m>k}(e^{-r\tau_k^*}\hat{V}(X_{\tau_k}) - c \sum_{1\le i\le k} e^{-r\tau_i^*}) + 1_{m\le k}(e^{-r\tau_m^*}\hat{V}(X_{\tau_m}) - c \sum_{1\le i \le m} e^{-r\tau_i^*}), \quad \textit{by } (2.2) \\
=& e^{-r(\tau_k^* \wedge \tau_m^*)}\hat{V}(X_{\tau_{k\wedge m}}) - c \sum_{1 \le i \le (k\wedge m)} e^{-r\tau_i^*} \\
=& S_k
\end{aligned}$$

Therefore we conclude that $S_k = e^{-r(\tau_k^* \wedge \tau_m^*)} \hat{V}(X_{\tau_{k\wedge m}}) - c \sum_{1 \le i \le (k\wedge m)} e^{-r\tau_i^*}$ is a martingale with respect to $\mathcal{F}_k^{\hat{\tau}^*}$.

Theorem 2.2.3. *Optimal Strategy*
Assume that $\hat{V}$ is a fixed point of $\mathcal{J}$, then $V = \hat{V}$. Moreover, $\hat{\tau} = \{\tau_1^*, \tau_2^*, \dots\}; \tau = \tau_m^*$ provides an optimal strategy.

Proof. First we prove that $\hat{V} \ge V$.
We have $\hat{V} \in F \implies \hat{V} \ge g(x)$, so $\forall k$, we have:

$$\begin{aligned}
& e^{-r\tau_k} g(X_{\tau_k}) - c \sum_{1\le i\le k} e^{-r\tau_i} \le e^{-r\tau_k}\hat{V}(X_{\tau_k}) - c \sum_{1\le i\le k} e^{-r\tau_i} \\
\implies & \mathbb{E}_x[e^{-r\tau_k} g(X_{\tau_k}) - c \sum_{1\le i\le k} e^{-r\tau_i}] \le \mathbb{E}_x[e^{-r\tau_k}\hat{V}(X_{\tau_k}) - c \sum_{1\le i\le k} e^{-r\tau_i}]
\end{aligned}$$

Taking supremum over all stopping time $\tau \in \mathcal{S}^{\hat{\tau}}_{\tau_k}$, we have:

$$\begin{aligned}
&\sup_{\tau\in\mathcal{S}^{\hat{\tau}}_{\tau_k}} \mathbb{E}_x[e^{-r\tau}g(X_\tau) - c\sum_{\tau_i\leq\tau} e^{-r\tau_i}] \leq \sup_{\tau\in\mathcal{S}^{\hat{\tau}}_{\tau_k}} \mathbb{E}_x[e^{-r\tau}\hat{V}(X_\tau) - c\sum_{\tau_i\leq\tau} e^{-r\tau_i}] \leq \hat{V}(x) \\
\Longrightarrow &\sup_{\hat{\tau}} \sup_{\tau\in\mathcal{S}^{\hat{\tau}}_{\tau_k}} \mathbb{E}[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau\}}] \leq \hat{V}(x) \\
\Longrightarrow &V(x) \leq \hat{V}(x)
\end{aligned}$$

by the Optional Sampling Theorem and the supermartingale property as proved in Lemma 2.2.1.
Next we prove that $\hat{V} \leq V$. By the martingale property as in Lemma 2.2.2 and the Optional Sampling Theorem, we have:

$$\begin{aligned}
\hat{V}(x) &= \mathbb{E}_x[e^{-r\tau^*_m}\hat{V}(X_{\tau^*_m}) - c\sum_{1\leq i\leq m} e^{-r\tau^*_i}] \\
&= \mathbb{E}_x[e^{-r\tau^*_m}g(X_{\tau^*_m}) - c\sum_{1\leq i\leq m} e^{-r\tau^*_i}] \\
&\leq \sup_{\hat{\tau}} \sup_{\tau\in\mathcal{S}^{\hat{\tau}}_{\tau_k}} \mathbb{E}_x[e^{-r\tau}g(X_\tau) - c\sum_{0<\tau_i\leq\tau} e^{-r\tau}] \\
&\leq V(x)
\end{aligned}$$

Hence the fixed point $\hat{V}$ equals the value function V. Furthermore, the strategy $\hat{\tau} = \{\tau^*_1, \tau^*_2, \dots\}; \tau = \tau^*_m$ is an optimal strategy.

Remark 2.2.1. *In the proof of 2.2.3, we take supremum only over finite stopping times when applying the Optional Sampling Theorem. It follows from Remark 2.1.2 that it is sufficient to consider all finite stopping times.*

Corollary 2.2.3.1. *Uniqueness of the Fixed Point*
The fixed point $\hat{V}$ is unique.

Proof. Suppose $\hat{V}_1$ and $\hat{V}_2$ are two fixed points of $\mathcal{J}$, then: $\hat{V}_1 = V$ and $\hat{V}_2 = V$, by the uniqueness of the supremum V, we have $\hat{V}_1 = \hat{V}_2$.

2.3 How to Find the Fixed Point

2.3.1 Existence of the Fixed Point

We have proven in the previous sections that if a fixed point of $\mathcal{J}$ exists, then it is unique. In this section we will prove the existence of such a fixed point, and we will further suggest one way to find the fixed point, as well as how fast this procedure would be.

Let f, F be defined as in Definition 2.2.1, we will now prove that we can use iteration from the lower bound to reach the fixed point.

Lemma 2.3.1. *$x \in \mathbb{R}_+$, $f_1 \in F$, $f_2 \in F$. If $f_1 \leq f_2$, then $\mathcal{J}f_1 \leq \mathcal{J}f_2$.*

Proof.

$$\begin{aligned}\mathcal{J}f_1(x) &= \max(g(x),\ \sup_{t\geq 0}\mathbb{E}_x[e^{-rt}(f_1(X_t)-c)]) \\ &\leq \max(g(x),\ \sup_{t\geq 0}\mathbb{E}_x[e^{-rt}(f_2(X_t)-c)]) \\ &= \mathcal{J}f_2(x)\end{aligned}$$

by monotonicity of the expectation operator.

Now we define a sequence of functions recursively by:

$$\begin{aligned} f_0 &= g(x) \\ f_{n+1} &= \mathcal{J}f_n, \quad n \geq 1 \end{aligned}$$

Lemma 2.3.2. *$\{f_n\}_{n\geq 0}$ is an increasing sequence.*

Proof. Clearly,

$$f_1 = \mathcal{J}f_0 \geq g(x) = f_0$$

Now assume $f_k \geq f_{k-1}$ for some k, then: $f_{k+1} = \mathcal{J}f_k \geq \mathcal{J}f_{k-1} = f_k$. Thus the statement follows by induction.

Lemma 2.3.3. *$\forall f \in F, \mathcal{J}f \in F$. Define: $f_\infty := \lim_{n\to\infty} f_n$, then f_∞ exists and $f_\infty \in F$.*

Proof. Taking $t = 0$, we have $\sup_{t\geq 0}\mathbb{E}_x[e^{-rt}g(X_t)] \geq g(x)$, so $(\mathcal{J}f)(x) \geq g(x)$. The upper bound is obvious.

The sequence $\{f_n\}_{n\geq 0}$ is bounded and thus it has a finite limit, for the upper and lower bounds, we have:

$$\begin{aligned} f_\infty &\geq f_0 = g(x) \\ f_\infty &= \lim_{n\to\infty} f_n = sup f_n \leq M \end{aligned}$$

And thus $f_\infty \in F$.

Theorem 2.3.4. *The Existence of the Fixed Points*
The function $f_\infty \in F$ is a fixed point of $\mathcal{J}$.

Proof. $\forall n$, we have $f_\infty \geq f_{n+1} \geq f_n$. By Lemma 2.3.1:

$$\mathcal{J}f_\infty \geq \mathcal{J}f_{n+1} \geq \mathcal{J}f_n = f_{n+1}$$

Taking limit on both sides, we have:

$$\lim \mathcal{J}f_\infty \geq f_\infty$$

So $\mathcal{J} f_\infty \geq f_\infty$.

For the other direction, fix x and let:

$$t_\infty = \inf\{t : \mathbb{E}_x[e^{-rt} f_\infty(X_t)] - c \text{ attains its maximum}\}$$

Then:

$$\begin{aligned} f_{n+1}(x) &= \mathcal{J} f_n(x) \\ &= \max(g(x), \ \sup_{t\geq 0} \mathbb{E}_x[e^{-rt}(f_n(X_t) - c)] \\ &\geq \max(g(x), \ \mathbb{E}_x[e^{-rt_\infty}(f_n(X_{t_\infty}) - c)]) \end{aligned}$$

Taking limit on both sides, by DCT, we have:

$$\begin{aligned} \lim_{n\to\infty} f_{n+1}(x) = f_\infty &\geq \max(g(x), \ \lim_{n\to\infty} \mathbb{E}_x[e^{-rt_\infty}(f_n(X_{t_\infty}]) - c)]) \\ &= \max(g(x), \ \mathbb{E}_x[e^{-rt_\infty}(f_\infty(X_{t_\infty}]) - c)]) \\ &= \mathcal{J} f_\infty \end{aligned}$$

So we have $f_\infty \geq \mathcal{J} f_\infty$, and thus f_∞ is a fixed point of $\mathcal{J}$.

Corollary 2.3.4.1. *f_∞ is the unique fixed point of $\mathcal{J}$, $f_\infty = V$*

Proof. This corollary follows directly from Theorem 2.5.2 and Theorem 2.3.4.

Remark 2.3.1. *We can also start the iteration from the upper bound, i.e.:*

$$\begin{aligned} &f_0 = M \\ &f_{n+1} = \mathcal{J} f_n, \quad n \geq 0 \end{aligned}$$

Then $f_\infty := \lim_{n\to\infty} f_n$ is also the fixed point of $\mathcal{J}$. The proof follows similarly.

Moreover, by monotonicity of $\mathcal{J}$, starting the iteration from any function $f \in F$, we will eventually reach the fixed point.

2.3.2 An Example

In this section we take the perpetual American put option for example. In this specific example let us assume that the underlying asset $(X_t)_{t\geq 0}$ follows a geometric Brownian motion which solves:

$$\begin{cases} dX_t &= \mu X dt + \sigma X dW_t \\ X_0 &= x \end{cases} \tag{2.3}$$

where W_t is a standard Brownian motion, $\mu > 0$, $\sigma > 0$, and $|\mu - \frac{1}{2}\sigma^2| < \infty$. Our goal is to find an optimal strategy which will maximize the discounted payoff at time 0.

Let us define a collection of functions Q such that:

$$Q := \{q \text{ Borel-measurable}, (K - x)^+ \leq q \leq K\}$$

where the upper bound is the strike price. Note that here we ask $K > c > 0$, otherwise the problem becomes trivial. If we look at the associated optimal stopping problem we can see that the value function has the expression:

$$V(x) = \sup_{\hat{\tau}} \sup_{\tau \in \mathcal{S}^{\hat{\tau}}} \mathbb{E}[e^{-r\tau}(K - X_\tau)^+ - \sum_{i=1}^{\infty} ce^{-r\tau_i} 1_{\{\tau_i \leq \tau\}}]$$

Let us write the value of a perpetual American put option under the physical measure under continuous information as $V^{Ame} = \sup_\tau \mathbb{E}_x[e^{-r\tau}(K - X_\tau)^+]$, then V is bounded above by V^{Ame}. Note that to find the fixed point we can start the iteration here from both the lower and the upper bound. Let us start from the lower bound for instance. We define:

$$\begin{aligned} q_0 &= (K - x)^+ \\ q_n &= \mathcal{J} q_{n-1}, \quad n \geq 1 \end{aligned}$$

Property of the Value Function

Proposition 2.3.1. *$\forall q \in Q, \mathcal{J} q \in Q, q_\infty := \lim_{n\to\infty} q_n$ exists and $q_\infty \in Q$. q_∞ is the unique fixed point of $\mathcal{J}$ in Q.*

The proof for the fixed point follows directly from Corollary 2.3.4.1.

Lemma 2.3.5. *The operator preserves monotonicity: q_n decreases in x for all $n \geq 0$, q_∞ decreases in x.*

Proof. Let us assume $0 < x_1 < x_2$, then $q_0(x_1) \geq q_0(x_2)$.

We can prove this property by a simple argument according to the Markov property: with probability 1, the sample path starting from x_2 dominates the sample path starting from x_1, otherwise if they ever intersect, they stay the same.

So for $0 < x_1 < x_2$, we have:

$$\mathbb{E}_{x_1}[e^{-rt}(q_0(X_t) - c)]) \geq \mathbb{E}_{x_2}[e^{-rt}(q_0(X_t) - c)])$$

And thus:

$$\begin{aligned} q_1(x_1) &= \mathcal{J} q_0(x_1) \\ &= \max((K - x_1)^+, \ \sup_{t\geq 0} \mathbb{E}_{x_1}[e^{-rt}(q_0(X_t) - c)]) \\ &\geq \max((K - x_2)^+, \ \sup_{t\geq 0} \mathbb{E}_{x_2}[e^{-rt}((q_0(X_t) - c)]) \\ &= \mathcal{J} q_0(x_2) = q_1(x_2) \end{aligned}$$

Similarly, $q_n(x_1) \geq q_n(x_2)$ by induction. Taking limit on both sides, we see the fixed point decreases in x. Therefore the operator $\mathcal{J}$ preserves monotonicity.

Lemma 2.3.6. *The operator preserves convexity: q_n is convex in x for all $n \geq 0$, q_∞ is convex in x.*

Proof. For $n = 0$, we have $\sup_{t\geq 0} \mathbb{E}_x[e^{-rt}((K - X_t)^+ - c)]$ convex, since the second derivative of an European option w.r.t x is positive, obviously $(K - x)^+$ is convex, we also have the function $\max(\cdot, \cdot)$ convex, therefore q_1 is convex. For $n \geq 1$ we have $\mathbb{E}_x[e^{-rt}(q_n(X_t) - c)]$ convex since the expectation does not change the convexity. Therefore we have q_n convex in x for all n. Taking limit on both sides, the pointwise limit of a convex function is convex, therefore we have q_∞ convex.

Lemma 2.3.7. *q_n is uniformly Lipschitz (1) on $\mathbb{R}_+$ for all $n \geq 0$.*

Proof. It is obvious that q_0 is Lipschitz (1). By Lemma 2.3.6, the convexity of q_n implies that q_n is locally Lipschitz for all n:

Assume $0 < a < x < y < b < c$, then by convexity of q_n:

$$\frac{q_n(a) - q_n(0)}{a - 0} \leq \frac{q_n(y) - q_n(x)}{y - x} \leq \frac{q_n(c) - q_n(b)}{c - b}$$
$$\implies |\frac{q_n(y) - q_n(x)}{y - x}| \leq \max(|\frac{q_n(a) - q_n(0)}{a - 0})|, \ |\frac{q_n(c) - q_n(b)}{c - b}|)$$

Since q_n decreases in x,

$$\max(|\frac{q_n(a) - q_n(0)}{a - 0})|, \ |\frac{q_n(c) - q_n(b)}{c - b}|) = |\frac{q_n(a) - q_n(0)}{a - 0})|$$

for $a < K$:

$$\begin{aligned} |\frac{q_n(y) - q_n(x)}{y - x}| &\leq |\frac{q_n(a) - q_n(0)}{a - 0})| \\ &\leq |\frac{(K - a) - K}{a - 0})| = 1 \end{aligned}$$

By the fact that q_n is bounded by $(K - x)^+$ and K, and that $q_n(0) = K$ for all n. Therefore we can see that the Lipschitz constant is 1 globally.

Proposition 2.3.2. *q_∞ converges to V^{Ame} as the cost c goes down to 0, and q_∞ converges to $(K - x)^+$ as the cost c goes up to the strike K.*

Proof. First we claim that when $c = 0$, $V = V^{Ame}$, as we are exposed to the continuous filtration. Let us first define a sequence $\{C_n\}_{n\geq 1} = \frac{K}{n}$. As $n \to \infty$, $C_n \to 0$. Let:

$$V_n(x) = \max((K - x)^+, \sup_t \mathbb{E}_x[e^{-rt}((K - x)^+ - C_n)])$$

Take $\epsilon > 0$, choose $n_0 \in \mathbb{N}$ such that $n_0 = \frac{K}{\epsilon}$, then for all $n > n_0$ we have:

$$\begin{aligned} |V_n(x) - V^{Ame}(x)| =& |\max((K - x)^+, \sup_t \mathbb{E}_x[e^{-rt}((K - x)^+ - C_n)]) \\ & - \max((K - x)^+, \sup_t \mathbb{E}_x[e^{-rt^*}(K - x)^+])| \\ \leq & |e^{-rt^*} C_n| = |\frac{e^{-rt^*} K}{n}| \\ \leq & \frac{K}{n} < \frac{K}{n_0} = \epsilon \end{aligned}$$

We notice that as $n \to \infty$, $K - C_n \to K$. Let:

$$V_n(x) = \max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}((K-x)^+ - (K - C_n)])$$

Take $\epsilon > 0$, for all $n > n_0$ we have:

$$\begin{aligned}|V_n(x) - (K-x)^+| =& |\max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}((K-x)^+ - (K - C_n))]) - (K-x)^+| \\ =& |\max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}((K-x)^+ - K + \frac{K}{n})]) - (K-x)^+| \\ \leq& |\frac{e^{-rt^*}K}{n}| < \frac{K}{n_0} = \epsilon\end{aligned}$$

so the proof is complete.

We can see that the fixed point q_∞, as well as all the intermediate iterations are decreasing, convex, bounded below by $(K-x)^+$ and above by the value of the perpetual American put option by definition. In addition, the sequence q_n is an increasing sequence.

A Numerical Example

As we can also see from the numerical example shown below, these properties are indeed what we observe when starting the iteration from the lower bound.

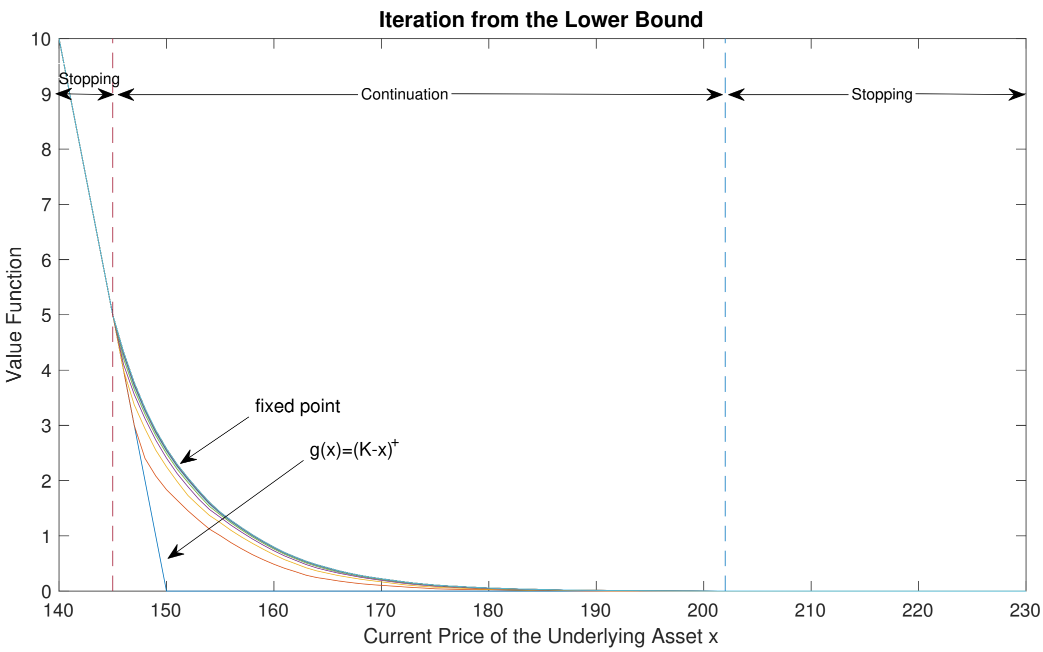

Figure 2.1: Finding the fixed point: $K = 150$, $x = 100$, $\mu = 0.15$, $r = 0.1$, $\sigma = 0.1$, $c = 0.01$, $T = 10$, Number of iterations $N = 40$, Number of spatial steps $N_x = 100$, Number of time steps $N_t = 50$.

Property of the Optimal Strategy

We can see from Figure 2.1 that, when the current value of the underlying process x is small, we are in the stopping area where we do not make observations. As x exceeds some certain threshold, we are in the continuation area and therefore make an observation. However, as x grows larger, we are in the stopping area again as the gap between the fixed point and the lower bound cannot compensate the cost of observation anymore. Therefore, for the particular choice of parameters used in Figure 2.1, the structure of the optimal strategy with respect to the starting point x should be to stop, to continue, and to stop.

However, does the structure of optimal strategy always behave like this, or it is only for this specific choice of parameters? We will prove later the structure depends on the drift term of the underlying process and the constants K and c.

We will also prove that the first optimal observation time τ_1 is always strictly larger than some $\epsilon > 0$. Then we will further prove the rate of convergence for this iteration.

For each step of the iteration, we can find a deterministic solution for the next step. We have:

$$\begin{aligned} q_n(x) &= \mathcal{J} q_{n-1}(x) \\ &= \max((K-x)^+, \ \sup_{t\geq 0} \mathbb{E}_x[e^{-rt}(q_{n-1}(X_t)-c)]) \\ &= \max((K-x)^+, \ \sup_{T\geq 0}(u(T,x)-e^{-rT}c) \end{aligned}$$

where $u(t,x)$ solves the following PDE by the Feynman–Kac formula :

$$\begin{aligned} u_t(t,x) - \mu x u_x(t,x) - \frac{1}{2}\sigma^2 x^2 u_x x(t,x) + ru(t,x) = 0 \\ u(0,x) = q_{n-1}(x) \end{aligned}$$

At every step, starting from the known function q_{n-1}, we take the largest $u(T,x)$ over all possible time horizon T, and further take the maximum with $(K-x)^+$ to get the next q_n. Note that there exists a unique non-negative solution for each iteration. Note that this is also the method we use to perform our numerical examples.

Now let us prove that there is always a lower bound for the continuation set for all iterations.

Lemma 2.3.8. *For all $n \geq 1$, there exist $a_n : 0 < a_n \leq K$, such that $\forall x \in C$ continuation set C, $x \geq a_n$.*

Proof. For $x < K$, we have x strictly larger than some $a_n > b_0$, where b_0 is the point where the value function of the perpetual American option under continuous observation under the physical measure starts to differ from $(K-x)$. This follows from the fact that q_n is bounded above by V^{Ame} by definition.

It follows from the convexity and monotonicity property, and the fact that q_n is bounded below by $(K-x)^+$ that $0 < a_n \leq K$.

Lemma 2.3.9. *The upper bound of the continuation set depends on the drift term of the underlying process X_t:*

Case 1: $\mu - \frac{1}{2}\sigma^2 > 0$*:*

For all $n \geq 1$, there exist $b_n : b_n \geq K$, such that $\forall x \in C$, $x \leq b_n$. That is to say, $C = \{x : a_n \leq x \leq b_n\}$.

Case 2: $\mu - \frac{1}{2}\sigma^2 < 0$*:*

Such b_n's do not exist. $C = \{x : x \geq a_n\}$.

Case 3: $\mu - \frac{1}{2}\sigma^2 = 0$*:*

When $K > 2c$, it is the same as in Case 2; when $K \leq 2c$, it is the same as in Case 1.

Proof. The behaviour of the functions depends on the asymptotic behaviour of the underlying process. Let us first claim that, by the law of the iterated logarithm, for the underlying process which solves (2.3):

1. when $\mu - \frac{1}{2}\sigma^2 > 0$, as $t \to \infty$, $X_t \to \infty$ with probability 1.
2. when $\mu - \frac{1}{2}\sigma^2 < 0$, as $t \to \infty$, $X_t \to 0$ with probability 1.
3. when $\mu - \frac{1}{2}\sigma^2 = 0$, as $t \to \infty$, X_t does not have a limit with probability 1.

For any $x_0 > 0$,

$$\begin{aligned}
\lim_{t\to\infty} \mathbb{P}(X_t > x_0) &= \lim_{t\to\infty} \mathbb{P}(xe^{(\mu-\frac{1}{2}\sigma^2)t+\sigma W_t} > x_0) \\
&= \lim_{t\to\infty} \mathbb{P}((\mu - \frac{1}{2}\sigma^2)t + \sigma W_t > \ln \frac{x_0}{x}) \\
&= \lim_{t\to\infty} \mathbb{P}(\frac{W_t}{\sqrt{t}} > \frac{\ln \frac{x_0}{x} - (\mu - \frac{1}{2}\sigma^2)t}{\sigma\sqrt{t}}) \\
&= \lim_{t\to\infty} \Phi((\mu - \frac{1}{2}\sigma^2)t - \frac{\ln \frac{x_0}{x}}{\sigma\sqrt{t}}) \\
&= \begin{cases} 1, & \text{if } \mu - \frac{1}{2}\sigma^2 > 0 \\ 0, & \text{if } \mu - \frac{1}{2}\sigma^2 < 0 \\ \frac{1}{2}, & \text{if } \mu - \frac{1}{2}\sigma^2 = 0 \end{cases}
\end{aligned}$$

where Φ is the CDF of the standard normal distribution.

For $x > K$, We let $\mathbb{E}_x[e^{-rT}(q_n(X_T) - c)] > 0$ for some T, which is $\mathbb{E}_x[q_n(X_T)] > c$. As proven above, by monotonicity of q_n, if there exists some b_n, such that $\mathbb{E}_{x_b}[q_n(X_T)] = c$. Then for $x > b_n$, we have $q_n = 0$. In other words, we want to know if there exists some b_n such that, starting from b_n, the expectation of the payoff minus cost at any time in the future in negative.

When $\mu - \frac{1}{2}\sigma^2 > 0$, the Brownian motion increases in time, we have $q_n(X_t) - c \leq V^{Ame}(X_t) - c = -c$, as $X_t \to \infty$ with probability 1, since the term $(\mu - \frac{1}{2}\sigma^2)t$ dominates σW_t. We can write:

$$\mathbb{E}[e^{-rt}(q_x(X_t) - c)] \leq \mathbb{E}[e^{-rt}(V^{Ame}(X_t) - c)] \leq V^{Ame}(x) - c$$

since $V^{Ame}(x) - c$ is a supermartingale. Take x sufficiently large such that $V^{Ame}(x) < c$, we can write:

$$\lim_{t \to \infty} \mathbb{E}[e^{-rt}(q_x(X_t) - c)] < 0$$

Which means by the monotonicity property, there exists such b_n. Therefore the structure of the optimal strategy would be $D \to C \to D$.

When $\mu - \frac{1}{2}\sigma^2 < 0$, the Brownian motion decreases in time, we have $\lim_t q_n(X_t) - c \geq \lim_t g(X_t) - c = \lim_{t\to\infty}(K - X_t)^+ - c = K - c$, as $X_t \to 0$ with probability 1. By DCT the function will converge to 0 but will never reach 0. Therefore there is no stopping region coming after the continuation region, we always observe once x exceeds some point. The structure of the optimal strategy would be $D \to C$.

When $\mu - \frac{1}{2}\sigma^2 = 0$, the process X_t does not have a limit with probability 1. For t large, with probability $\frac{1}{2}$, it will end up being ∞ and with probability $\frac{1}{2}$ it would be 0. And thus the function would have a bound like:

$$\begin{aligned}
&e^{-rt}((K - X_t)^+ - c) \leq e^{-rt}(q_n(X_t) - c) \leq e^{-rt}(V^{Ame}(X_t) - c) \\
\implies &\mathbb{E}[e^{-rt}((K - X_t)^+ - c)] \leq \mathbb{E}[e^{-rt}(q_n(X_t) - c)] \leq \mathbb{E}[e^{-rt}(V^{Ame}(X_t) - c)] \\
\implies &e^{-rt}(\frac{1}{2}K - c)] \leq \mathbb{E}[e^{-rt}(q_n(X_t) - c)] \leq e^{-rt}(\frac{1}{2}K - c) \\
\implies &\mathbb{E}[e^{-rt}(q_n(X_t) - c)] = e^{-rt}(\frac{1}{2}K - c)
\end{aligned}$$

We can see that, when $K > 2c$, the limit $e^{-rt}(\frac{1}{2}K - c)$ would be positive, so it would be the same as in Case 2. When $K < 2c$, it would be the same in Case 1 that there exist some t such that $\mathbb{E}[e^{-rt}(q_n(X_t) - c)]$ is negative. When $K - 2c$, the continuation area and the stopping area will coincide after the first time $q_n(x) - c = 0$.

Now let us assume that we are in the continuation set, then how does the next observation time t change with x? Let us define:

$$\begin{aligned}
t^*(x, q_n) &= \inf\{t : q_n(x) = \mathbb{E}_x[e^{-rt}(q_{n-1}(x) - c)]\} \\
t^*(x) &= \inf_n t^*(x, q_n)
\end{aligned}$$

As we know from Lemma 2.3.7, all q_n's are uniformly Lipschitz (1), which means the initial condition cannot grow faster than linearly, so that the speed that u grows in t cannot be too fast either: for a small $T - t$, we know that in this backward heat equation $|u(t, x) - u(T, x)|$ is bounded by $\tilde{C}\sqrt{T - t}$, for some constant $\tilde{C}$. However, for a fixed x, to make sure that x is in the continuation region, u at some optimal stopping time $t^*(x, q_n)$ has to exceed the discounted one step cost $ce^{-rt^*(x,q_n)}$. And thus we obtain a lower bound for the optimal stopping time t for all x and all q_n. We can also motivate this result as follows:

Lemma 2.3.10. *$\forall x \in C$, the optimal stopping time $t^*(x, q_n)$ satisfies:*

$$t^*(x, q_n) \geq (\frac{c}{\tilde{C}x})^2$$

for some constant $\tilde{C}$.

Proof. Let us assume $x \in C$. Fix some time horizon T, let:

$$u^T(0,x) = \mathbb{E}_x[e^{-rT} q_{n-1}(X_T)]$$

We have:

$$\begin{aligned}
&|u^T(0,x) - u^0(0,x)| \\
=&|\mathbb{E}_x[e^{-rT} q_{n-1}(X_T)] - q_{n-1}(x)| \\
\le&\mathbb{E}_x[|e^{-rT} q_{n-1}(X_T) - q_{n-1}(x)|] \\
=&\mathbb{E}_x[|e^{-rT} q_{n-1}(X_T) - e^{-rT} q_{n-1}(x) + e^{-rT} q_{n-1}(x) - q_{n-1}(x)|] \\
\le&e^{-rT}\mathbb{E}_x[|q_{n-1}(X_T) - q_{n-1}(x)|] + q_{n-1}(x)|1 - e^{-rT}|
\end{aligned}$$

where the second term $q_{n-1}(x)|1 - e^{-rT}| \le \tilde{C}_1 T$ for some $\tilde{C}_1$ for small T. Let us take a look at the first term. By the Lipschitz property of q_{n-1}, we have:

$$\begin{aligned}
\mathbb{E}_x[|q_{n-1}(X_T) - e^{-rT} q_{n-1}(x)|] \le&\mathbb{E}_x[|X_T - x|] \\
=&x\mathbb{E}[|e^{(\mu - \frac{1}{2}\sigma^2)T + \sigma W_T} - 1|]
\end{aligned}$$

By Cauchy-Schwartz inequality, we can write:

$$\begin{aligned}
\mathbb{E}[|e^{(\mu-\frac{1}{2}\sigma^2)T+\sigma W_T} - 1|] \le&(\mathbb{E}[|e^{(\mu-\frac{1}{2}\sigma^2)T+\sigma W_T} - 1|]^2)^{\frac{1}{2}} \\
=&(\mathbb{E}[|e^{(2\mu-\sigma^2)T+2\sigma W_T} - 2e^{(\mu-\frac{1}{2}\sigma^2)T+\sigma W_T} + 1)|]^{\frac{1}{2}} \\
=&(1 - (\mu - \frac{1}{2}\sigma^2)T - 2(1 - (\mu - \frac{1}{2}\sigma^2)T) + 1)^{\frac{1}{2}} \\
\le&\tilde{C}_2\sqrt{T}
\end{aligned}$$

for some $\tilde{C}_2$. Let t^* be the smallest T that $u^T(0,x) - e^{-rT}c$ attains its maximum, then:

$$\begin{aligned}
ce^{-rt^*(x,q_n)} &\le |u^{t^*(x,q_n)}(0,x) - u^0(0,x)| \\
&\le x(\tilde{C}_2 e^{-rt^*(x,q_n)}\sqrt{t^*(x,q_n)} + \tilde{C}_1 t^*(x,q_n)) \\
&\le x\tilde{C}\sqrt{t^*(x,q_n)} \\
\implies t^*(x,q_n) &\ge (\frac{c}{\tilde{C}x})^2
\end{aligned}$$

for small $t^*(x,q_n)$, for some constant .

Lemma 2.3.11. *On set C, $t^*(x) > \epsilon$, for some $\epsilon > 0$.*

Proof. We know from Lemma 2.3.9 that if the continuation area has an upper bound b_n, the first optimal observation time $t^*(x,q_n) \ge (\frac{c}{b_n})^2$.

When the continuation area does not have an upper bound, let us assume the process starts from a large $x >> K$. Define function $f : t \to \mathbb{E}[e^{-rt}(q_n(X_t) - c)]$. We can see that when $t = 0$, $f(t) = c < 0$. As t increases, by the almost sure continuity of the sample path, $f(t)$ increases until it reaches 0 at some $t'(x)$. The optimal first observation time $t^*(x) > t'(x)$.

Define $\epsilon(x) = \inf_n \epsilon(x,q_n)$, we obtain the lower bound for all $t^*(x,q_n)$. It is obvious that $t^*(x) > \epsilon(x)$.

Lemma 2.3.12. *J is a contraction mapping on $(\{q_n\}_{n\geq 0}, \|\cdot\|_\infty)$.*

Proof. Let $q_i, q_j \in \{q_n\}_{n\geq 0}$, we claim that:

$$d(\mathcal{J}q_i, \mathcal{J}q_j) \leq \beta d(q_i, q_j)$$

for all $x \in \mathbb{R}_+$, and for some $\beta \in [0, 1)$.

Fix x, we have:

$$\begin{aligned}
&d(\mathcal{J}q_i, \mathcal{J}q_j)\\
=&d(\max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}(q_i(X_t)-c)]), \max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}(q_j(X_t)-c)]))\\
\leq&\|\sup_t \mathbb{E}_x[e^{-rt}(q_i(X_t)-c)] - \sup_t \mathbb{E}_x[e^{-rt}(q_j(X_t)-c)]\|_\infty\\
=&\|\mathbb{E}_x[e^{-rt_i^*}(q_i(X_{t_i^*})-c)] - \mathbb{E}_x[e^{-rt_j^*}(q_j(X_{t_j^*})-c)]\|_\infty
\end{aligned}$$

where $t^*(x, q_i)$ is the first time $\mathbb{E}_x[e^{-rt}(q_{i-1}(X_t) - c)]$ attains its maximum, define $t^*(x, q_j)$ similarly. Without loss of generality, let us assume $t^*(x, q_i) \leq t^*(x, q_j)$:

$$\begin{aligned}
&\|\mathbb{E}_x[e^{-rt^*(x,q_i)}(q_i(X_{t^*(x,q_i)})-c)] - \mathbb{E}_x[e^{-rt^*(x,q_j)}(q_j(X_{t^*(x,q_j)})-c)]\|_\infty\\
\leq&\|\mathbb{E}_x[e^{-rt^*(x,q_i)}(q_i(X_{t^*(x,q_i)}) - q_j(X_{tt^*(x,q_i)}))]\|_\infty\\
\leq&e^{-rt^*(x,q_i)}\|q_i(x)-q_j(x)\|_\infty \leq e^{-r\epsilon}\|q_i(x)-q_j(x)\|_\infty
\end{aligned}$$

where $\epsilon > 0$ is the lower bound for all $t^*(x)$.

Remark 2.3.2. *q_n is Cauchy and its limit is also in $\{q_n\}_{n\geq 0}$. By the Banach fixed point theorem, $\mathcal{J}$ is a contraction mapping on $(\{q_n\}_{n\geq 0}, \|\cdot\|_\infty)$ with modulus $\beta = e^{-r\epsilon} < 1$. The fixed point exists and is unique, which agrees with our claim before. Furthermore, we can start the iteration from an arbitrary element $q_n \in Q$ to obtain the fixed point.*

Remark 2.3.3. *We can also apply Blackwell's sufficient conditions for a contraction that $\mathcal{J}$ satisfies monotonicity and the discounting conditions.*

Blackwell's sufficient conditions: *Let $X \subset \mathbb{R}^l$ and $B(X)$ is the space of bounded functions $f : X \to \mathbb{R}$, T is a contraction with modulus β if:*

1. [Monotonicity] $f, g \in B(X)$, and $f \leq g$ for all $x \in X$. $\implies Tf \leq Tg$ for all $x \in X$.

2. [Discounting] There exists some $\beta \in (0, 1)$, such that $T(f + a)(x) \leq Tf(x) + \beta a$ for all $f \in B(X)$, $a \geq 0$, $x \in X$.

As we have proven the monotonicity in Lemma 2.3.5, it suffices to check the discounting condition here. Let $a \geq 0$:

$$\begin{aligned}
&\mathcal{J}(q_n + a) - \mathcal{J}(q_n)\\
=&\max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}(q_n(X_t)+a-c)]) - \max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}(q_n(X_t)-c)])\\
\leq&\sup_t \mathbb{E}_x[e^{-rt}(q_n(X_t)+a-c)] - \sup_t \mathbb{E}_x[e^{-rt}(q_n(X_t)-c)] = e^{-rt^*(x)}a\\
\leq&e^{-r\epsilon}a
\end{aligned}$$

where $\beta = e^{-r\epsilon} < 1$. Therefore $\mathcal{J}$ is a contraction.

Theorem 2.3.13. *Rate of convergence*

The sequence q_n converges uniformly to q_∞, and the rate of convergence is exponential.

Proof. First we want to show that $q_1 - q_0$ is bounded by some constant β_0:

$$\begin{aligned} q_1 - q_0 &= \max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}(q_0(X_t) - c)]) - q_0(x) \\ &= \max((K-x)^+, \sup_t \mathbb{E}_x[e^{-rt}((K-X_t)^+ - c)]) - (K-x)^+ \\ &\leq e^{-rt^*}(K-c) = \beta_0 \end{aligned}$$

Since $(K-x)^+ \geq 0$ and $(K-X_t)^+ - c \leq K - c$, and t^* is the first time $\mathbb{E}_x[e^{-rt}(K-X_t)^+]$ attains its maximum.

We now claim that:

$$q_n \leq q_\infty \leq q_n + \frac{\beta_0}{1-\beta}\beta^n$$

where β is the modulus of the contraction mapping $\mathcal{J}$. We know that q_n is an increasing sequence, so the left inequality holds. We can write:

$$\begin{aligned} q_\infty - q_n &= q_\infty - \mathcal{J}q_{n-1} \\ &= \mathcal{J}q_\infty - \mathcal{J}q_{n-1} \\ &\leq \beta(q_\infty - q_{n-1}) \\ &= \beta(q_\infty - q_n + q_n - q_{n-1}) \\ \implies q_\infty - q_n &\leq \frac{\beta}{1-\beta}(q_n - q_{n-1}) \\ &\leq \frac{\beta^n}{1-\beta}(q_1 - q_0) \\ &\leq \frac{\beta_0}{1-\beta}\beta^n \end{aligned}$$

Therefore we say that the rate of convergence is exponential.

A Numerical Example

In the following numerical examples we used the same parameter set as in Figure 2.1.

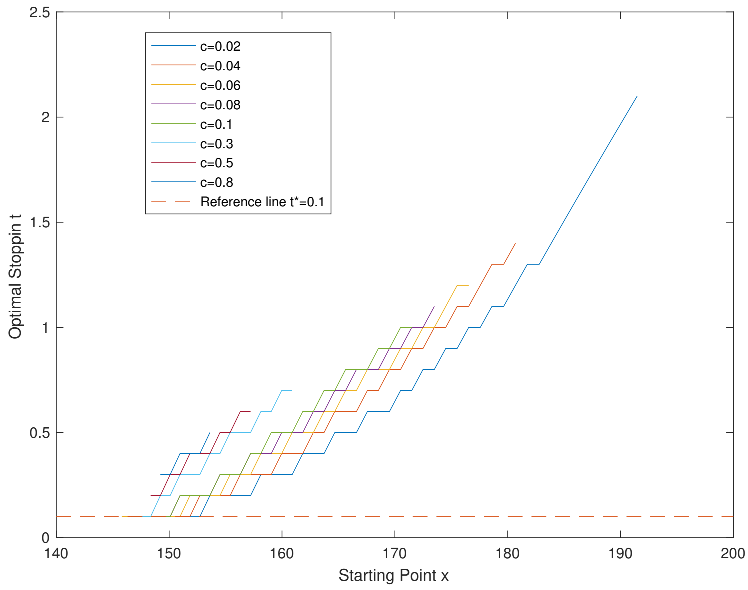

Figure 2.2: Optimal t^* for the next observation on set C

We can see from the figure above that:

(i) On set C, the optimal time for the next observation is bounded below by some positive number;

(ii) when the cost c grows larger, the optimal time also grows larger;

(iii) when the cost c grows larger, the continuation area shrinks.

All the above observations agree with our theoretical approaches before, based on this specific example.

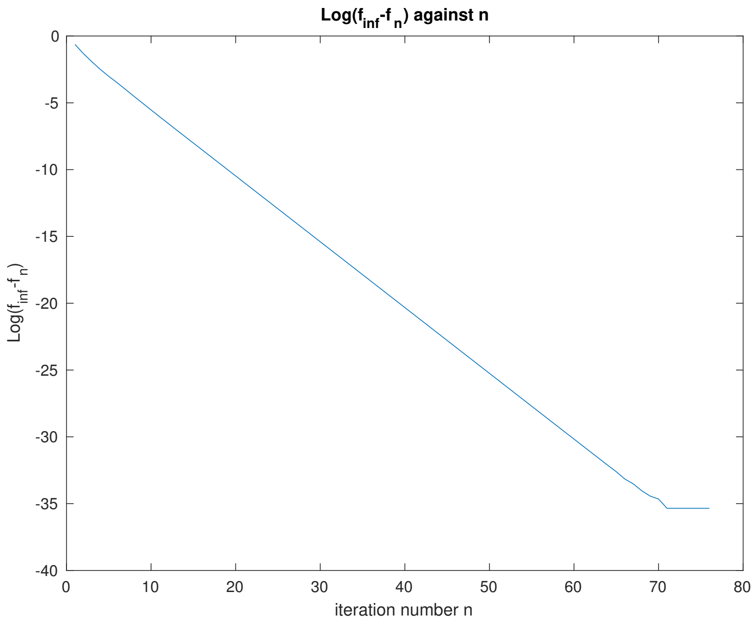

Figure 2.3: Rate of convergence

We can see that $\frac{1}{n}\ln(q_\infty - q_n)$ appears to be linearly, which means the sequence is of exponential convergence rate. Which agrees with Theorem 2.3.13.

2.4 When Number of Observations Is Restricted

It is natural to ask what would the optimal strategy be if we were restricted on the number of observations we can make. We will now prove that we can still use fixed point iteration to find the optimal strategy in this case, depending on the maximal number of observations we can make.

Let us define a sequence V_n when the underlying process can be observed for at most n times:

Definition 2.4.1.

$$V_0(x) = g(x)$$
$$V_n(x) = \sup_{\hat{\tau}} \sup_{\tau \in \mathcal{S}^{\hat{\tau}}, \tau \leq \tau_n} \mathbb{E}[e^{-r\tau} g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i} 1_{\{\tau_i \leq \tau\}}], \quad n \geq 1$$

Theorem 2.4.1.

$$V_n = f_n$$

Proof. First we know that $V_0 - f_0$, let us assume $V_{n-1} = f_{n-1}$, for some $n \geq 1$. Fix x, fix

$\hat{\tau}, \tau$, let $\tau'_k = \tau_{k+1} - \tau_1$, and $\tau' = \tau - \tau_1$. Then:

$$V_n(x) = \sup_{(\hat{\tau},\tau),\tau\le\tau_n} \mathbb{E}[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau\}}]$$
$$= \sup_{(\hat{\tau},\tau)} \mathbb{E}[e^{-r(\tau\wedge\tau_n)}g(X_{\tau\wedge\tau_n}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le(\tau\wedge\tau_n)\}}]$$

We have:

$$\mathbb{E}_x[e^{-r(\tau\wedge\tau_n)}g(X_{\tau\wedge\tau_n}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le(\tau\wedge\tau_n)\}}]$$
$$= \mathbb{E}_x[1_{\{\tau=0\}}(e^{-r(\tau\wedge\tau_n)}g(X_{\tau\wedge\tau_n}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le(\tau\wedge\tau_n)\}})]$$
$$+ \mathbb{E}_x[1_{\{\tau\ge\tau_1\}}e^{-r\tau_1}\mathbb{E}_{X_{\tau_1}}[e^{-r(\tau'\wedge\tau'_{n-1})}g(X_{\tau\wedge\tau_n}) - \sum_{i=1}^{\infty} ce^{-r\tau'_i}1_{\{\tau'_i\le(\tau'\wedge\tau'_{n-1})\}}]]$$
$$-1_{\{\tau\ge\tau_1\}}ce^{-r\tau_1}$$
$$= 1_{\{\tau<\tau_1\}}g(x) + \mathbb{E}_x[1_{\{\tau\ge\tau_1\}}e^{-r\tau_1}\mathbb{E}_{X_{\tau_1}}[e^{-r(\tau'\wedge\tau'_{n-1})}g(X_{\tau\wedge\tau_n}) - \sum_{i=1}^{\infty} ce^{-r\tau'_i}1_{\{\tau'_i\le(\tau'\wedge\tau'_{n-1})\}} - c]]$$
$$= 1_{\{\tau\ge\tau_1\}}g(x) + 1_{\{\tau\ge\tau_1\}}e^{-r\tau_1}\mathbb{E}_x[V_{n-1}(X_{\tau_1}) - c]$$
$$= 1_{\{\tau\ge\tau_1\}}g(x) + 1_{\{\tau\ge\tau_1\}}e^{-r\tau_1}\mathbb{E}_x[f_{n-1}(X_{\tau_1}) - c]$$
$$\le \max(g(x), \sup_{t>0}\mathbb{E}_x[e^{-rt}(f_{n-1}(X_t) - c)])$$
$$= \mathcal{J}f_{n-1}(x)$$
$$= f_n(x)$$

Take supremum over $\hat{\tau}, \tau$ on both sides, we have $V_n \le f_n$.

For the other direction, let us fix x, let $t_n = t(x, f_{n-1})$ be the optimal t for step n, on set C, we have $\mathcal{J}f_{n-1}(x) = \mathbb{E}_x[e^{-rt_n}(f_n(x) - c)]$. For given $\epsilon > 0$, let τ be ϵ-optimal in $V_{n-1}(X_{t_n})$, such that:

$$\tau \le \tau_{n-1}$$
$$V_{n-1}(X_{t_n}) \le \mathbb{E}_{X_{t_n}}[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau\}}] + \epsilon$$

Let $\tau' = t_n + \tau$, $\tau'_{k+1} = t_n + \tau_k$:

$$\mathcal{J}f_{n-1}(x) - \epsilon = \mathbb{E}_x[[e^{-rt_n}(f_{n-1}(X_{t_n}) - c)] - \epsilon$$
$$= \mathbb{E}_x[e^{-rt_n}(V_{n-1}(X_{t_n}) - c)] - \epsilon$$
$$\le \mathbb{E}_x[e^{-rt_n}(\mathbb{E}_{X_{t_n}}[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau\}}] - c)]]$$
$$= \mathbb{E}_x[e^{-r\tau'}g(X_{\tau'}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau'\}}]$$
$$\le V_n(x)$$

Hence we can conclude that $V_n = f_n$.

So when the maximal number of observations is fixed as n, we can still use the fixed point iteration to find our value function: it equals to the n^{th} iteration f_n. And similarly, after doing n iterations, we already have an corresponding optimal strategy.

2.5 Stopping Between Observation Times Is Allowed

In practice, we usually face problems where one observes the underlying process discretely but stops the process continuously. The question being: is it possible find the optimal strategy in a similar way? We will now prove that, in the case where we allow stopping between observation times, we can always reduce it to the case where we do not. We will now characterize this situation and motivate our strategy.

After making an observation, we take a look at the information we have up to now, and decide whether to continue observing the underlying process. We have 3 choices:

(i) To decide to stop observing and stop the process immediately;

(ii) To decide to stop observing and stop the process in a deterministic future time;

(iii) To decide to make another observation and face the same choices after making the next observation.

Let us define the value function W which is associated with the case where stopping is allowed anytime:

$$W : \mathbb{R}_+ \to \mathbb{R}_+$$
$$W(x) = \sup_{\hat{\tau}} \sup_{\tau \in \mathcal{T}^{\hat{\tau}}} \mathbb{E}[e^{-r\tau} g(X_\tau) - \sum_{i=1}^{\infty} c e^{-r\tau_i} 1_{\{\tau_i \leq \tau\}}]$$

Let set F be defined as in Definition 2.2.1. Now we introduce a new function h:

Definition 2.5.1. *Function h*

$$h : \mathbb{R}_+ \to \mathbb{R}_+$$
$$h(x) = \sup_{t \geq 0} \mathbb{E}_x[e^{-rt} g(X_t)]$$

Let us now introduce an operator $\mathcal{H}$ which acts on F:

Definition 2.5.2. *Operator $\mathcal{H}$*
Let be $\mathcal{H}$ be an operator acting on F, such that $\forall f \in F$:

$$(\mathcal{H}f)(x) = \max(h(x), \ \sup_{t>0} \mathbb{E}_x[e^{-rt}(f(X_t) - c)])$$

Let us assume that $\hat{W}$ is a fixed point of the operator $\mathcal{H}$ (the existence comes in the later section), i.e. $\hat{W} = \mathcal{H}\hat{W}$. We can thus write:

$$\hat{W} : \mathbb{R}_+ \to \mathbb{R}_+$$
$$\hat{W}(x) = \max(h(x), \ \sup_{t \geq 0} \mathbb{E}_x[e^{-rt}(\hat{W}(X_t) - c)])$$

Now we will prove that $\hat{W} = W$, and again we will provide an optimal strategy.

To generalize the previous case, first we will prove the following lemma:

Lemma 2.5.1. *For a bounded function $g : \mathbb{R}_+ \to \mathbb{R}_+$, define function h, such that $h(x) = \sup_{t\geq 0} \mathbb{E}_x[e^{-rt}g(X_t)]$. Let $\hat{\tau} = \{\tau_1, \tau_2, \dots\}$ be given, then:*

$$\sup_{\tau\in\mathcal{T}^{\hat{\tau}}} \mathbb{E}_x[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau\}}]$$
$$= \sup_{\tau'\in\mathcal{S}^{\hat{\tau}}} \mathbb{E}_x[e^{-r\tau'}h(X_{\tau'})] - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau'\}}]$$

Where $\mathcal{T}^{\hat{\tau}}$, $\mathcal{S}^{\hat{\tau}}$ are defined as in Definition 2.1.1

Proof. Define $t' := \inf\{t : \mathbb{E}_{X_{\tau'}}[e^{-rt}g(X_t)]$ attains its maximum$\}$.Let $\tau^{opt} = \inf\{\tau \in \mathcal{T}^{\hat{\tau}} : \mathbb{E}_x[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau\}}]$ attains its maximum$\}$, define now $\tau' = \tau_k$, where $\tau_k \leq \tau^{opt} < \tau_{k+1}$. Then:

$$\begin{aligned}
&\mathbb{E}_x[e^{-r\tau'}h(X_{\tau'}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau'\}}]\\
=&\mathbb{E}_x[e^{-r\tau_k}\sup_{t\geq 0}\mathbb{E}_{X_{\tau_k}}[e^{-rt}g(X_t)] - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau_k\}}]\\
=&\mathbb{E}_x[e^{-r\tau_k}\mathbb{E}_{X_{\tau_k}}[e^{-rt'}g(X_{t'})] - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau_k\}}]\\
=&\mathbb{E}_x[\mathbb{E}_x[e^{-r(t'+\tau_k)}g(X_{t'+\tau_k})|\mathcal{G}^{\hat{\tau}}_{\tau_k}] - \sum_{i=1}^{\infty} c1_{\{\tau_i\leq\tau_k\}}]\\
=&\mathbb{E}_x[e^{-r(t'+\tau_k)}g(X_{t'+\tau_k}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau_k\}}]\\
\leq&\mathbb{E}_x[e^{-r\tau^{opt}}g(X_{\tau^{opt}}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau_k\}}]\\
=&\sup_{\tau\in\mathcal{T}^{\hat{\tau}}} \mathbb{E}_x[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau\}}]
\end{aligned}$$

By the strong Markov property. For the other direction, define now $\tau^{opt\prime} = \inf\{\tau \in \mathcal{S}^{\hat{\tau}} : \mathbb{E}_x[e^{-r\tau'}h(X_{\tau'})] - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\leq\tau'\}}]$ attains its maximum$\}$.

Define $t' := \inf\{t : \mathbb{E}_{X_{\tau^{opt'}}}[e^{-rt}g(X_t)] \text{ attains its maximum}\}$, then:

$$\begin{aligned}
&\mathbb{E}_x[e^{-r\tau}g(X_\tau) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau\}}] \\
=&\mathbb{E}_x[e^{-r(\tau^{opt'}+t')}g(X_{\tau^{opt'}+t'}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau\}}] \\
=&\mathbb{E}_x[e^{-r\tau^{opt'}}\mathbb{E}_{X_{\tau^{opt'}}}[e^{-rt'}g(X_{t'})] - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau\}}] \\
=&\mathbb{E}_x[e^{-r\tau^{opt'}}h(X_{\tau^{opt'}}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau\}}] \\
\le&\mathbb{E}_x[e^{-r\tau^{opt'}}h(X_{\tau^{opt'}}) - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau^{opt'}\}}] \\
=&\sup_{\tau'\in\mathcal{S}^{\hat{\tau}}} \mathbb{E}_x[e^{-r\tau'}h(X_{\tau'})] - \sum_{i=1}^{\infty} ce^{-r\tau_i}1_{\{\tau_i\le\tau'\}}]
\end{aligned}$$

Now we can conclude that LHS=RHS by construction, and we have the optimal τ and τ' defined as in the proof. Note that the optimal τ' needs to be the last observation time before the optimal τ.

This is the key property for proving that the 3-choice problem reduces to the 2-choice problem. Let us take the fixed point of $\mathcal{H}$, $\hat{W}$, as defined in Definition 2.5.2, then using the supermartingale and martingale property proved in Lemma 2.2.1 and Lemma 2.2.2 respectively, we can prove $\hat{W} = W$, and give an optimal strategy.

Let us again define the continuation set C' and stopping set D':

Definition 2.5.3. *Continuation and Stopping set*

$$\begin{aligned}
C' &:= \{x : \hat{W}(x) > \sup_{t\ge 0} \mathbb{E}_x[e^{-rt}g(X_t)]\} \\
D' &:= \{x : x \notin C\}
\end{aligned}$$

Then define a sequence of observation time $\hat{\tau}^{*\prime} = \tau_1^{*\prime}, \tau_2^{*\prime}, ...$ and a stopping time $\tau^{*\prime}$, such that:

Definition 2.5.4. *Observation Time and Stopping Time*
Define the observation times by the following recursive construction:

$$t^{*\prime}(x) = \begin{cases} \inf\{t : \hat{W}(x) = \mathbb{E}_x[e^{-rt}(\hat{W}(X_t) - c)]\}, & x \in C \\ \infty, & x \in D \end{cases}$$

Define:

$$\begin{aligned}
\tau_0^{*\prime} &:= 0 \\
\tau_{k+1}^{*}{}' &:= \tau_k^{*\prime} + t^{*\prime}(X_{\tau_k^*}), \quad k \ge 0
\end{aligned}$$

Let $m \in 0, 1, ...$ be the index of the last observation time, i.e.:

$$m = \max\{k : \tau_k^{*\prime} < \infty\}$$
$$i.e. \quad {\tau_m^*}' < \infty, \quad {\tau_{m+1}^*}' = \infty, ...$$

Then:

$$\tau^{*\prime} := {\tau_m^*}' + \inf\{t : \hat{W}(X_{{\tau_m^*}'}) = \mathbb{E}_{X_{{\tau_m^*}'}}[e^{-rt} g(X_{{\tau_m^*}'+t})]\}$$

Theorem 2.5.2. *Allowing Stopping at Anytime*
The fixed point $\hat{W}$ equals W. Moreover, $\{\hat{\tau}^{*\prime}, \tau^{*\prime}\}$ provides an optimal strategy.

Proof. The proof of $\hat{W} \geq W$ is the same as in Theorem 2.2.3. The other direction follows directly from Lemma 2.5.1. Knowing the fact that the observation sequence in Definition 2.5.4is defined as in Definition 2.2.4, and that the stopping time is defined as in Lemma 2.5.1, the claim that $\{\hat{\tau}^{*\prime}, \tau^{*\prime}\}$ provides an optimal strategy follows from Theorem 2.2.3 and Lemma 2.5.1.

Remark 2.5.1. *We can easily see from the definition that $\forall f \in F, \mathcal{J}f \leq \mathcal{H}f$, which means that based on the same information, our payoff would be higher if given more freedom on choosing the stopping time.*

Also, by simply substituting g in the operator $\mathcal{J}$ by h, we construct $\mathcal{H}$ and thus can easily prove that $\hat{W} = W$ using the same reasoning in Section 2.2. The difference is that the space would now be bounded below by h, which is obvious from the definition of $\mathcal{H}$.

We can therefore conclude that, in this chapter, we only need to consider the case where stopping is allowed only immediately after making an observation.

Chapter 3

The Quickest Detection Problem

In this chapter we will take a look at the quickest detection problem with discrete costly observations. We will first discuss the classical quickest detection problem, then formulate it under our settings. We will see that it can be regarded as a specific example of the previous case, which we can still use the fixed point approach to solve, and find an optimal strategy accordingly.

3.1 The Classical Quickest Detection Problem

In the classical quickest detection problem [1], one observes the trajectory of a Brownian motion $(X_t)_{t\geq 0}$, whose drift changes from 0 to μ at some random time θ, and thus solves:

$$dX_t = \mu 1_{\{\theta \leq t\}} dt + \sigma dW_t \tag{3.1}$$

where $\mu > 0, \sigma > 0$, and $(W_t)_{t\geq 0}$ is a standard Brownian motion. The random time θ is independent of W_t with a known distribution, usually it is assumed that θ takes value 0 with known probability π, and given $\theta > 0$,it follows an exponential distribution with known parameter $\lambda > 0$. i.e.:

$$\begin{aligned} &\mathbb{P}(\theta = 0) = \pi; \\ &\mathbb{P}(\theta > t) = (1 - \pi)e^{-\lambda t} \end{aligned} \tag{3.2}$$

The random time θ is called the "disorder time". Our goal in the detection problem is to detect as quickly as possible when the change has happened, while try not to declare the change before it actually appeared. In other words:

(i) After time θ, the delay in declaration of disorder should be as short as possible;

(ii) A false alarm (declaring the disorder before it actually appears) should happen as rarely as possible.

Let us assume we are under an infinite time horizon.

By continuously observing the process X_t, our information is the natural filtration generated by X_t: $\mathcal{F}_t^X = \sigma\{X_s : 0 \leq s \leq t\}$, the goal is to identify a stopping time $\tau \in \mathcal{F}_t^X$, as our declaration that the drift has changed. Let us introduce a posteriori probability process Π_t:

$$\Pi_t = \mathbb{P}(\theta \leq t | \mathcal{F}_t^X), \;\; \Pi_0 = \pi$$

Naturally, we would ask this stopping time τ to be close to θ in some sense. We can for instance form the value function for the quickest detection problem in the Bayesian setting as:

$$V(\pi) = \inf_\tau (\mathbb{P}(\tau < \theta) + c\mathbb{E}_\pi[\tau - \theta]^+)$$

where the term $\mathbb{P}(\tau < \theta)$ is the probability of a "false alarm", $\mathbb{E}_\pi[\tau - \theta]^+$ is the delayal penalty, and $c > 0$ is a known constant represents the relative weight between those two terms in our valuation, or the running cost in our delay.

We would like to formulate the value function in terms of Π_t:

$$\begin{aligned}
V(\pi) &= \inf_\tau (\mathbb{P}(\tau < \theta) + c\mathbb{E}_\pi[\tau - \theta]^+) \\
&= \inf_\tau (\mathbb{E}_\pi[1 - \Pi_\tau] + c\mathbb{E}_\pi[\tau - \theta]^+) \\
&= \inf_\tau (\mathbb{E}_\pi[1 - \Pi_\tau] + c\mathbb{E}_\pi[1_{\tau > \theta} \int_\theta^\tau dt]) \\
&= \inf_\tau (\mathbb{E}_\pi[1 - \Pi_\tau] + c\mathbb{E}_\pi[\int_0^\infty 1_{\tau > t} 1_{\theta \leq t} dt]) \\
&= \inf_\tau (\mathbb{E}_\pi[1 - \Pi_\tau] + c\mathbb{E}_\pi[\int_0^\tau \Pi_t dt]) \\
&= \inf_\tau \mathbb{E}_\pi[1 - \Pi_\tau + c \int_0^\tau \Pi_t dt]
\end{aligned}$$

It is verified in [1] that the posteriori probability process Π_t is a diffusion process.

3.2 Quickest Detection with Discrete Costly Observations

In this section we will formulate the quickest detection problem with discrete costly observations. Let us assume that our underlying process $(X_t)_{t \geq 0}$ is a Brownian motion with drift changing from 0 to μ at a random time θ, where the dynamics of X_t and distribution of θ are as described in Section 3.1. Our goal is to provide a best estimation of the unknown parameter . Suppose that $\mathbb{P}(\theta = 0) = \pi$ is know at time 0, also the starting point of X_t is know, $X_0 = x$.

Similarly, instead of observing the process X_t continuously, we can now only observe it at discrete random time points. Based on the information we have, we can find an optimal time to stop observing the process and declare the disorder, i.e. the best estimation of θ. Note that the information we have depends on the sequence of observing times we choose.

From the realted discussion in Chapter 2 and Section 3.1, we can formulate this detection problem into an optimal stopping problem.

3.2.1 Formulation of the Problem

Let us assume that the underlying process X_t follows (3.1) and θ satisfies (3.2).

Let us define a sequence of observation times $\hat{\tau} = \{\tau_k\}_{k\geq 1}$ the same way as in Chapter 2, such that each element of the sequence is measurable with respect to the filtration generated by all the previous observations. Let us then define the filtration associated with the specific sequence $\hat{\tau}$ up to time t as:

$$\mathcal{G}_t^{\hat{\tau}} = \sigma\{X_{\tau_1}, X_{\tau_2}, \ldots, X_{\tau_i}; \tau_1, \tau_2, \ldots, \tau_i; \tau_i \leq t\}$$

Let us define the set of stopping times $\mathcal{S}^{\hat{\tau}}$ and $\mathcal{T}^{\hat{\tau}}$ as in 2.5.3.

Let us introduce the posteriori probability process Π_t under discrete observations as:

$$\Pi_t = \mathbb{P}(\theta \leq t | \mathcal{G}_t^{\hat{\tau}}),\ \ \Pi_0 = \pi$$

Let us assume that whenever making an observation, we have a fixed cost $d > 0$. We want to choose a sequence of observation times, as well as a stopping time, so that when we stop, we can keep the probability of a "false alarm" as small as possible, and the delay in making a late declaration as short as possible, and we also want to make as few observations as possible since our cost increases with number of observations. We can thus write the value function below as U, which minimizes the summation of all three factors described. According to the reformulation in the previous section, we can write the value function as a function of the known parameter π:

Definition 3.2.1. *Value Function U*

$$U : [0,1] \to [0,1]$$

$$U(x) = \inf_{\hat{\tau}} \inf_{\tau \in \mathcal{S}^{\hat{\tau}}} \mathbb{E}_\pi[1 - \Pi_\tau + c\int_0^\tau \Pi_t dt + d\sum_{i=1}^{\infty} 1_{\{\tau_i \leq \tau\}}]$$

Where c represents the relative weight between a false alarm and a delay, or to interpret as a running cost of delay, and $d > 0$ *is the observation cost.*

Remark 3.2.1. *Again we do not ask the stopping time* τ *to be finite, in other words,* $\tau = \infty$ *can be an optimal stopping time. However, one will never make an observation at infinity.*

Therefore, in the detection problem, we can consider the cases where the optimal stopping time is finite and infinite separately. In the following sessions it is sufficient to consider when all stopping times are finite.

Remark 3.2.2. *From Definition 3.2.1 we can see that, U is bounded by* 0 *and above by* $1 - \pi$.

Remark 3.2.3. *We can see from Definition 3.2.1 that the filtration is generated by the discrete observation of* X_t*, while the value function is defined in terms of* Π_t*, which means that our information set in the value function comes from another process, which is not Markovian. In the later sections we will prove that this specific problem can be solved using the fixed point approach, for the process* Π_t *has the Markov property and is admissible to the filtration at previous observation time.*

Remark 3.2.4. *From Definition 3.2.1 we can see that we formulate the problem in the discrete case now, however, as we will see later the process Π_t is piecewise continuous. This is to keep the consistency with the main problem. Later we will see that this formulation indeed works, and give a general requirement for the continuous part of the underlying process.*

3.2.2 Properties of Π_t

We now study the properties of the posterior probability process Π_t. First let us define the conditional likelihood ratio process that θ has already appeared. For any sequence $\hat{\tau}$:

Definition 3.2.2. *Conditional Likelihood Ratio Process*

$$\Phi_t^{\hat{\tau}} = \frac{\mathbb{P}(\theta \leq t|\mathcal{G}_t^{\hat{\tau}})}{\mathbb{P}(\theta > t|\mathcal{G}_t^{\hat{\tau}})}$$
$$\Phi_0^{\hat{\tau}} = \frac{\pi}{1-\pi}$$

By Appendix A in [3], we can derive a recursive formula for the posterior process $\Phi^{\hat{\tau}}$ at any time t:

$$\Phi_t^{\hat{\tau}} = \begin{cases} j(\Delta\tau_n, \Phi_{\tau_{n-1}}^{\hat{\tau}}, \frac{\Delta X_{\tau_n}}{\sqrt{\Delta\tau_n}}) & \text{if } t = \tau_n \\ e^{\lambda(t-\tau_{n-1})}(\Phi_{\tau_{n-1}}^{\hat{\tau}} + 1) - 1 & \text{if } \tau_{n-1} \leq t < \tau_n \end{cases} \tag{3.3}$$

where $\Delta\tau_n = \tau_n - \tau_{n-1}$, $\Delta X_{\tau_n} = X_{\tau_n} - X_{\tau_{n-1}}$, and:

$$\begin{aligned} & j(\Delta t, \phi, z) \\ & = \exp\{\mu z\sqrt{\Delta t} + (\lambda - \frac{\mu^2}{2})\Delta t\}\phi + \int_0^{\Delta t} \lambda \exp\{(\lambda + \frac{\mu z}{\sqrt{\Delta t}})u - \frac{\mu^2 u^2}{2\Delta t}\}du \end{aligned}$$

Here we will follow the derivation of [3] to derive (3.3). Since we define this recursive relation for any sequence $\hat{\tau}$, it is easier to assume that the sequence is deterministic. Taking conditional expectations iteratively, the result of observing at stopping times follows directly.

Suppose $\tilde{X}$ is a standard Brownian motion on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, whose drift changes from 0 to μ at time $\tilde{\theta}$, where $\mathbb{P}(\tilde{\theta} = 0) = \pi$ and $\tilde{\mathbb{P}}(\tilde{\theta} \in dt|\tilde{\theta} > 0) = \lambda e^{-\lambda t}dt$. Let $0 = t_0 < t_1 < ...$ be the deterministic infinite sequence of observation times. Let:

$$L_t(u, x_0, x_1, \dots) := \prod_{l\geq 1, t_l \leq t} \frac{1}{\sqrt{2\pi(t_l - t_{l-1})}} t \exp\{\frac{(x_l - x_{l-1} - \mu(t_l - t_{l-1} \vee u)^+)^2}{2(t_l - t_{l-1})}\}$$

Then we have:

$$\tilde{\mathbb{P}}(\tilde{X} \in dx_l \text{ for all } l \geq 1, t_l \leq t) = L_t(\tilde{\theta}, x_0, x_1, \dots) \prod_{l\geq 1, t_l \leq t} dx_l$$

We can write the conditional likelihood of observations $\tilde{X}_{t_0}, \tilde{X}_{t_1}, \dots$ given $\tilde{\theta} = u$ as:

$$\begin{aligned} L_t(u) &:= L_t(u, \tilde{X}_{t_0}, \tilde{X}_{t_1}, \dots) \\ &= \prod_{l\geq 1, t_l \leq t} \frac{1}{\sqrt{2\pi(t_l - t_{l-1})}} t \exp\{\frac{(\tilde{X}_{t_l} - \tilde{X}_{t_{l-1}} - \mu(t_l - t_{l-1} \vee u)^+)^2}{2(t_l - t_{l-1})}\} \end{aligned}$$

This process $\tilde{X}$ can be constructed by a change of measure from another process X which is a standard Brownian motion on another probability space $(\Omega, \mathcal{F}, \mathbb{P})$, i.e. the drift never changes. Let us assume there is another random variable θ on the same probability space with $\mathbb{P}(\theta = 0) = \pi$, and $\mathbb{P}(\theta \in dt|\theta > 0) = \lambda e^{-\lambda t}$. Then we can write:

$$\begin{aligned}\mathbb{P}(X \in dx_l \text{ for all } l \geq 1, t_l \leq t) &= L_t(\infty, x_0, x_1, \dots) \prod_{l\geq 1, t_l \leq t} dx_l \\ &= \prod_{l\geq 1, t_l\leq t} \frac{1}{\sqrt{2\pi(t_l - t_{l-1})}} t \exp\{\frac{(x_l - x_{l-1})^2}{2(t_l - t_{l-1})}\}\end{aligned}$$

since the drift stays 0.

Let $\mathbb{F} = (\mathcal{F}_t)_{tgeq0}$ be the discrete time filtration generated by observations of X at the infinite sequence $0 = t_0 < t_1 < ...$, and let $\mathbb{G} = (\mathcal{G}_t)_{tgeq0}$, where $\mathcal{G}_t = \mathcal{F}_t \wedge \sigma(\theta)$.

Define $\tilde{\mathbb{P}}$ on $\mathcal{G}_\infty$ by:

$$\begin{aligned}\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}} &= Z_t(\theta) := \frac{L_t(\theta)}{L_t(\infty)} \\ &= \exp\{\sum_{l=1}^{\infty} 1_{\{t_l \leq t\}} (\frac{(X_{t_l} - X_{t_{l-1}})\mu(t_l - \theta \vee t_{l-1})^+}{t_l - t_{l-1}} - \frac{\mu^2((t_l - \theta \vee t_{l-1})^+)^2}{2(t_l - t_{l-1})})\}\end{aligned}$$

So we can see that, under $\tilde{\mathbb{P}}$, conditioning on θ, $X_{t_l} - X_{t_{l-1}}$ for $l \geq 1$ are independent, and normally distributed with mean $\mu(t_l - \theta \vee t_{l-1})^+$ and variance $t_l - t_{l-1}$.

We can easily see from the expression that $Z_0(\theta) = 1$, and thus $\tilde{\mathbb{P}}$ and $\mathbb{P}$ are identical on $\mathcal{G}_0$, so θ has the same distribution on $\tilde{\mathbb{P}}$ and $\mathbb{P}$. Under $\tilde{\mathbb{P}}$, X has the distribution of a standard Brownian motion whose drift changes from 0 to μ at θ.

By Bayes' theorem, we define:

$$\begin{aligned}\Phi_t &:= \frac{\tilde{\mathbb{P}}(\theta \leq t|\mathcal{F}_t)}{\tilde{\mathbb{P}}(\theta > t|\mathcal{F}_t)} \\ &= \frac{\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}|\mathcal{F}_t]}{\mathbb{E}^P[Z_t(\theta)1_{\{\theta>t\}}|\mathcal{F}_t]}\end{aligned}$$

On the set $\{\theta > t\}$, we have $(t_l - \theta \vee t_{l-1})^+ = (t_l - \theta)^+ = 0$ for all $l \geq 1, t_l \leq t$. So $Z_t(\theta)1_{\{\theta>t\}} = 1_{\{\theta>t\}}$. So:

$$\begin{aligned}\mathbb{E}^P[Z_t(\theta)1_{\{\theta>t\}}|\mathcal{F}_t] &= \mathbb{P}(\theta > t|\mathcal{F}_t) \\ &= \mathbb{P}(\theta > t) \\ &= (1 - \pi)e^{-\lambda t}\end{aligned}$$

So:

$$\begin{aligned}\Phi_t &= \frac{\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}|\mathcal{F}_t]}{\mathbb{E}^P[Z_t(\theta)1_{\{\theta>t\}}|\mathcal{F}_t]} \\ &= \frac{\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}|\mathcal{F}_t]}{(1-\pi)e^{-\lambda t}} \\ &= \frac{e^{\lambda t}}{1-\pi}\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}]\end{aligned}$$

We can write:

$$\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}] = \pi Z_t(0) + (1-\pi)\int_0^t \lambda e^{-\lambda t} Z_t(u)du$$

Let us assume that $t_{n-1} \leq t < t_n$ for some $n \geq 1$. Then $Z_t(u) = Z_{t_{n-1}}(u)$ for all u, $Z_{t_{n-1}(u)} = 1$ for all $t_{n-1} \leq u < t_n$.

So we have:

$$\begin{aligned}\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}] &= \pi Z_{t_{n-1}}(0) + (1-\pi)\int_0^t \lambda e^{-\lambda t} Z_{tn-1}(u)du \\ &= \pi Z_{t_{n-1}}(0) + (1-\pi)\int_0^{t_{n-1}} \lambda e^{-\lambda t} Z_{tn-1}(u)du + (1-\pi)\int_{t_{n-1}}^t \lambda e^{-\lambda t} Z_{tn-1}(u)du \\ &= \frac{1-\pi}{e^{\lambda t_{n-1}}}\Phi_{t_{n-1}} + (1-\pi)(e^{-\lambda t_{n-1}} - e^{-\lambda t})\end{aligned}$$

It follows that:

$$\begin{aligned}\Phi_t &= \frac{e^{\lambda t}}{1-\pi}\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}] \\ &= e^{\lambda(t-t_{n-1})}\Phi_{t_{n-1}} + e^{\lambda(t-t_{n-1})} - 1 \\ &= e^{\lambda(t-t_{n-1})}(\Phi_{t_{n-1}} + 1) - 1\end{aligned}$$

which would be the case where $t_{n-1} \leq t < t_n$ in (3.3).

Now we will derive Φ_{t_n} conditioning on $\Phi_{t_{n-1}}$. For $t_{n-1} \leq u$, we have $Z_{t_{n-1}}(u) = 1$, so:

$$\begin{aligned}&Z_{t_n}(u) \\ =&Z_{t_{n-1}}(u)\exp\{\frac{(X_{t_n} - X_{t_{n-1}})\mu(t_n - u\vee t_{n-1})^+}{t_n - t_{n-1}} - \frac{\mu^2((t_n - u\vee t_{n-1})^+)^2}{2(t_n - t_{n-1})})\}\end{aligned}$$

So we can write:

$$\begin{aligned}\Phi_t =&\frac{e^{\lambda t}}{1-\pi}\mathbb{E}^P[Z_t(\theta)1_{\{\theta\leq t\}}] \\ =&\frac{e^{\lambda t_n}}{1-\pi}[(\pi Z_{t_{n-1}}(0) + (1-\pi)\int_0^{t_{n-1}} \lambda e^{-\lambda u} Z_{t_{n-1}}(u)du)\exp{(X_{t_n} - X_{t_{n-1}})\mu - \frac{\mu^2}{2}(t_n - t_{n-1}))} \\ &+ (1-\pi)\int_{t_{n-1}}^{t_n} \lambda e^{\lambda u} Z_{t_{n-1}}(u)\exp\{\frac{(X_{t_n} - X_{t_{n-1}})\mu(t_n - u)}{t_n - t_{n-1}}) - \frac{\mu^2(t_n-u)^2}{2(t_n - t_{n-1})}du\}] \\ =&\exp\{(X_{t_n} - X_{t_{n-1}})\mu - \frac{\mu^2}{2}(t_n - t_{n-1})\}e^{\lambda(t_n - t_{n-1})}\Phi_{t_{n-1}} \\ &+ \int_{t_{n-1}}^{t_n} \lambda e^{\lambda(t_n - u)}\exp\{\{\frac{(X_{t_n} - X_{t_{n-1}})\mu(t_n - u)}{t_n - t_{n-1})} - \frac{\mu^2(t_n-u)^2}{2(t_n - t_{n-1})}\}du\end{aligned}$$

Apply a change of variables $w = u - t_n$, we can write:

$$\begin{aligned}\Phi_t =&\exp\{(X_{t_n} - X_{t_{n-1}})\mu - \frac{\mu^2}{2}(t_n - t_{n-1})\}e^{\lambda(t_n - t_{n-1})}\Phi_{t_{n-1}} \\ &+ \int_{t_n - t_{n-1}}^0 \lambda e^{-\lambda w}\exp\{\{\frac{(X_{t_n} - X_{t_{n-1}})\mu(-w)}{t_n - t_{n-1})} - \frac{\mu^2 w^2}{2(t_n - t_{n-1})}\}du \\ =&\exp\{(X_{t_n} - X_{t_{n-1}})\mu - \frac{\mu^2}{2}(t_n - t_{n-1})\}e^{\lambda(t_n - t_{n-1})}\Phi_{t_{n-1}} \\ &+ \int_0^{t_n - t_{n-1}} \lambda\exp\{\{(\lambda + \frac{(X_{t_n} - X_{t_{n-1}})\mu}{t_n - t_{n-1})})w - \frac{\mu^2 w^2}{2(t_n - t_{n-1})}\}du\end{aligned}$$

which would be the case where t coincides with the observation times in (3.3).

Lemma 3.2.1. *For any sequence $\hat{\tau}$, the posterior probability process $\Pi_t^{\hat{\tau}}$ is piecewise deterministic Markovian.*

Proof. By (3.3) we see that the current value of the process $\Pi_{\tau_k}^{\hat{\tau}}$, at τ_i's, depends only on the previous $\Pi_{\tau_{k-1}}^{\hat{\tau}}$, the increment in $\hat{\tau}$ and the increment in X. It does not depend on the path coming before $\tau_k - 1$. Between τ_i's, the evolution of the process is exponential. Therefore, the process has random jump points, between which it is deterministic. So it is a piecewise deterministic Markov process.

Note that this process is not time-homogeneous.

Lemma 3.2.2. *Define a discrete-time filtration $\mathcal{F}^{\hat{\tau}}$ indexed by k, by setting $\mathcal{F}_k^{\hat{\tau}} = \mathcal{G}_{\tau_k}^{\hat{\tau}}$. For any sequence $\hat{\tau}$, $\forall \tau_k \in \hat{\tau}$, $\Pi_{\tau_k}^{\hat{\tau}} \in \mathcal{F}_k^{\hat{\tau}}$. Furthermore, $\Pi_t^{\hat{\tau}} \in \mathcal{F}_k^{\hat{\tau}}$, for all $t < \tau_{k+1}$.*

Proof. By (3.3), we can see that every $\Pi_{\tau_k}^{\hat{\tau}}$ can be written as:

$$\Pi_{\tau_k}^{\hat{\tau}} = f(\Pi_{\tau_{k-1}}^{\hat{\tau}}, \tau_k - \tau_{k-1}, X_{\tau_k} - X_{\tau_{k-1}})$$

where f is a function that can be easily derived from (3.3). We see that $\Pi_{\tau_{k-1}}^{\hat{\tau}} \in \mathcal{F}_{k-1}^{\hat{\tau}}$, $\tau_k - \tau_{k-1} \in \mathcal{F}_{k-1}^{\hat{\tau}}$, and $X_{\tau_k} - X_{\tau_{k-1}} \in \mathcal{F}_k^{\hat{\tau}}$. Therefore $\Pi_{\tau_k}^{\hat{\tau}} \in \mathcal{F}_k^{\hat{\tau}}$.
Based on the filtration $\mathcal{F}^{\hat{\tau}}$, we will also know what τ_{k+1} is. Since the evolution of Π between τ_k and τ_{k+1} is deterministic, we will also have full information to determine $\Pi_t, t < \tau_{k+1}$.

3.2.3 A Fixed Point Approach

Let us prove in this section that we can use a fixed point approach to find the value function of the detection problem with discrete costly observations, and provide an optimal strategy. Furthermore, we can provide the value function and an optimal strategy when stopping between observation times is allowed. Also when the number of observation is restricted, we can use the fixed point approach with limited iterations to solve the problem.

Let us first define a set of functions L on which we will define our operator:

Definition 3.2.3. *Set L*

$$L := \{l \ \ Borel-measurable : [0,1] \to [0,1], \ \ 0 \le l(\pi) \le 1 - \pi\}$$

Let us now introduce an operator $\mathcal{L}$ to characterize our choices as described above:

Definition 3.2.4. *Operator $\mathcal{L}$*
Let $\mathcal{L}$ be an operator acting on L, such that $\forall l \in L$:

$$(\mathcal{L}l)(\pi) = \min(1-\pi, \ \inf_{t\ge 0} \mathbb{E}_\pi[l(\Pi_t) + c\int_0^t \Pi_u du + d])$$

Let us assume that $\hat{U}$ is a fixed point of the operator $\mathcal{L}$, i.e. $\hat{U} = \mathcal{L}\hat{U}$. We can thus write:

$$\hat{U}(\pi) = \min(1-\pi, \ \inf_{t\geq 0} \mathbb{E}_\pi[\hat{U}(\Pi_t) + c\int_0^t \Pi_u du + d])$$

Similarly, we will now prove that if the fixed point $\hat{U}$ exists, it equals the value function U. Furthermore, it is also the unique fixed point. We will then provide an associated optimal strategy. Finally, we will prove the existence of the fixed point and how to find it.

Let us first define the continuation and stopping set, and a sequence of specific observation times $\hat{\tau}^* = \tau_1^*, \tau_2^*, ...$, which we will prove to be an optimal strategy:

Definition 3.2.5. *Continuation and Stopping set*

$$C := \{\pi : \hat{U}(\pi) < 1-\pi\}$$
$$D := \{\pi : \pi \notin C\}$$

Definition 3.2.6. *Observation Times Define the observation times by the following recursive construction:*

$$t^*(x) = \begin{cases} \inf\{t : \hat{U}(x) = \mathbb{E}_\pi[\hat{U}(\Pi_t) + c\int_0^t \Pi_u du + d]\}, & \pi \in C \\ \infty, & \pi \in D \end{cases}$$

Define:

$$\tau_0^* := 0$$
$$\tau_{k+1}^* := \tau_k^* + t^*(\Pi_{\tau_k^*}), \quad k \geq 0$$

Let $m \in 0, 1, ...$ be the index of the last observation time, i.e.:

$$m = \max\{k : \tau_k^* < \infty\}$$
$$i.e. \quad \tau_m^* < \infty, \quad \tau_{m+1}^* = \infty, ...$$

We will prove the uniqueness of the fixed point by the following claims:

Claim 3.2.3. *Let $\hat{\tau}$ be any sequence of stopping times, then $k \to \hat{U}(\Pi_{\tau_k}) + c\int_0^{\tau_k} \Pi_u du + dk$ is a submartingale with respect to $\mathcal{F}_k^{\hat{\tau}}$.*

Claim 3.2.4. *Let $\hat{\tau}^*, m$ be defined as in Definition 3.2.6, then $k \to \hat{U}(\Pi_{\tau^*_{k\wedge m}}) + c\int_0^{\tau^*_{k\wedge m}} \Pi_u du + d\sum_{1\leq i\leq (k\wedge m)}$ is a martingale with respect to $\mathcal{F}_k^{\hat{\tau}^*}$.*

Claim 3.2.5. *$\hat{U} = U$, $\hat{U}$ is the unique fixed point of $\mathcal{L}$.*

Remark 3.2.5. *From the proof we can see that, we do not require the process which we observe to be Markovian, for any optimal stopping problem with discrete observations, whether it is to maximize or to minimize some payoff functions, if satisfying the following criterias, we can use a fixed point approach and the super/sub-margingale argument to find its unique fixed point:*

For any sequence $\hat{\tau}$:

(i) The underlying process $\Pi^{\hat{\tau}}$ *formulated in the value function is Markovian;*

(ii) The process X_{τ_k} *which we observe, along with the sequence* $\tau_1, \tau_2, ..., \tau_k$ *generates a filtration* $\mathcal{F}_k^{\hat{\tau}}$*,* $\forall k$*,* $\Pi_{\tau_k}^{\hat{\tau}} \in \mathcal{F}_k^{\hat{\tau}}$*.*

As we have already proven in Section 2.5, when stopping at anytime is allowed, the case is even simpler here. We know that the posterio probability process Π_t is deterministic until the next observation, therefore the optimal strategy for declaring the disorder time between the observation times can be naturally defined.

Chapter 4

Conclusion

What we have done:

In this project we study the optimal stopping problem where the underlying process is continuous, but one can only make discrete observations with constant costs. We further assume that the observation time points are not deterministic, one can choose the random times at which one observes to optimize the value function.

We first formulate this problem under the assumption that the underlying process is Markovian, and one can only stop the process at the observation time points. We then define an operator associated with this problem, characterizing our two choices after making each observation: to stop or to continue. By defining a specific sequence of observation times, we prove that the fixed point of this operator equals the value function, and it is unique, assuming the existence of the fixed point. We further conclude that this sequence of observation times provides an optimal strategy. We then prove the existence of the fixed point, and a constructive way to find it: the iterative procedure.

We discuss a specific example where the gain function is the payoff function of a put option, which one can naturally compare with a perpetual American put option under continuous information. We discuss the structure of the stopping set and the continuation set, and prove that it depends on the choice of the set of parameters. We further prove that in this example, the sequence formed by starting the iteration from the lower bound has the contraction property, and this sequence converges uniformly to the fixed point, and the rate of convergence is exponential.

We prove that in the sequence constructed by iterating from the lower bound, each element equals the value function where one can make at most finitely many observation.

We also prove that when one can observe the underlying process discretely but can stop the process at any time, it reduces to the previous case where one can only stop the process at the observation times.

We formulate the quickest detection problem under discrete costly observations, and discuss the property of the underlying process Π_t. We can easily see that it can also be solved by our approach.

Why this is meaningful:

Under our setting one can observe a continuous process discretely, but one can choose when to observe. This problem can have various applications. Here we provide 2 possible candidates:

We can for example take a look at real options. Assuming we want to purchase a firm in the future whose value cannot be observed directly from the market, so as to maximize its discounted value. We would need to make an investigation whenever we want to know its current value, and the investigation costs. We want to make the decision whether to purchase the company now, or make another observation. This corresponds to our problem where one can only stop the process at the observation points.

We can also look at the fishing problem. Assuming we have the ability to catch all of the fishes in a pool, and the amount of fish follows some Markov process. We cannot observe the amount of the fishes continuously, but we can make investigations which is costly. We make two decisions: the time points at which we make those investigations and the time point at which we fish. This corresponds to our problem where one can stop the process at any time.

What can be further studied:

We are interested in the property of the value function in the detection problem, is the fixed point concave? What does the structure of the continuation and the stopping set look like? We also want to compare the optimal strategy with the existing literature.

We would also like to discuss the case where τ takes the value ∞ and if it changes the sufficient conditions.

Bibliography

[1] Goran Peskir & Albert Shiryaev (2006), *Optimal Stopping and Free-Boundary Problems*, Birkhäuser, Basel

[2] Hannah Dyrssen & Erik Ekström (2018), *Sequential testing of a Wiener process with costly observations*, Sequential Analysis, 37:1, 47-58, DOI:10.1080/07474946.2018.1427973

[3] Erhan Bayraktar & Ross Kravitz (2015). *Quickest Detection with Discretely Controlled Observations*, Sequential Analysis, 34:1, 77-133, DOI: 10.1080/07474946.2015.995993