

Stopping in Continuous Time

Duarte Gonçalves

University College London

Topics in Economic Theory

Buying a Product

Suppose you are considering buying a laptop.

Read another review? Try it in a shop? Ask someone who has one?

Each tells you something about whether it is right for you.

At any point, you can buy, walk away, or keep looking.

Buying now may be a mistake.

Walking away now may also be a mistake.

Finding out more takes time.

When is it worth waiting?

Same question for investment, entry, exit, and filing a patent.

Time is often observed, data that comes for free. What can we make of it?

Are there patterns that informative about market forces or private information?

Need theory to make sense of the data.

This Lecture

We will cover most of technical tools you'll need in this lecture.

This seems daunting, but it will pay off:

By the end of it, it'll be as if you're able to speak a new language and have gained theory superpowers.

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

Overview

1. Buying a Product, or An Introduction to Brownian Motion

- Random Walk Limit
- Processes and Information
- Brownian Motion
- Properties of the Brownian Motion

2. Stopping Times and Martingales

3. Exit Probability and Expected Stopping Time

4. Stochastic Calculus

5. Itô's Formula

6. Generators and Dynkin's Formula

7. DPP, HJB, and Optimal Stopping

8. Optimal Search, Optimal Pricing, and Comparative Statics

9. Stopping Problems: A Working Method

Product Information

Setup (Branco, Sun, & Villas-Boas 2012 MnSc)

Expected net utility x : expected valuation minus price.

Outside option: 0 . Risk-neutral DM.

Product has many attributes; each is good or bad for the consumer.

Each new attribute raises or lowers expected utility.

Cost $c > 0$ per unit of time spent inspecting attributes.

Stop and choose whether to buy: buy if $x \geq 0$, walk away otherwise.

Scrap value: $v(x) := \max\{x, 0\} = x^+$.

Continue: pay flow cost $c > 0$ to inspect another attribute.

Choice conditional on stopping is easy. When to stop is not.

Inspecting Attributes

One attribute every Δ units of time; independent $\xi_k \in \{-1, 1\}$, equally likely.

$$X_t^\Delta = x + \sigma\sqrt{\Delta} \sum_{k=1}^{\lfloor t/\Delta \rfloor} \xi_k.$$

Each attribute has small impact on overall valuation.

X_t^Δ is just a random walk.

At $s < t$,

$$\mathbb{E}[X_t^\Delta - X_s^\Delta] = 0, \quad \text{Var}(X_t^\Delta - X_s^\Delta) = \sigma^2(t - s).$$

Why $\sqrt{\Delta}$? Increments of order Δ would make the variance vanish when looking at continuous time limit with $\Delta \downarrow 0$.

As $\Delta \downarrow 0$, the random walk converges in distribution to

$$X_t = x + \sigma B_t,$$

where B_t is a **Brownian motion**.

Definition

A **stochastic process** $X = (X_t)_{t \geq 0}$ is a family of random variables on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

Fix t : X_t is a random variable.

Fix ω : $t \mapsto X_t(\omega)$ is a **sample path**.

Important: continuous time does not mean continuous paths.

Brownian motion: continuous paths.

Poisson process: jumps at random times.

We'll mainly use continuous processes; **càdlàg** means right-continuous with left limits.

Information over Time

Definition

A **filtration** $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ satisfies $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ for $s \leq t$. A process X is **adapted** if X_t is $\mathcal{F}_t$ -measurable for every t .

$\mathcal{F}_t$: events whose occurrence can be determined by time t ; information available at t .

Natural filtration: $\mathcal{F}_t^X = \sigma(X_s : 0 \leq s \leq t)$.

Adapted: current state is observable; future innovations are not.

Throughout, augment filtrations to satisfy the **usual conditions**: completeness and right-continuity (cf. Pham 2009, ch. 1 for technical details).

Information matters: feasible decision under a filtration may be infeasible under another.

Brownian Motion

Definition

An adapted process B is a **standard Brownian motion** wrt $\mathbb{F}$ if:

$B_0 = \mathbf{0}$ a.s.; paths are continuous a.s.;

for $0 \leq s < t$, $B_t - B_s \sim N(\mathbf{0}, t - s)$ and is independent of $\mathcal{F}_s$.

$$B_t \sim N(\mathbf{0}, t); \text{Cov}(B_s, B_t) = \min\{s, t\}.$$


Increments are independent; levels are not.

For $X_t = x + \mu t + \sigma B_t$,

$$X_{t+h} \mid \mathcal{F}_t \sim N(X_t + \mu h, \sigma^2 h).$$

μ : drift per unit time; σ^2 : variance per unit time.

Brownian Motion



Brownian Motion: Useful Properties

Symmetry: $(-B_t)_{t \geq 0}$ is Brownian motion.

Time scaling: for $a > 0$, $(a^{-1/2}B_{at})_{t \geq 0}$ is Brownian motion.

Restarting: for fixed s , $(B_{s+t} - B_s)_{t \geq 0}$ is Brownian motion independent of $\mathcal{F}_s$.

Later: restarting also works at an a.s. finite *stopping time*.

Markov property: given B_s , the earlier path adds no information about future levels.

Economic use: the current state can summarise the history for future decisions/problems.

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
 - Stopping Time
 - Martingales
 - Martingale Convergence
 - Optional Stopping
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

Why Search If Expected Utility Does Not Change?

Suppose consumer starts off indifferent between buying and not, $x = 0$.

Then at any $t > 0$, $\mathbb{E}[X_t - X_0] = \mathbb{E}[(x + \sigma B_t) - (x + \sigma B_0)] = \sigma \mathbb{E}[B_t - B_0] = 0$.

The expected value of the product is the same as at the start!

This neglects the option value: decision changes depending on X_t .

Why Search If Expected Utility Does Not Change?

Suppose consumer searches for a fixed time T , then buys iff $X_T \geq 0$.

$$\mathbb{E}_0[v(X_T) - cT] = \sigma \sqrt{\frac{T}{2\pi}} - cT.$$

For small T , searching is strictly better than stopping.

Information is valuable because good news can induce a purchase; bad news can induce rejection.


Near indifference, the benefit is of order $\sqrt{T}$; the cost is of order T .

Best *fixed* search duration:

$$T^* = \frac{\sigma^2}{8\pi c^2}, \quad \mathbb{E}_0[v(X_{T^*}) - cT^*] = \frac{\sigma^2}{8\pi c}.$$

But why commit to the same duration? Should the DM not stop after a stream of just good (or just bad) news?

Responding to the Information



Good enough news: buy. Bad enough news: walk away.

Between the boundaries: keep learning.

Choose a rule that responds to discoveries, rather than a date in advance.

Choosing When to Stop

Definition

A random time $\tau : \Omega \rightarrow [0, \infty]$ is an $\mathbb{F}$ -**stopping time** if

$$\{\tau \leq t\} \in \mathcal{F}_t \quad \forall t \geq 0.$$

By time t , the DM knows whether stopping has occurred.

Examples: a fixed date; first time $X_t \geq b$; first exit from (a, b) .

Non-example: time of the highest value of B over $[0, 1]$.

Choosing a stopping time $\neq$ choosing a date in advance.

A **stopping rule** specifies what to do after every possible history.

Definition

$$\mathcal{F}_\tau := \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t \forall t \geq 0\}.$$

Describes information available when stopping occurs.

For adapted continuous X , X_τ is $\mathcal{F}_\tau$ -measurable on $\{\tau < \infty\}$.

If S, T are stopping times, so are $S \wedge T$, $S \vee T$, and $S + T$ (under the usual conditions).

For continuous adapted X , first hitting a closed set is a stopping time.

Martingales

Definition

An adapted, integrable process M is a **martingale** if

$$\mathbb{E}[M_t \mid \mathcal{F}_s] = M_s \quad (0 \leq s \leq t).$$

Replace $=$ by $\geq$ for a submartingale, or by $\leq$ for a supermartingale.

Brownian motion is a martingale: Independence of Brownian increments means future news is independent of $\mathcal{F}_t$.

So is $X_t = x + \sigma B_t$: no predictable improvement in expected utility.

Given current expected net utility X_t , future news satisfies

$$X_{t+h} = X_t + \sigma(B_{t+h} - B_t), \quad B_{t+h} - B_t \perp \mathcal{F}_t.$$

Markov property: two review histories ending at the same X_t give the same future problem.

But $v(X_t)$ is a **submartingale**:

$$\mathbb{E}[v(X_t) \mid \mathcal{F}_s] \geq v(\mathbb{E}[X_t \mid \mathcal{F}_s]) = v(X_s).$$

The Consumer's Stopping Problem

$X_t = x + \sigma B_t$; $c > 0$, $\sigma > 0$; no discounting for now.

Stopping times: $\mathbb{T}$.

Consumer's optimal stopping problem:

$$V(x) := \sup_{\tau \in \mathbb{T}} \mathbb{E}^x [v(X_\tau) - c\tau].$$

Benefit: further information can change the purchase decision.

Cost: searching takes time.

$V(x) \geq x^+$: stopping immediately is feasible.

Guess: walk away below a , buy above b , continue in between.

How can we evaluate this rule? How can we prove it is optimal?

Martingale Convergence

Theorem (Martingale Convergence)

Let M be a càdlàg martingale. M is uniformly integrable if and only if M_t converges a.s. and in L^1 to random variable M_∞ with $M_t = \mathbb{E}[M_\infty \mid \mathcal{F}_t] \forall t$.

Martingales converge to limiting distribution.

Uniform integrability of M : $\lim_{k \rightarrow \infty} \sup_t \mathbb{E}[|M_t| \mathbf{1}_{\{|M_t| \geq k\}}] = 0$.

Boundedness is sufficient for uniform integrability.

Optional Stopping

Theorem (Optional Stopping)

Let M be a càdlàg martingale and $S \leq T$ stopping times. If (1) T is a.s. finite and $M_{t \wedge T}$ is uniformly integrable, (2) M is square-integrable and $\mathbb{E}[M_T] < \infty$, or (3) $|M_{t \wedge T}| \leq C$ a.s. for all t , then

$$\mathbb{E}[M_T \mid \mathcal{F}_S] = M_S.$$

A very useful and common use: $M_{t \wedge T}$ is bounded, then $\mathbb{E}[M_T] = M_0$.

On the conditions:

Martingale property extends to information conditional on stopping.

Boundedness is sufficient for uniform integrability.

a.s. finite stopping alone **not** sufficient.

$T = \inf\{t > 0 : B_t = 1\}$ is a.s. finite, but $\mathbb{E}[B_T] = 1 \neq 0$.

Here $\mathbb{E}[T] = \infty$; stopped family not uniformly integrable.

Condition (2) will make sense later, it's here for future reference.

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
 - Exit Probability
 - Expected Stopping Time
 - In Pursuit of Optimality
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

Purchase Probability / First-Exit Probabilities

Characterise Choice Probabilities via Optional Stopping:

$$X_t = x + \sigma B_t, a < x < b, \tau_{a,b} = \inf\{t : X_t \notin (a, b)\}.$$

Brownian motion exits a bounded interval a.s.; continuous paths imply $X_{\tau_{a,b}} \in \{a, b\}$.

Stopped process $X_{t \wedge \tau_{a,b}}$ is bounded.

First-exit/hitting probabilities: $p(x) := \mathbb{P}^x(X_{\tau_{a,b}} = b)$. Then, by optional stopping,

$$x = \mathbb{E}^x[X_{\tau_{a,b}}] = p(x)b + (1 - p(x))a \implies p(x) = \frac{x - a}{b - a}.$$

In our model, if $a < 0 < b$, then consumer buys when stopping at $X_t = b$ and doesn't when stopping at $X_t = a$.

Exit probability is independent of σ for *fixed* boundaries.

Volatility changes how quickly the process gets there.

Expected Search Duration / Stopping Time

Characterise Expected Stopping Time via Optional Stopping:

$B_t^2 - t$ is a martingale. $B_{t+s} - B_t \perp\!\!\!\perp B_t$, hence

$$\begin{aligned} 0 &= \mathbb{E}[(B_{t+s} - B_t)B_t \mid \mathcal{F}_t] \\ &= \mathbb{E}[B_{t+s}B_t \mid \mathcal{F}_t] - B_t^2 \\ &= \text{Cov}(B_{t+s}, B_t) - B_t^2 \\ &= t - B_t^2. \end{aligned}$$

Hence: $(X_t - x)^2 - \sigma^2 t$ is a martingale.

Optional sampling at $\tau_{a,b} \wedge T$:

$$\sigma^2 \mathbb{E}^x[\tau_{a,b} \wedge T] = \mathbb{E}^x[(X_{\tau_{a,b} \wedge T} - x)^2] \leq \max\{(b - x)^2, (x - a)^2\}.$$

Monotone convergence gives $\mathbb{E}^x[\tau_{a,b}] < \infty$.

Let $T \rightarrow \infty$; bounded convergence on the right gives

$$\sigma^2 \mathbb{E}^x[\tau_{a,b}] = \mathbb{E}^x[(X_{\tau_{a,b}} - x)^2]$$

Expected Search Duration / Stopping Time

Characterise Expected Stopping Time via Optional Stopping:

$B_t^2 - t$ is a martingale.

Hence: $(X_t - x)^2 - \sigma^2 t$ is a martingale.

Optional sampling gives

$$\begin{aligned}\sigma^2 \mathbb{E}^x[\tau_{a,b}] &= \mathbb{E}^x[(X_{\tau_{a,b}} - x)^2] \\ &= p(x)(b - x)^2 + (1 - p(x))(a - x)^2 \\ &= \frac{(x - a)(b - x)^2 + (b - x)(a - x)^2}{b - a} \\ &= (x - a)(b - x).\end{aligned}$$

Expected search time: $\mathbb{E}^x[\tau_{a,b}] = \frac{(x - a)(b - x)}{\sigma^2}$

Increases the closer x is to $(b + a)/2$.

Decreases with σ , for fixed boundaries.

Greater σ implies realised BM path more volatile $\implies$ reach thresholds faster.

Evaluating a Stopping Time

Fix $a < 0 < b$ and stop when X exits (a, b) .

If $x \in (a, b)$, the expected payoff is

$$\begin{aligned}W_{a,b}(x) &= \mathbb{E}^x[V(X_{\tau_{a,b}})] - c\mathbb{E}^x[\tau_{a,b}] \\&= b\frac{x-a}{b-a} - \frac{c}{\sigma^2}(x-a)(b-x).\end{aligned}$$

Restrict to symmetric boundaries $a = -b$ and set $x = 0$:

$$W_{-b,b}(0) = \frac{b}{2} - \frac{cb^2}{\sigma^2}.$$

Maximising over $b > 0$ gives

$$b^* = \sigma^2/(4c) \text{ and } W_{-b^*,b^*}(0) = \frac{\sigma^2}{16c}.$$

More costly search: require less evidence before deciding.

More informative reviews: it is worth waiting for stronger evidence.

But Is Stopping Time Actually Optimal?

Stopping time τ_{-b^*, b^*} clearly feasible and optimal within policies with fixed symmetric thresholds.

Still missing: could an asymmetric or history-dependent rule do better?

Need a method that compares our candidate with every stopping time.

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
 - Stochastic Integration
 - Properties of the Itô Integral
 - Local Martingales
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

The General Stopping Problem

X_t Markovian, time-homogeneous process; discount rate $r \geq 0$.

Flow payoff $f(x)$ while continuing; terminal payoff / scrap value $v(x)$.

$$V(x) = \sup_{\tau \in \mathbb{T}} \mathbb{E}^x \left[\int_0^\tau e^{-rt} f(X_t) dt + e^{-r\tau} v(X_\tau) \right].$$

Continuation region: $C = \{x : V(x) > v(x)\}$.

Stopping region: $D = \{x : V(x) = v(x)\}$.

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

Existence and optimality require proof; sup need not be attained in general.

Want something like Bellman equation:


$$W(X_t) = \max\{v(X_t), f(X_t)\Delta + e^{-r\Delta} \mathbb{E}[W(X_{t+\Delta})]\}.$$

Characterise value function and derive implications for optimal (stopping) policy.

Suggestive of differential equation: $0 = \lim_{\Delta \downarrow 0} f(X_t) + e^{-r\Delta} \mathbb{E}[W(X_{t+\Delta}) - W(X_t)]/\Delta$.

But what exactly is this 'derivative' $\lim_{\Delta \downarrow 0} \mathbb{E}[W(X_{t+\Delta}) - W(X_t)]/\Delta$?

Ordinary Integration




Approximate $\int_0^T H_t dt$ by $\sum_i H_{t_i}(t_{i+1} - t_i)$.

Multiply a height by a time increment; add the rectangle areas.

Search costs are of this kind: $\int_0^\tau c dt = c\tau$.

Stochastic Integrating



Now use instead $\sum_i H_{t_i}(B_{t_{i+1}} - B_{t_i})$.

H_{t_i} is fixed before reading the next review; increment may have either sign.

In the consumer problem, $H_t = \sigma V'(X_t)$ is value's sensitivity to news.

Integrand Cannot Anticipate the Future

For simple process,

$$H_s = \sum_i H_i \mathbf{1}_{\{t_i < s \leq t_{i+1}\}}, \quad H_i \text{ is } \mathcal{F}_{t_i}\text{-measurable,}$$

define

$$\int_0^t H_s \, dB_s := \sum_i H_i (B_{t_{i+1} \wedge t} - B_{t_i \wedge t}).$$



H_i can depend on all past news, but not on the next innovation.

Predictable processes: measurable wrt σ -field generated by càdlàg adapted processes.

Continuous adapted integrands are predictable, e.g., $V'(X_t)$.

This is same information restriction as in choosing a feasible stopping rule: cannot anticipate the future.

Why the Endpoint Matters



Take $H = B$. The two sums differ by

$$\sum_i B_{t_{i+1}} \Delta B_i - \sum_i B_{t_i} \Delta B_i = \sum_i (\Delta B_i)^2.$$

Using right endpoint lets weight respond to very same increment it multiplies.

Why Ordinary Calculus Is Not Enough

For a smooth path x , $df(x_t) = f'(x_t) dx_t$.

BM paths a.s. not differentiable:

Heuristically: $dB_t = \sqrt{dt}\epsilon_t$, $\epsilon_t \sim N(0, 1)$.

$$\lim_{dt \downarrow 0} \left\| \frac{B_{t+dt} - B_t}{dt} \right\| = \lim_{dt \downarrow 0} |\epsilon_t|(dt)^{-1/2} = \infty \text{ a.s.}$$

Applying this blindly to $f(x) = x^2$ and Brownian motion would give

$$B_t^2 = 2 \int_0^t B_s dB_s = 2 \sum_i B_{t_i} (B_{t_{i+1}} - B_{t_i}).$$

But while a nonanticipating sum $\sum B_{t_i} (B_{t_{i+1}} - B_{t_i})$ has mean zero, and so

$$\mathbb{E}[2 \sum_i B_{t_i} (B_{t_{i+1}} - B_{t_i})] = 0,$$

we know $\mathbb{E}[B_t^2] = t$.

Something is missing.

Exact identity is instead:

$$B_t^2 = 2 \sum_k B_{t_i} \Delta B_i + \sum_i (\Delta B_i)^2.$$

Brownian squared increments do not disappear as the partition becomes finer.

Proposition (Brownian Quadratic Variation)

For deterministic partitions of $[0, t]$ with mesh $|\pi| \rightarrow 0$, $[B]_t := \lim_{|\pi| \downarrow 0} \sum_i (\Delta B_i)^2 = t$ in L^2 .

Proof

The sum has expectation t . Independence of increments gives

$$\text{Var} \left(\sum_i (\Delta B_i)^2 \right) = 2 \sum_i (\Delta t_i)^2 \leq 2t|\pi| \rightarrow 0.$$

Continuous finite-variation paths have zero quadratic variation.

Brownian paths have infinite total variation a.s.; ordinary Riemann–Stieltjes integration cannot cover general adapted integrands.

For the consumer's valuation, $[X]_t = \sigma^2 t$: accumulated variance of the news.

Equivalent definition: B_t is BM iff B is continuous martingale with $B_0 = 0$ and $[B]_t = t$.

Itô Integration

Start with $H_s = \sum_i H_i \mathbf{1}_{\{t_i < s \leq t_{i+1}\}}$, with H_i $\mathcal{F}_{t_i}$ -measurable.

Define

$$\int_0^t H_s \, dB_s := \sum_k H_k (B_{t_{i+1} \wedge t} - B_{t_i \wedge t}).$$

Integrand is chosen **before** next increment is observed.

BM does not have bounded variation (not Riemann-Stieltjes integral),

but Itô integral is well-defined in probability for predictable square integrable processes, $\mathbb{E} \int_0^T H_s^2 \, ds < \infty$.

Extend by mean-square limits to predictable H with $\mathbb{E} \left[\int_0^t H_s^2 \, ds \right] < \infty$.

Predictable: measurable wrt σ -field generated by left-continuous adapted processes.

In particular, continuous adapted integrands are predictable.

For a formal derivation of the Itô integral and its properties see Øksendal (2013) Ch. 3.

Properties of the Itô Integral

Proposition (Itô Isometry and Martingale Property)

If H is predictable square-integrable, then $M_t = \int_0^t H_s dB_s$, $t \leq T$, is square-integrable martingale and

$$\mathbb{E}[M_t] = 0, \quad \mathbb{E}[M_t^2] = \mathbb{E} \left[\int_0^t H_s^2 ds \right] = \mathbb{E}[M]_t.$$

Intuition

For simple predictable process H , write integral as $M_t = \sum_{t_i} H_{t_i} \Delta B_{t_i}$.

Then, $\mathbb{E}[M_{t+s} | \mathcal{F}_t] = \mathbb{E} \left[\sum_{t_i} H_{t_i} \Delta B_{t_i} | \mathcal{F}_t \right] = \sum_{t_i \leq t} H_{t_i} \Delta B_{t_i} + \sum_{t_i > t} H_{t_i} \mathbb{E}[\Delta B_{t_i} | \mathcal{F}_t] = 0$,
and then clearly: $\mathbb{E}[M_t] = \mathbb{E}[\sum_i H_i \Delta B_i] = \sum_i H_i \mathbb{E}[\Delta B_i] = 0$.

Moreover: $(\sum_i H_{t_i} \Delta B_{t_i})^2 = \sum_i H_{t_i}^2 (\Delta B_{t_i})^2 + \sum_{i \neq j} H_{t_i} H_{t_j} \Delta B_{t_i} \Delta B_{t_j}$.

Cross terms have zero expectation: for $t_i < t_j$, $\mathbb{E}[\Delta B_{t_i} \Delta B_{t_j}] = \mathbb{E}[\Delta B_{t_i} \mathbb{E}[\Delta B_{t_j} | \mathcal{F}_{t_i}]] = 0$.

Then, $\mathbb{E}[\sum_{i \neq j} H_{t_i} H_{t_j} \Delta B_{t_i} \Delta B_{t_j}] = \sum_{i \neq j} H_{t_i} H_{t_j} \mathbb{E}[\Delta B_{t_i} \Delta B_{t_j}] = 0$.

For diagonal terms, $\sum_i \mathbb{E}[H_{t_i}^2 (\Delta B_{t_i})^2] = \sum_i \mathbb{E}[H_{t_i}^2] \Delta t_i = \mathbb{E} \left[\int_0^t H_s^2 ds \right]$.

Properties of the Itô Integral

Proposition (Itô Isometry and Martingale Property)

If H is predictable square-integrable, then $M_t = \int_0^t H_s dB_s$, $t \leq T$, is a square-integrable martingale and

$$\mathbb{E}[M_t] = 0, \quad \mathbb{E}[M_t^2] = \mathbb{E} \left[\int_0^t H_s^2 ds \right] = \mathbb{E}[M]_t.$$

Isometry lets us pass to limits and define the integral consistently.

Variance of M_t accumulates at rate H_t^2 .

For any *deterministic* H , M_t is normally distributed with variance $\mathbb{E}[M]_t$.

Note: if $H_t := f(t, B_t)$ and $|f(t, x)| \leq C(1 + |x|^p) \forall x, t$, then H_t is predictable square-integrable.

Proposition (Properties of the Itô Integral)

For square-integrable predictable H and K , $M := \int H dB$ and $N := \int K dB$,

Linearity: $\int_a^b \alpha H_s + \beta K_s dB_s = \alpha \int_a^b H_s dB_s + \beta \int_a^b K_s dB_s$.

Additivity: $\int_a^b H_s dB_s = \int_a^c H_s dB_s + \int_c^b H_s dB_s$.

Unbiasedness: $\mathbb{E} \left[\int_0^t H_s dB_s \right] = 0$.

Continuous Martingale: $\int_0^t H_s dB_s = \mathbb{E} \left[\int_0^{t+h} H_s dB_s \mid \mathcal{F}_t \right]$.

Itô Isometry: $\mathbb{E}[M, N]_t = \mathbb{E} \left[\int_0^t H_s K_s dB_s \right] = \mathbb{E} \left[\int_0^t H_s dB_s \int_0^t K_s dB_s \right] = \mathbb{E}[M]_t \mathbb{E}[N]_t$.

Isometry will also allow us to keep track of multiple sources of uncertainty in learning and control problems.

Local Martingales

Definition

M is a **local martingale** if there are stopping times $T_n \uparrow \infty$ a.s. such that for each n $M_{t \wedge T_n}$ is a martingale.

If $\int_0^t H_s^2 ds < \infty$ a.s., Itô integral exists locally.

It need not have mean zero without an integrability condition.

Useful sufficient conditions:

bounded local martingale $\implies$ martingale;

$\mathbb{E} \int_0^t H_s^2 ds < \infty \implies$ square-integrable martingale;

nonnegative local martingale $\implies$ supermartingale.

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

Definition Itô Process

An **Itô process** has the following form, with μ, σ predictable,

$$X_t = X_0 + \int_0^t \mu_s ds + \int_0^t \sigma_s dB_s, \quad dX_t = \mu_t dt + \sigma_t dB_t.$$

Stochastic differential equation (SDE): $dX_t = \mu(t, X_t) dt + \sigma(t, X_t) dB_t, \quad X_0.$

Strong solution: solution constructed on a given filtered probability space carrying B :

$$X_t = X_0 + \int_0^t \mu(s, X_s) ds + \int_0^t \sigma(s, X_s) dB_s.$$

Standard sufficient conditions for existence and uniqueness of strong solution on $[0, T]$ with $\sup_{t \in [0, T]} \mathbb{E}|X_t|^2 < \infty$: uniform Lipschitz continuity in x and linear growth,

$$|\mu(t, x) - \mu(t, y)| + |\sigma(t, x) - \sigma(t, y)| \leq L|x - y|, \quad \forall x, y, t,$$

$$\mu(t, x)^2 + \sigma(t, x)^2 \leq C(1 + |x|^2), \quad \forall x, t,$$

and X_0 $\mathcal{F}_0$ -measurable and $\mathbb{E}|X_0|^2 < \infty$. We'll assume these conditions and μ, σ continuous in t , unless otherwise stated.

Theorem (Itô's Lemma)

If $dX_t = \mu_t dt + \sigma_t dB_t$ and $f \in C^{1,2}$, then

$$df(t, X_t) = \left(f_t + \mu_t f_x + \frac{1}{2} \sigma_t^2 f_{xx} \right) dt + \sigma_t f_x dB_t.$$

Derivatives are evaluated at (t, X_t) .

Intuition

2nd-order Taylor expand of f at (t, X_t) : $df = f_t dt + \frac{1}{2} f_{tt} (dt)^2 + f_x dX + \frac{1}{2} f_{xx} (dX)^2 + \dots$

As $dt \downarrow 0$, (dt) , $dt dB_t$, and higher-order terms vanish faster.

Replace $(dt)^2 = dt dB_t = \text{higher-order terms} = 0$ and $(dB_t)^2 = dt$ (quadratic variation of BM) and rearrange terms.

$$\implies df(t, X_t) = \left(f_t + \mu_t f_x + \frac{1}{2} \sigma_t^2 f_{xx} \right) dt + \sigma_t f_x dB_t.$$

Theorem (Itô's Lemma)

If $dX_t = \mu_t dt + \sigma_t dB_t$ and $f \in C^{1,2}$, then

$$df(t, X_t) = \left(f_t + \mu_t f_x + \frac{1}{2} \sigma_t^2 f_{xx} \right) dt + \sigma_t f_x dB_t.$$

Derivatives are evaluated at (t, X_t) .

Ordinary calculus misses $\frac{1}{2} \sigma_t^2 f_{xx}'' dt$ because it discards squared increments.

Itô calculus acknowledges $(dB_t)^2 = dt$.

Heuristically: $(dB)^2 = dt$, $dt dB = (dt)^2 = 0$.

These are bookkeeping rules for limits of sums, not actual identities.

Itô integration by parts: for $f \in C^{1,2}$, we can do

$$\int_0^t X_s dY_s = X_t Y_t - X_0 Y_0 - \int_0^t Y_s dX_s - [X, Y]_t$$

Theorem (Itô's Lemma)

If $dX_t = \mu_t dt + \sigma_t dB_t$ and $f \in C^{1,2}$, then

$$df(t, X_t) = \left(f'_t + \mu_t f'_x + \frac{1}{2} \sigma_t^2 f''_{xx} \right) dt + \sigma_t f'_x dB_t.$$

Derivatives are evaluated at (t, X_t) .

Itô's lemma $\approx$ chain rule of stochastic differentiation.

Drift term: predictable change due to time ($f'_t dt$), drift ($\mu_t f'_x dt$), and curvature ($\frac{1}{2} \sigma_t^2 f''_{xx} dt$).

Diffusion term: unpredictable changes due to Brownian innovations ($\sigma_t f'_x dB_t$).

Extends to multidimensional Itô processes, $dX_t = \mu_t dt + \Sigma_t dB_t$:

$$df = \left(f'_t + \nabla_x f \cdot \mu_t + \frac{1}{2} \text{tr}(\Sigma_t^T \nabla_{xx}^2 f \Sigma_t) \right) dt + \nabla_x f \Sigma_t dB_t.$$

For a formal treatment: Øksendal (2013, ch. 4).

Itô's Formula: Two Examples

Squared Brownian motion: $f(x) = x^2$,

$$d(B_t^2) = 2B_t dB_t + dt, \quad B_t^2 - t = 2 \int_0^t B_s dB_s.$$

Explains the martingale used to compute expected exit times.

Exponential martingale: $f(t, x) = e^{\lambda x - \lambda^2 t/2}$,

$$df(t, B_t) = \lambda f(t, B_t) dB_t.$$

Drift cancels.

Later: exponential martingales will change probabilities and describe likelihood ratios.

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
 - Infinitesimal Generator
 - Dynkin's Formula
 - Exit Problems as Differential Equations
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

Definition (Infinitesimal Generator)

The generator of a time-homogeneous Markov process for functions in its domain is

$$\mathcal{L}f(x) = \lim_{h \downarrow 0} \frac{\mathbb{E}^x[f(X_h)] - f(x)}{h}.$$

For an Itô process $dX_t = \mu(X_t) dt + \sigma(X_t) dB_t$ and $f \in \mathcal{C}^2$, it is given by

$$\mathcal{L}f(x) := \mu(x)f'(x) + \frac{1}{2}\sigma(x)^2f''(x).$$

Intuition

$$\begin{aligned}\mathcal{L}f(x) &= \lim_{dt \downarrow 0} \frac{\mathbb{E}^x[f(X_{dt})] - f(x)}{dt} = \lim_{dt \downarrow 0} \frac{\mathbb{E}^x[f(x) + (\mu(x)f'(x) + \frac{1}{2}\sigma^2(x)f''(x)) dt + \sigma f'(x) dB_t] - f(x)}{dt} \\ &= \mu(x)f'(x) + \frac{1}{2}\sigma^2(x)f''(x) + \lim_{dt \downarrow 0} \frac{\sigma(x)f'(x)\mathbb{E}^x[dB_t]}{dt} \\ &= \mu(x)f'(x) + \frac{1}{2}\sigma^2(x)f''(x)\end{aligned}$$

Definition (Infinitesimal Generator)

The generator of a time-homogeneous Markov process for functions in its domain is

$$\mathcal{L}f(x) = \lim_{h \downarrow 0} \frac{\mathbb{E}^x[f(X_h)] - f(x)}{h}.$$

For an Itô process $dX_t = \mu(X_t) dt + \sigma(X_t) dB_t$ and $f \in \mathcal{C}^2$, it is given by

$$\mathcal{L}f(x) := \mu(x)f'(x) + \frac{1}{2}\sigma(x)^2 f''(x).$$

Consumer problem: $X_t = x + \sigma B_t$. Hence, $\mathcal{L}V(x) = \frac{1}{2}\sigma^2 V''(x)$.

This will be a key concept for the differential characterisation of the value function, i.e., for Hamiltonian-Jacobi-Bellman equations (HJBs).

Definition (Infinitesimal Generator)

The generator of a time-homogeneous Markov process for functions in its domain is

$$\mathcal{L}f(x) = \lim_{h \downarrow 0} \frac{\mathbb{E}^x[f(X_h)] - f(x)}{h}.$$

For an Itô process $dX_t = \mu(X_t) dt + \sigma(X_t) dB_t$ and $f \in \mathcal{C}^2$, it is given by

$$\mathcal{L}f(x) := \mu(x)f'(x) + \frac{1}{2}\sigma(x)^2f''(x).$$

Extensions:

Discounting: replaces $\mathcal{L}$ by $(\mathcal{L} - r)$, e.g., $(\mathcal{L} - r)f(x) = \mu(x)f'(x) + \frac{1}{2}\sigma(x)^2f''(x) - rf(x)$.

Time-dependent coefficients: spatial infinitesimal generator

$$\mathcal{L}_t f(x) := \mu(t, x)f'(x) + \frac{1}{2}\sigma(t, x)^2f''(x).$$

Time-dependent function and coefficients: extended infinitesimal generator

$$(\partial_t + \mathcal{L}_t)f(t, x) := f'_t(t, x) + \mu(t, x)f'_x(t, x) + \frac{1}{2}\sigma(t, x)^2f''_{xx}(t, x).$$

References: Øksendal (2013, ch. 7); Pham (2009, ch. 1.3).

Theorem (Dynkin's Formula)

Let X_t be a time-homogeneous Itô process with $dX_t = \mu(X_t) dt + \sigma(X_t) dB_t$ and $X_0 = x$, $\tau \in \mathbb{T}$ a stopping time, $f \in \mathcal{C}^2$, and discount rate $r \geq 0$.

If (1) f is compactly supported or (2) for all t $|X_{t \wedge \tau}| \leq C$ a.s. and $\mathbb{E}[\tau] < \infty$, then,

$$\mathbb{E}^x[e^{-r\tau}f(X_\tau)] = f(x) + \mathbb{E}^x \int_0^\tau e^{-rs}(\mathcal{L} - r)f(X_s) ds.$$

Proof Sketch

Apply Itô to $e^{-rt}f(X_t)$ and take expectations. Stochastic integral disappears if it is a true martingale stopped at τ .

Reference: Øksendal (2013, chs. 7, 10.2).

Extensions: for potentially time-inhomogeneous coefficients $-\mu(t, x), \sigma(t, x)$ – and time-dependent functions $f \in \mathcal{C}^{1,2}$, under suitable conditions,

$$\mathbb{E}^{t,x}[e^{-r(\tau-t)}f(\tau, X_\tau)] = f(t, x) + \mathbb{E}^{t,x} \int_t^\tau e^{-r(s-t)}(\partial_s + \mathcal{L}_s - r)f(s, X_s) ds.$$

Exit Problems as Differential Equations

Back to the consumer problem:

Let X_t be a time-homogeneous Itô process and τ be first exit from (a, b) .

Exit probability: if $p(x) = \mathbb{P}^x(X_\tau = b)$, then by Dynkin's formula, for $x \in (a, b)$,

$$p(x) = \mathbb{E}^x[p(X_\tau \wedge h)] = p(x) + \mathbb{E}^x \left[\int_0^{\tau \wedge h} \mathcal{L}p(X_s) ds \right] \implies 0 = \mathbb{E}^x \left[\int_0^{\tau \wedge h} \mathcal{L}p(X_s) ds \right]$$

$$0 = \mathcal{L}p(x) = \lim_{h \downarrow 0} \frac{\mathbb{E}^x \left[\int_0^{\tau \wedge h} \mathcal{L}p(X_s) ds \right]}{h}$$

Moreover, at $p(b) = 1$ and at $p(a) = 0$.

p solves the following boundary-value problem:

$$\mathcal{L}p = 0 \text{ on } (a, b), \quad p(a) = 0, \quad p(b) = 1.$$

Exit Problems as Differential Equations

Back to the consumer problem:

Let X_t be a time-homogeneous Itô process and τ be first exit from (a, b) .

Exit probability: if $p(x) = \mathbb{P}^x(X_\tau = b)$, then p solves the boundary-value problem:

$$\mathcal{L}p = 0 \text{ on } (a, b), \quad p(a) = 0, \quad p(b) = 1.$$

Expected exit time: if $m(x) = \mathbb{E}^x[\tau]$, then by similar arguments, m solves the boundary-value problem:

$$\mathcal{L}m + 1 = 0 \text{ on } (a, b), \quad m(a) = m(b) = 0.$$

For $\mu = 0$, constant $\sigma > 0$:

$$\begin{aligned} \frac{1}{2}\sigma^2 p''(x) &= 0 \text{ on } (a, b), \quad p(a) = 0, \quad p(b) = 1 \implies p(x) = \frac{x-a}{b-a}. \\ \frac{1}{2}\sigma^2 m''(x) &= -1 \text{ on } (a, b), \quad m(a) = m(b) = 0 \implies m(x) = \frac{(x-a)(b-x)}{\sigma^2}. \end{aligned}$$

Dynkin verifies the solutions (these are $\mathcal{C}^2$); regularity is needed for the converse.

Exit Problems as Differential Equations

Logic extends to more complicated exit problems:

Let X_t be an Itô process and $\tau := \inf t > 0 \mid X_t \notin (\underline{x}_t, \bar{x}_t)$, where $\underline{x}_t, \bar{x}_t \in \mathcal{C}^1$.

Exit probability: if $p(t, x) = \mathbb{P}^{t,x}(X_\tau = \bar{x}_\tau)$, then

$$(\partial_t + \mathcal{L}_t)p = 0 \text{ on } (\underline{x}_t, \bar{x}_t), \quad p(t, \underline{x}) = 0, \quad p(t, \bar{x}) = 1.$$

Expected exit time: if $m(x) = \mathbb{E}^{t,x}[\tau]$, then

$$(\partial_t + \mathcal{L}_t)m = -1 \text{ on } (\underline{x}_t, \bar{x}_t), \quad m(t, \underline{x}) = m(t, \bar{x}) = 0.$$

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
 - Dynamic Programming
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

The General Stopping Problem

State X is a time-homogeneous Itô process; discount rate $r \geq 0$.

Flow payoff $f(x)$ while continuing; terminal payoff / scrap value $v(x)$.

$$V(x) := \sup_{\tau \in \mathbb{T}} \mathbb{E}^x \left[\int_0^\tau e^{-rt} f(X_t) dt + e^{-r\tau} v(X_\tau) \right]. \quad (\text{Value})$$

Continuation region: $C = \{x : V(x) > v(x)\}$.

Stopping region: $D = \{x : V(x) = v(x)\}$.

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

Want something like Bellman equation:

$$W(X_t) = \max\{v(X_t), f(X_t)\Delta + e^{-r\Delta} \mathbb{E}[W(X_{t+\Delta})]\}.$$

Characterise value function and derive implications for optimal (stopping) policy.

Suggestive of differential equation: $0 = \lim_{\Delta \downarrow 0} f(X_t) + e^{-r\Delta} \mathbb{E}[W(X_{t+\Delta}) - W(X_t)]/\Delta$.

This is exactly where the dynamic programming principle and HJBs come in, leveraging Itô's lemma and Dynkin's formula.

Theorem (Dynamic Programming Principle)

Let X_t be a time-homogeneous Itô process satisfying the usual conditions and $r \geq 0$. Then, for any stopping time $\theta \in \mathbb{T}$,

$$V(x) = \sup_{\tau} \mathbb{E}^x \left[\int_0^{\tau \wedge \theta} e^{-rs} f(X_s) ds + \mathbf{1}_{\{\tau < \theta\}} e^{-r\tau} V(X_\tau) + \mathbf{1}_{\{\tau \geq \theta\}} e^{-r\theta} V(X_\theta) \right].$$

When τ says to stop before θ , stop optimally; otherwise continue up to θ and, from then onward, continue optimally. This is decidedly Bellman-like.

If θ occurs before leaving the continuation region,

$$V(x) = \mathbb{E}^x \left[\int_0^{\theta} e^{-rs} f(X_s) ds + e^{-r\theta} V(X_\theta) \right].$$

References: Pham (2009, ch. 3.3, 5.2).

Extensions: More generally, for $dX_t = \mu(t, X_t) dt + \sigma(t, X_t) dB_t$ under suitable regularity for f, v (e.g., continuous and bounded/bounded growth), for any $\theta \in \mathbb{T}$ with $\theta \geq t$,

$$V(t, x) = \sup_{\tau > t} \mathbb{E}^{t,x} \left[\int_t^{\tau \wedge \theta} f(s, X_s) ds + \mathbf{1}_{\{\tau < \theta\}} v(\tau, X_\tau) + \mathbf{1}_{\{\tau \geq \theta\}} V(\theta, X_\theta) \right].$$

The Stopping Equation

Suppose V is smooth locally in the continuation region $C := \{x \mid V(x) > v(x)\}$.

Apply Dynkin's formula to the DPP for a short, locally stopped interval:

$$0 = \mathbb{E}^x \int_0^\theta e^{-rs} [f + (\mathcal{L} - r)V](X_s) ds.$$

Dividing by the expected interval length and taking limits yields an Hamiltonian-Jacobi-Bellman equation (HJB):

$$\mathcal{L}V - rV + f = 0 \quad \text{in } C.$$

At the boundary of C , ∂C , immediate stopping must beat waiting; by definition, $v = V$, a boundary condition (BC).

Theorem (Differential Characterisation of the Value Function 1)

Let X_t be a time-homogeneous Itô process satisfying the usual conditions and $r \geq 0$. Let V satisfy (Value), $C := \{x \mid V > v\}$, and $\tau := \inf\{t > 0 \mid X_t \notin C\}$. Suppose

- (1) $V \in \mathcal{C}^2(C) \cap C^0(\bar{C})$,
- (2) $\mathbb{E}^x[|V(X_\theta)| + \int_0^\theta |(\mathcal{L} - r)V(X_t)| dt] < \infty$ for any $x \in C$ and bounded $\theta \in \mathbb{T}$ s.t. $\theta \leq \tau$ a.s.,
- (3) $\tau < \infty$ a.s. for all $X_0 = x \in C$.

Then,

$$(\mathcal{L} - r)V + f = 0, \quad \text{on } C, \quad (\text{HJB})$$

$$V - v = 0, \quad \text{on } \partial C. \quad (\text{BC})$$

Variational inequality: combine both conditions: $\max\{v - V, (\mathcal{L} - r)V + f\} = 0$.

References: Øksendal (2013, ch. 10, 11.2), Pham (2008, ch. 3.4, 5.2).

Sufficient condition for (2): C bounded and the stated Lipschitz conditions on μ, σ .

Not helpful with our consumer problem, but we'll make use of it later on.

Theorem (Differential Characterisation of the Value Function 2)

Let X_t be a time-homogeneous Itô process satisfying the usual conditions and $r \geq 0$. Let V satisfy (Value), $C := \{x \mid V > v\}$, and $\tau := \inf\{t > 0 \mid X_t \notin C\}$. Further assume that f, v are Lipschitz continuous, $\sigma \geq \varepsilon > 0$, and $\tau < \infty$ a.s. for all $X_0 = x \in C$. Then $V \in \mathcal{C}^2(C) \cap \mathcal{C}^0$ and

$$(\mathcal{L} - r)V + f = 0, \quad \text{on } C, \quad (\text{HJB})$$

$$V - v = 0, \quad \text{on } \partial C. \quad (\text{BC})$$

Extensions: More generally, for $dX_t = \mu(t, X_t) dt + \sigma(t, X_t) dB_t$ under suitable regularity for f, v (e.g., bounded/bounded growth) and $V \in \mathcal{C}^{1,2}(C) \cap \mathcal{C}^0(\bar{C})$, we get

$$\max\{v - V, (\partial_t + \mathcal{L}_t - r)V + f\} = 0$$

with $\mathcal{L}_t = \mu(t, x)\partial_x + \frac{1}{2}\sigma(t, x)^2\partial_{xx}$.

Value Matching and Smooth Fit

Are we done? Not yet... We don't know what the continuation region C is:
it's a **free-boundary problem**.

(In all fairness, oftentimes we also don't really care what C or V are. Later, we'll learn how to extract comparative statics straight from the HJB.)

Value matching: at the boundary, assuming V is continuous, $V = v$. This is our boundary condition.

We need more conditions to pin-down the stopping boundary.

Smooth fit: under some regularity conditions, $V' = v'$ at ∂C .

Intuition: A small change in the state around the stopping threshold does not create a first-order gain from changing the stopping decision.

Theorem (Smooth-fit principle)

Let X_t be a time-homogeneous Itô process satisfying the usual conditions and $r \geq 0$. Let V satisfy (Value), $C := \{x \mid V > v\}$, and $\tau := \inf\{t > 0 \mid X_t \notin C\}$. Further assume that f, v are Lipschitz continuous, $\sigma \geq \varepsilon > 0$, and $\tau < \infty$ a.s. for all $X_0 = x \in C$. Then $V \in \mathcal{C}^2(C) \cap \mathcal{C}^0$ and

$$(\mathcal{L} - r)V + f = 0, \quad \text{on } C, \quad (\text{HJB})$$

$$V - v = 0, \quad \text{on } \partial C. \quad (\text{BC})$$

If, in addition, v is continuously differentiable in an open neighbourhood of ∂C , then

$$V' - v' = 0, \quad \text{on } \partial C. \quad (\text{SF})$$

Smooth fit is not automatic: jumps, degeneracy, or nonsmooth payoffs can invalidate it.

Ways to go about it:

Guess and verify. Assume HJB, BC, SF, and verify that it works.

Check if problem satisfies known conditions. See Peskir & Shiryaev (2009, ch. IV.9.1.), Pham (2009, ch. 5).

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

Back to the Consumer Problem

Primitives: $f = -c$, $v(x) = x^+$, $r = 0$, $\mathcal{L} = \frac{1}{2}\sigma^2\partial_{xx}$.

Optimal Stopping Problem: $V(x) := \sup_{\tau \in \mathbb{T}} \mathbb{E}^x [-c\tau + v(X_\tau)]$,

with the understanding that $\tau = \infty$ the payoff is $-\infty$.

Since $\mathbb{P}^x(\tau = \infty) > 0 \implies$ payoff is $-\infty$, any optimal stopping time must have $\tau < \infty$ a.s.

Guess $C = (a, b)$, with $a < 0 < b$.

$$\frac{1}{2}\sigma^2 V'' - c = 0 \implies W(x) = \frac{c}{\sigma^2}x^2 + K_1x + K_0 \quad \text{in } (a, b).$$

Four unknowns (K_0, K_1, a, b) ; four boundary conditions:

$$W(a) = 0, \quad W'(a) = 0, \quad W(b) = b, \quad W'(b) = 1.$$

Subtract derivative conditions: $2c(b - a)/\sigma^2 = 1$.

Value matching then gives $a = -b$.

$$b = \frac{\sigma^2}{4c}, \quad K_1 = \frac{1}{2}, \quad K_0 = \frac{\sigma^2}{16c}.$$

Unique function satisfying (HJB), (BC), and (SF) $\implies W = V$.

Proposition

Let $b = \sigma^2/(4c)$. Then, $\tau_{-b,b}$ solves the consumer's stopping problem.

Walk away when $X_t \leq -b$; buy when $X_t \geq b$. Continue otherwise.

Symmetric boundaries come from the driftlessness.

Purchase Probabilities and Search Duration

The optimal purchase probability and search duration are:

$$\mathbb{P}_x(\text{buy}) = \frac{x+b}{2b}, \quad \mathbb{E}_x[\tau^*] = \frac{b^2 - x^2}{\sigma^2}.$$

Higher c : narrower continuation region and less expected search.

Higher σ : faster movement, but wider optimal boundaries.

$$\mathbb{E}_0[\tau^*] = \frac{\sigma^2}{16c^2}.$$

Higher σ speeds up the same path distribution, *holding the thresholds fixed*, leading to shorter search duration.

But as volatility is greater, search is more valuable, so consumer expands continuation region.

Here optimal policy response dominates: faster information arrival produces endogenously *longer* search.

Sharp distinction between a change in the process from a change in the optimal rule.

Pricing When Consumers Can Search

Firm sets price $p \geq 0$ and faces marginal production cost k .

Initial net utility $x = u_0 - p$; demand is

$$D(p) = \begin{cases} 1, & p \leq u_0 - b, \\ (u_0 - p + b)/(2b), & u_0 - b < p < u_0 + b, \\ 0, & p \geq u_0 + b. \end{cases}$$

Seller maximises $(p - k)D(p)$ over $p \geq k$.

If $u_0 - k \leq -b$: no strictly profitable sale.

When $u_0 - k > -b$, concave problem, unique solution:

$$p^* = \begin{cases} (u_0 + k + b)/2, & -b < u_0 - k < 3b, \\ u_0 - b, & u_0 - k \geq 3b. \end{cases}$$

In the second case, price so that the consumer buys immediately.

Proposition

Let $u_0 - k \geq -b$. The optimal price p^* , purchase probability q^* , profit π^* , and consumer surplus u^* in the consumer problem with optimal pricing are given by

$$p^* := \begin{cases} (u_0 + k + b)/2, & -b \leq u_0 - k < 3b, \\ u_0 - b, & u_0 - k \geq 3b. \end{cases},$$

$$q^* := D(p^*),$$

$$\pi^* := (p^* - k)q^*,$$

$$u^* := q^*b - c \frac{b^2 - (u_0 - p^*)^2}{\sigma^2}$$

where $b = \sigma^2/(4c)$.

Who Benefits from Lower Search Costs?

Think about the emergence of a new technology that lowers the search costs.

What happens to prices and demand?

Who benefits from the lower search costs?

Who Benefits from Lower Search Costs?

Cheaper Search When Prior Expectations Are Favourable

Suppose $u_0 > k$: prior expected valuation exceeds production cost.

Very high search costs: $c \geq 3\sigma^2/[4(u_0 - k)]$.

Seller induces immediate purchase.

Cheaper search makes investigating more attractive: seller cuts the price to keep the consumer from searching.

Demand stays at one; price and profit fall; consumer surplus rises.

Who Benefits from Lower Search Costs?

Cheaper Search When Prior Expectations Are Favourable

Suppose $u_0 > k$: prior expected valuation exceeds production cost.

Very high search costs: $c \geq 3\sigma^2/[4(u_0 - k)]$.

Demand stays at one; price and profit fall; consumer surplus rises.

Intermediate search costs: $\sigma^2/[4(u_0 - k)] < c < 3\sigma^2/[4(u_0 - k)]$.

Consumer searches, but initially expects buying at the chosen price to be worthwhile.

Search can reveal reasons *not to buy*: cheaper search lowers demand at a fixed price.

Demand also becomes less sensitive to price, so seller raises the price.

Lost sales outweigh the higher markup; the price increase outweighs the consumer's direct benefit from cheaper search.

Demand falls; price rises; profit and consumer surplus fall.

Who Benefits from Lower Search Costs?

Cheaper Search When Prior Expectations Are Favourable

Suppose $u_0 > k$: prior expected valuation exceeds production cost.

Very high search costs: $c \geq 3\sigma^2/[4(u_0 - k)]$.

Demand stays at one; price and profit fall; consumer surplus rises.

Intermediate search costs: $\sigma^2/[4(u_0 - k)] < c < 3\sigma^2/[4(u_0 - k)]$.

Demand falls; price rises; profit and consumer surplus fall.

Very low search costs: $c < \sigma^2/[4(u_0 - k)]$.

Seller charges more than the consumer's prior expected valuation.

Buying initially looks unattractive at the chosen price.

Search is worthwhile because favourable discoveries can justify a purchase.

Cheaper search raises demand at a fixed price.

Seller optimally reprices upward and the purchase probability still falls.

The higher markup more than compensates for lost sales.

Direct benefit of cheaper search outweighs price increase for the consumer.

Demand falls; price, profit, and consumer surplus rise.

Who Benefits from Lower Search Costs?

Cheaper Search When Prior Expectations Are Unfavourable

Suppose $u_0 \leq k$: prior expected valuation is below production cost. Profitable trade *requires* favourable discoveries.

High search costs: $c \geq \sigma^2/[4(k - u_0)]$.

At any price above production cost, investigating is too costly: consumer walks away.

Cheaper search has no effect on trade until it crosses the threshold.

Demand, profit, and consumer surplus remain zero. The optimal price is not uniquely determined.

Who Benefits from Lower Search Costs?

Cheaper Search When Prior Expectations Are Unfavourable

Suppose $u_0 \leq k$: prior expected valuation is below production cost. Profitable trade *requires* favourable discoveries.

High search costs: $c \geq \sigma^2/[4(k - u_0)]$.

Demand, profit, and consumer surplus remain zero. The optimal price is not uniquely determined.

Low search costs: $0 < c < \sigma^2/[4(k - u_0)]$.

Consumer is willing to investigate an initially unattractive offer: favourable discoveries may make buying worthwhile.

Cheaper search encourages investigation before walking away, raising demand at a fixed price.

Seller raises the price, but this does not offset the increase in purchases or the consumer's direct benefit from cheaper search.

Demand, price, profit, and consumer surplus all rise.

Cheaper Search: Search Duration

High search costs: no search.

If $u_0 > k$ and $c \geq 3\sigma^2/[4(u_0 - k)]$, seller induces immediate purchase.

If $u_0 < k$ and $c \geq \sigma^2/[4(k - u_0)]$, consumer walks away immediately.

Expected search duration remains zero until search costs fall below the relevant threshold.

Once the consumer searches, cheaper search increases expected duration.

Further information is less costly: consumer requires stronger evidence before buying or walking away.

Seller raises the price, lowering initial expected net utility, moving the consumer towards rejection.

Nevertheless, both decision boundaries move farther from the consumer's initial expected net utility.

Consumer therefore takes longer, on average, to reach either decision.

Who Benefits from Lower Search Costs?

Think about the emergence of a new technology that lowers the search costs.

What happens to prices and demand?

Who benefits from the lower search costs?

Despite the very simple setting, we get rich and nuanced comparative statics.

Profits and consumer surplus can increase or decrease and don't react the same way to cheaper search.

Overview

1. Buying a Product, or An Introduction to Brownian Motion
2. Stopping Times and Martingales
3. Exit Probability and Expected Stopping Time
4. Stochastic Calculus
5. Itô's Formula
6. Generators and Dynkin's Formula
7. DPP, HJB, and Optimal Stopping
8. Optimal Search, Optimal Pricing, and Comparative Statics
9. Stopping Problems: A Working Method

Stopping Problems: A Working Method

1. Specify the state, information, feasible stopping rules, and payoffs.
2. Write the generator.
3. Get variational inequality (HJB) and (BC).
4. Use value matching and, when justified, smooth fit to characterise the continuation region.
5. Recover behaviour: exit probabilities, duration, comparative statics.

Next: Applications to market entry and exit and strategic stopping.

Stopping in Continuous Time

Duarte Gonçalves

University College London

Topics in Economic Theory

References: Economic Applications

Branco, Sun, & Villas-Boas (2012 MnSc), *Optimal Search for Product Information*.

References: Mathematical Background

Øksendal (2003), *Stochastic Differential Equations: An Introduction with Applications*, chs. 7 and 10.

Diffusions, generators, Dynkin's formula, and optimal stopping.

Pham (2009), *Continuous-time Stochastic Control and Optimization with Financial Applications*, chs. 1, 3, and 5.

Stochastic calculus, dynamic programming, and stopping.

Karatzas & Shreve (1998 printing), *Brownian Motion and Stochastic Calculus*, 2nd ed.

Further foundations and precise versions of the probabilistic results.

Further Results and Proofs

Versions, stopping-time information, and strong Markov property.

Reflection, time scaling, and drifted exit probabilities.

General verification and the Snell envelope.

Finite horizons and failure of smooth fit.

Modifications and Indistinguishability

X, Y are **modifications** if $\mathbb{P}(X_t = Y_t) = 1$ for each fixed t .

They are **indistinguishable** if $\mathbb{P}(X_t = Y_t \ \forall t \geq 0) = 1$.

Order of quantifiers matters: an uncountable union of null events need not be null.

Example: $U \sim U(0, 1)$, $X_t = \mathbf{1}_{\{t=U\}}$ and $Y_t = 0$.

At every fixed t , $X_t = Y_t$ a.s.; yet every path differs somewhere.

Continuous modifications are indistinguishable.

Proof: equality simultaneously at rational times, then continuity.

Restarting at a Random Time

Theorem (Strong Markov Property of Brownian Motion)

For an a.s. finite stopping time τ , the process

$$(B_{\tau+t} - B_{\tau})_{t \geq 0}$$

is a Brownian motion independent of $\mathcal{F}_{\tau}$.

Brownian motion restarts after hitting a boundary.

It does not restart after an arbitrary random time, such as a future maximum.

Supports dynamic programming after state-dependent decisions.

For strong Markov diffusion, new law depends on state X_{τ} .

Reflection and One-Sided Hitting Times

Let $\tau_a = \inf\{t \mid B_t = a\}$ for $a > 0$.

Reflect a path after its first hit of a : symmetry and strong Markov preserve its probability law.

$$\mathbb{P}(\tau_a \leq t) = \mathbb{P}\left(\max_{s \leq t} B_s \geq a\right) = 2\mathbb{P}(B_t \geq a) = 2\left[1 - \Phi\left(\frac{a}{\sqrt{t}}\right)\right].$$

Differentiate to obtain the density

$$f_{\tau_a}(t) = \frac{a}{\sqrt{2\pi t^3}} e^{-a^2/(2t)}, \quad t > 0.$$

$\tau_a < \infty$ a.s., but $\mathbb{E}\tau_a = \infty$.

A stopping rule can eventually succeed and still have an infinite expected waiting time.

Time Scaling

For $X_t = x + \sigma B_t$ and fixed boundaries $a < x < b$,

$$\tau_{a,b}^\sigma \stackrel{d}{=} \frac{1}{\sigma^2} \tau_{a,b}^1.$$

Multiplying volatility speeds up the same path distribution.

For the consumer problem, time change $s = \sigma^2 t$ gives

$$\sup_{\tau} \mathbb{E}[(x + \sigma B_\tau)^+ - c\tau] = \sup_{\rho} \mathbb{E}[(x + B_\rho)^+ - (c/\sigma^2)\rho].$$

Policies depend on c/σ^2 in Brownian time.

Calendar duration also divides by σ^2 ; remember both effects.

Exit Probabilities with Drift

$X_t = x + \mu t + \sigma B_t$; first exit τ from (a, b) .

For $\mu \neq 0$, solve $\mu p' + \frac{1}{2}\sigma^2 p'' = 0$ with $p(a) = 0, p(b) = 1$:

$$p(x) = \frac{1 - e^{-2\mu(x-a)/\sigma^2}}{1 - e^{-2\mu(b-a)/\sigma^2}}.$$

Wald's identity, justified by $\mathbb{E}\tau < \infty$, gives

$$\mathbb{E}^x \tau = \frac{a + (b - a)p(x) - x}{\mu}.$$

As $\mu \rightarrow 0$, recover $(x - a)/(b - a)$ and $(x - a)(b - x)/\sigma^2$.

Scale function converts drifted exit problem into linear one.

The Snell Envelope

Finite horizon T ; continuous adapted reward Y with $\mathbb{E} \sup_{t \leq T} |Y_t| < \infty$.

Define

$$S_t = \operatorname{ess\,sup}_{\tau \in \mathbb{T}, t \leq \tau \leq T} \mathbb{E}[Y_\tau \mid \mathcal{F}_t].$$

S is the smallest supermartingale dominating Y .

$\tau^* = \inf\{t \mid S_t = Y_t\}$ is optimal; $S_{t \wedge \tau^*}$ is a martingale.

For a Markov stopping problem,

$$Y_t = \int_0^t e^{-rs} f(X_s) ds + e^{-rt} v(X_t).$$

The free-boundary method is a concrete way to construct this payoff bound.

Finite Horizons

With a deadline T , time is part of the state.

$$V(t, x) = \sup_{t \leq \tau \leq T} \mathbb{E}_{t,x} \left[\int_t^\tau e^{-r(s-t)} f(s, X_s) ds + e^{-r(\tau-t)} v(\tau, X_\tau) \right].$$

Stopping equation becomes

$$\max\{g - V, V_t + \mathcal{L}V - rV + f\} = 0, \quad V(T, x) = v(T, x).$$

Stationary payoffs do not imply constant boundaries when time remaining changes.

At a deadline there is no option to acquire more information.

Example (A Degenerate State Process)

Let $X_t = x + t$, $r = 1$, and $v(x) = 1$ for $x < 0$, $v(x) = 2$ for $x \geq 0$. Then

$$V(x) = \begin{cases} \max\{1, 2e^x\}, & x < 0, \\ 2, & x \geq 0. \end{cases}$$

For $-\log 2 < x < 0$, wait until X reaches 0; otherwise stop immediately.

At $b = -\log 2$, V is continuous and v is locally smooth.

But $V'_-(b) = 0$ and $V'_+(b) = 1$.

No Brownian variation across the boundary to enforce tangency.

Later: viscosity solutions handle stopping equations without classical smoothness.