

Stopping in Continuous Time

When should an economic agent act?

Duarte Gonçalves

University College London

Topics in Economic Theory

5 October 2026

Buying a Product

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.

Buying a Product

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.

Buying a Product

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.

Overview

1. Buying a Product: Brownian Uncertainty
2. Stochastic Calculus: Valuing Continuation
3. Optimal Stopping: Boundaries and Verification
4. Investment: The Value of Waiting
5. Entry and Exit: Hysteresis
6. Strategic Stopping: Races and Contests

Overview

1. Buying a Product: Brownian Uncertainty
2. Stochastic Calculus: Valuing Continuation
3. Optimal Stopping: Boundaries and Verification
4. Investment: The Value of Waiting
5. Entry and Exit: Hysteresis
6. Strategic Stopping: Races and Contests

Product Information

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

Expected net utility x : expected valuation minus price.

Outside option: 0 . Risk-neutral DM.

Each new attribute raises or lowers expected utility.

Cost $c > 0$ per unit of time spent searching.

Product Information

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

Expected net utility x : expected valuation minus price.

Outside option: 0 . Risk-neutral DM.

Each new attribute raises or lowers expected utility.

Cost $c > 0$ per unit of time spent searching.

If the DM stops at state x , choose whether to buy:

$$g(x) := \max\{x, 0\} = x^+.$$

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

Small Discoveries

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 discovery is small, but there are many discoveries.

Small Discoveries

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 discovery is small, but there are many discoveries.

At grid times $s < t$,

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

Why $\sqrt{\Delta}$? Increments of order Δ would make the variance vanish.

Small Discoveries

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 discovery is small, but there are many discoveries.

At grid times $s < t$,

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

Why $\sqrt{\Delta}$? Increments of order Δ would make the variance vanish.

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

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

Finite attribute exhaustion is abstracted from in this stationary model.

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**.

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**.

Continuous time does not mean continuous paths.

Brownian motion: continuous paths.

Poisson process: jumps at random times.

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

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 .

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

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 .

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.

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

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.

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 Paths



A simulated path on a fine grid; Brownian paths have no instantaneous velocity.
Typical change over h is of order $\sqrt{h}$, not h .

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 $\mathcal{F}_s$.

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 $\mathcal{F}_s$.

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

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

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 $\mathcal{F}_s$.

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

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

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

This is a property of the model, not a property of all learning problems.

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) .

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) .

Nonexample: 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.

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.

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

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.

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

But $g(X_t) = X_t^+$ is a **submartingale**:

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

Learning is useful because the DM can change the eventual action.

The Consumer's Stopping Problem

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

Feasible rules: $\mathbb{T}_1 := \{\tau : \tau \text{ stopping time, } \mathbb{E}[\tau] < \infty\}$.

Definition (Value)

$$V(x) := \sup_{\tau \in \mathbb{T}_1} \mathbb{E}_x[X_\tau^+ - c\tau].$$

Benefit: further information can change the purchase decision.

Cost: searching takes time.

The Consumer's Stopping Problem

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

Feasible rules: $\mathbb{T}_1 := \{\tau : \tau \text{ stopping time, } \mathbb{E}[\tau] < \infty\}$.

Definition (Value)

$$V(x) := \sup_{\tau \in \mathbb{T}_1} \mathbb{E}_x[X_\tau^+ - c\tau].$$

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?

Theorem (Optional Sampling)

Let M be a càdlàg martingale and $S \leq T$ bounded stopping times. Then

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

The result extends to a.s. finite $S \leq T$ if the stopped process $M_{t \wedge T}$ is uniformly integrable.

A bounded stopped process is sufficient for uniform integrability.

Theorem (Optional Sampling)

Let M be a càdlàg martingale and $S \leq T$ bounded stopping times. Then

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

The result extends to a.s. finite $S \leq T$ if the stopped process $M_{t \wedge T}$ is uniformly integrable.

A bounded stopped process is sufficient for uniform integrability.

A.s. finite stopping alone is **not** sufficient.

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

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

First Exit: Which Boundary?

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

Brownian motion exits a bounded interval a.s.; continuity implies $X_{\tau_{a,b}} \in \{a,b\}$.

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

First Exit: Which Boundary?

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

Brownian motion exits a bounded interval a.s.; continuity implies $X_{\tau_{a,b}} \in \{a, b\}$.

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

Optional Sampling

Write $p = \mathbb{P}_x(X_{\tau_{a,b}} = b)$. Then

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

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

Volatility changes how quickly the process gets there.

First Exit: How Long?

$B_t^2 - t$ is a martingale: use independent, mean-zero increments.

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

First Exit: How Long?

$B_t^2 - t$ is a martingale: use independent, mean-zero increments.

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 $\mathbb{E}_x[\tau_{a,b}] < \infty$.

First Exit: How Long?

$B_t^2 - t$ is a martingale: use independent, mean-zero increments.

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

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

Largest in the middle; decreases with σ for fixed boundaries.

Evaluating a Stopping Rule

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

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

$$J_{a,b}(x) = b \frac{x-a}{b-a} - \frac{c}{\sigma^2} (x-a)(b-x).$$

Evaluating a Stopping Rule

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

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

$$J_{a,b}(x) = b \frac{x-a}{b-a} - \frac{c}{\sigma^2} (x-a)(b-x).$$

At $x = 0$, restrict temporarily to symmetric boundaries $a = -b$:

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

Maximising over $b > 0$ gives $b = \sigma^2/(4c)$.

Evaluating a Stopping Rule

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

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

$$J_{a,b}(x) = b \frac{x-a}{b-a} - \frac{c}{\sigma^2} (x-a)(b-x).$$

At $x = 0$, restrict temporarily to symmetric boundaries $a = -b$:

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

Maximising over $b > 0$ gives $b = \sigma^2/(4c)$.

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: Brownian Uncertainty
2. Stochastic Calculus: Valuing Continuation
3. Optimal Stopping: Boundaries and Verification
4. Investment: The Value of Waiting
5. Entry and Exit: Hysteresis
6. Strategic Stopping: Races and Contests

Why Ordinary Calculus Is Not Enough

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

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

$$B_t^2 = 2 \int_0^t B_s dB_s.$$

Why Ordinary Calculus Is Not Enough

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

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

$$B_t^2 = 2 \int_0^t B_s dB_s.$$

But a nonanticipating sum $\sum B_{t_k}(B_{t_{k+1}} - B_{t_k})$ has mean zero.

$\mathbb{E}[B_t^2] = t$. Something is missing.

Why Ordinary Calculus Is Not Enough

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

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

$$B_t^2 = 2 \int_0^t B_s dB_s.$$

But a nonanticipating sum $\sum B_{t_k}(B_{t_{k+1}} - B_{t_k})$ has mean zero.

$\mathbb{E}[B_t^2] = t$. Something is missing.

Exact identity:

$$B_t^2 = 2 \sum_k B_{t_k} \Delta B_k + \sum_k (\Delta B_k)^2.$$

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

Proposition (Brownian Quadratic Variation)

For deterministic partitions π of $[0, t]$ with mesh $|\pi| \rightarrow 0$,

$$\sum_k (B_{t_{k+1}} - B_{t_k})^2 \longrightarrow t \quad \text{in } L^2.$$

Write $[B]_t = t$.

Proof

The sum has expectation t . Independence gives variance

$$2 \sum_k (t_{k+1} - t_k)^2 \leq 2t|\pi| \longrightarrow 0.$$

Continuous finite-variation paths have zero quadratic variation.

Brownian motion has infinite total variation on every nontrivial interval a.s.

Itô Integration

Start with $H_s = \sum_k H_k \mathbf{1}_{\{t_k < s \leq t_{k+1}\}}$, where H_k is $\mathcal{F}_{t_k}$ -measurable.

Define

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

The integrand is chosen **before** the next increment is observed.

Itô Integration

Start with $H_s = \sum_k H_k \mathbf{1}_{\{t_k < s \leq t_{k+1}\}}$, where H_k is $\mathcal{F}_{t_k}$ -measurable.

Define

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

The integrand is chosen **before** the next increment is observed.

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

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

In particular, continuous adapted integrands are predictable.

Proposition (Itô Isometry and Martingale Property)

If H is predictable and $\mathbb{E} \int_0^T H_s^2 ds < \infty$, then $M_t = \int_0^t H_s dB_s$, $t \leq T$, is a square-integrable martingale and

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

Also $[M]_t = \int_0^t H_s^2 ds$.

Properties of the Itô Integral

Proposition (Itô Isometry and Martingale Property)

If H is predictable and $\mathbb{E} \int_0^T H_s^2 ds < \infty$, then $M_t = \int_0^t H_s dB_s$, $t \leq T$, is a square-integrable martingale and

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

Also $[M]_t = \int_0^t H_s^2 ds$.

For simple H : cross terms vanish by conditioning; each squared increment has conditional expectation Δt .

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

Economic use: isolate an unpredictable term with zero conditional mean.

Local Martingales

Definition

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

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

It need not have mean zero without an integrability condition.

Local Martingales

Definition

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

If $\int_0^t H_s^2 ds < \infty$ a.s., the 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.

When taking expectations after Itô: check which case applies.

Itô Processes and Diffusions

An **Itô process** has the form

$$X_t = X_0 + \int_0^t b_s \, ds + \int_0^t a_s \, dB_s, \quad dX_t = b_t \, dt + a_t \, dB_t,$$

with suitable measurability and $\int_0^t (|b_s| + a_s^2) \, ds < \infty$ a.s.

dX_t is shorthand for the integral equation.

Itô Processes and Diffusions

An **Itô process** has the form

$$X_t = X_0 + \int_0^t b_s \, ds + \int_0^t a_s \, dB_s, \quad dX_t = b_t \, dt + a_t \, dB_t,$$

with suitable measurability and $\int_0^t (|b_s| + a_s^2) \, ds < \infty$ a.s.

dX_t is shorthand for the integral equation.

A time-homogeneous **diffusion** solves

$$dX_t = \mu(X_t) \, dt + \sigma(X_t) \, dB_t.$$

Global Lipschitz coefficients with linear growth give a unique, nonexplosive strong solution.

This solution is strong Markov; local versions require control of explosion and boundaries.

Theorem (Itô's Formula)

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

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

Derivatives are evaluated at (t, X_t) .

The new term is $\frac{1}{2} a_t^2 f_{xx} dt$.

Theorem (Itô's Formula)

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

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

Derivatives are evaluated at (t, X_t) .

The new term is $\frac{1}{2} a_t^2 f_{xx} dt$.

Taylor expansion: $\Delta X = O(\sqrt{\Delta t})$, so $(\Delta X)^2 = O(\Delta t)$.

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

These are bookkeeping rules for limits of sums, not ordinary algebra.

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.

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

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

The drift cancels; Gaussian moments verify integrability on finite horizons.

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

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

The drift cancels; Gaussian moments verify integrability on finite horizons.

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

Infinitesimal Generator

For $dX_t = \mu(X_t) dt + \sigma(X_t) dB_t$, define

Definition (Infinitesimal Generator)

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

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

Drift of $f(X_t)$, not just drift of X_t .

Infinitesimal Generator

For $dX_t = \mu(X_t) dt + \sigma(X_t) dB_t$, define

Definition (Infinitesimal Generator)

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

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

Drift of $f(X_t)$, not just drift of X_t .

Brownian state: $\mathcal{L}f = \frac{1}{2}\sigma^2f''$.

A martingale state can have a rising expected value after a convex transformation.

Discounting replaces $\mathcal{L}$ by $\mathcal{L} - r$.

Theorem (Discounted Dynkin Formula)

For $f \in C^2$ and a stopping time T , under conditions allowing optional sampling and passage to the limit,

$$\mathbb{E}_x[e^{-rT}f(X_T)] = f(x) + \mathbb{E}_x \int_0^T e^{-rs}(\mathcal{L}f - rf)(X_s) ds.$$

Proof

Apply Itô to $e^{-rt}f(X_t)$ and take expectations. The stochastic integral disappears if it is a true martingale stopped at T . For example, it suffices that $\mathbb{E} \int_0^T e^{-2rs} \sigma^2(X_s) f'(X_s)^2 ds < \infty$, with the drift and terminal terms integrable.

Exit Problems as Differential Equations

Let τ be first exit from a bounded interval (a, b) .

Exit probability: if $p(x) = \mathbb{P}_x(X_\tau = b)$, then

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

Exit Problems as Differential Equations

Let τ be first exit from a bounded interval (a, b) .

Exit probability: if $p(x) = \mathbb{P}_x(X_\tau = b)$, then

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

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

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

Exit Problems as Differential Equations

Let τ be first exit from a bounded interval (a, b) .

Exit probability: if $p(x) = \mathbb{P}_x(X_\tau = b)$, then

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

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

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

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

$$p(x) = \frac{x - a}{b - a}, \quad m(x) = \frac{(x - a)(b - x)}{\sigma^2}.$$

Dynkin verifies the solutions; regularity is needed for the converse.

Overview

1. Buying a Product: Brownian Uncertainty
2. Stochastic Calculus: Valuing Continuation
3. Optimal Stopping: Boundaries and Verification
4. Investment: The Value of Waiting
5. Entry and Exit: Hysteresis
6. Strategic Stopping: Races and Contests

The General Stopping Problem

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

Flow payoff $f(x)$ while continuing; terminal payoff $g(x)$.

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

Admissibility includes the integrability needed to define the objective.

The General Stopping Problem

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

Flow payoff $f(x)$ while continuing; terminal payoff $g(x)$.

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

Admissibility includes the integrability needed to define the objective.

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

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

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

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

Dynamic Programming Principle

For an intermediate bounded stopping time θ , the DPP is

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

Stop before θ , or reach θ and solve the remaining problem.

Dynamic Programming Principle

For an intermediate bounded stopping time θ , the DPP is

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

Stop before θ , or reach θ and solve the remaining problem.

If θ occurs before leaving the continuation region,

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

Requires the Markov property and feasible continuation rules; not just an algebraic identity.

The Stopping Equation

Suppose V is smooth locally in the continuation region.

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

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

The Stopping Equation

Suppose V is smooth locally in the continuation region.

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

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

Divide by the expected interval length and take the local limit:

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

In the stopping region, immediate stopping must beat waiting.

Remark (Variational Inequality / Stopping HJB)

$$\max\{g - V, \mathcal{L}V - rV + f\} = 0.$$

Value Matching and Smooth Fit

The continuation region is unknown: a **free-boundary problem**.

Value matching: at a boundary b where V is continuous,

$$V(b) = g(b).$$

Value Matching and Smooth Fit

The continuation region is unknown: a **free-boundary problem**.

Value matching: at a boundary b where V is continuous,

$$V(b) = g(b).$$

Smooth fit: for a regular boundary of a nondegenerate one-dimensional diffusion, with differentiable payoff and suitable regularity,

$$V'(b) = g'(b).$$

Intuition: arbitrarily small moves across the boundary should not create a first-order gain from waiting.

Value Matching and Smooth Fit

The continuation region is unknown: a **free-boundary problem**.

Value matching: at a boundary b where V is continuous,

$$V(b) = g(b).$$

Smooth fit: for a regular boundary of a nondegenerate one-dimensional diffusion, with differentiable payoff and suitable regularity,

$$V'(b) = g'(b).$$

Intuition: arbitrarily small moves across the boundary should not create a first-order gain from waiting.

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

Here we use it to find a candidate, then verify that candidate directly.

Back to the Consumer

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

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

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

Back to the Consumer

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

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

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

Four unknowns; four boundary conditions:

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

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

Value matching then gives $a = -b$.

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

Proposition

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

$$V(x) = \begin{cases} 0, & x \leq -b, \\ \frac{c}{\sigma^2}(x+b)^2, & -b < x < b, \\ x, & x \geq b, \end{cases} \quad \tau^* = \inf\{t \geq 0 : |X_t| \geq b\}$$

solves the consumer's stopping problem.

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

Continue even when expected net utility is positive, if $X_t < b$.

Proposition

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

$$V(x) = \begin{cases} 0, & x \leq -b, \\ \frac{c}{\sigma^2}(x+b)^2, & -b < x < b, \\ x, & x \geq b, \end{cases} \quad \tau^* = \inf\{t \geq 0 : |X_t| \geq b\}$$

solves the consumer's stopping problem.


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

Continue even when expected net utility is positive, if $X_t < b$.

Symmetric boundaries come from the driftless, undiscounted specification.

They are not a general property of learning and stopping.

The Value of Waiting



$V - g > 0$ only in the continuation region; slopes agree at the boundaries.

Verification: No Rule Can Do Better

Let v be the candidate. Check

$$v \geq g, \quad \frac{1}{2}\sigma^2 v'' - c \leq 0,$$

with equality in $(-b, b)$.

v is C^1 and piecewise C^2 with bounded first and second derivatives.

The corresponding extension of Itô's formula applies; there is no local-time term because v' is continuous.

Verification: No Rule Can Do Better

Let v be the candidate. Check

$$v \geq g, \quad \frac{1}{2}\sigma^2 v'' - c \leq 0,$$

with equality in $(-b, b)$.

v is C^1 and piecewise C^2 with bounded first and second derivatives.

The corresponding extension of Itô's formula applies; there is no local-time term because v' is continuous.

For any admissible τ , $\mathbb{E}\tau < \infty$ and $|v'| \leq 1$ give

$$\mathbb{E}_x[v(X_\tau) - c\tau] = v(x) + \mathbb{E}_x \int_0^\tau (\frac{1}{2}\sigma^2 v''(X_s) - c) ds \leq v(x).$$

Since $v \geq g$, every rule gives at most $v(x)$.

Verification: The Candidate Attains the Bound

Under τ^* the drift of $v(X_t) - ct$ is zero until stopping.

$$\mathbb{E}_x[\tau^*] = (b^2 - x^2)/\sigma^2 < \infty \text{ for } |x| < b.$$

At stopping, $v(X_{\tau^*}) = g(X_{\tau^*})$.

Verification: The Candidate Attains the Bound

Under τ^* the drift of $v(X_t) - ct$ is zero until stopping.

$$\mathbb{E}_x[\tau^*] = (b^2 - x^2)/\sigma^2 < \infty \text{ for } |x| < b.$$

At stopping, $v(X_{\tau^*}) = g(X_{\tau^*})$.

Therefore

$$\mathbb{E}_x[g(X_{\tau^*}) - c\tau^*] = v(x).$$

Upper bound for all rules + equality for one rule $\implies$ optimality.

Verification: The Candidate Attains the Bound

Under τ^* the drift of $v(X_t) - ct$ is zero until stopping.

$$\mathbb{E}_x[\tau^*] = (b^2 - x^2)/\sigma^2 < \infty \text{ for } |x| < b.$$

At stopping, $v(X_{\tau^*}) = g(X_{\tau^*})$.

Therefore

$$\mathbb{E}_x[g(X_{\tau^*}) - c\tau^*] = v(x).$$

Upper bound for all rules + equality for one rule $\implies$ optimality.

What was proved?

The value is finite and equals v .

An optimal stopping time exists.

No history-dependent rule improves on the threshold rule.

Purchase Probabilities and Search Duration

For $|x| < b$, where $b = \sigma^2/(4c)$,

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

Outside $(-b, b)$, stop immediately.

Purchase Probabilities and Search Duration

For $|x| < b$, where $b = \sigma^2/(4c)$,

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

Outside $(-b, b)$, stop immediately.

Higher c : narrower continuation region and less expected search.

Higher σ : faster movement, but wider optimal boundaries.

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

Here the boundary response dominates: more informative discoveries can lead to *longer* search.

Distinguish a change in the process from a change in the optimal rule.

Pricing When Consumers Can Search

Baseline expected valuation u , price p , marginal production cost k .

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

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

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

Pricing When Consumers Can Search

Baseline expected valuation u , price p , marginal production cost k .

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

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

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

If $u + b \leq k$: no strictly profitable sale.

Otherwise,

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

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

Lower Search Costs Can Hurt the Consumer

Hold the price fixed: lower c expands feasible gains from search.

But the seller adjusts its price. Write $d = u - k > 0$.

In the interior pricing region, $b > d/3$,

$$p^* = k + \frac{d+b}{2}, \quad x^* = \frac{d-b}{2}, \quad V(x^*) = \frac{(d+b)^2}{16b}.$$

Lower Search Costs Can Hurt the Consumer

Hold the price fixed: lower c expands feasible gains from search.

But the seller adjusts its price. Write $d = u - k > 0$.

In the interior pricing region, $b > d/3$,

$$p^* = k + \frac{d+b}{2}, \quad x^* = \frac{d-b}{2}, \quad V(x^*) = \frac{(d+b)^2}{16b}.$$

Since $\partial V(x^*)/\partial b = (1 - d^2/b^2)/16$,

$d/3 < b < d \implies \text{lower } c \implies \text{lower consumer surplus.}$

The price response can outweigh the direct benefit of cheaper information.

Individual and equilibrium comparative statics need not agree.

Overview

1. Buying a Product: Brownian Uncertainty
2. Stochastic Calculus: Valuing Continuation
3. Optimal Stopping: Boundaries and Verification
4. Investment: The Value of Waiting
5. Entry and Exit: Hysteresis
6. Strategic Stopping: Races and Contests

An Irreversible Investment

Setup (McDonald & Siegel 1986 QJE; Dixit & Pindyck 1994, ch. 5)

Pay a sunk cost $I > 0$ to acquire a project worth X_t .

The opportunity remains available if the firm waits.

Discount rate $r > 0$; project value follows geometric Brownian motion.

$$dX_t = \mu X_t dt + \sigma X_t dB_t, \quad X_0 = x > 0, \quad \sigma > 0, \quad r > \mu.$$

An Irreversible Investment

Setup (McDonald & Siegel 1986 QJE; Dixit & Pindyck 1994, ch. 5)

Pay a sunk cost $I > 0$ to acquire a project worth X_t .

The opportunity remains available if the firm waits.

Discount rate $r > 0$; project value follows geometric Brownian motion.

$$dX_t = \mu X_t dt + \sigma X_t dB_t, \quad X_0 = x > 0, \quad \sigma > 0, \quad r > \mu.$$

Constant-cost benchmark of the more general investment problem:

$$F(x) = \sup_{\tau} \mathbb{E}_x[e^{-r\tau}(X_{\tau} - I)].$$

Never investing is allowed and pays 0.

Positive NPV, $x > I$, is sufficient if the only alternative is *never* investing. What if the firm can wait?

Geometric Brownian Motion

Apply Itô to $\log X_t$:

$$d \log X_t = (\mu - \frac{1}{2}\sigma^2) dt + \sigma dB_t.$$

Therefore

$$X_t = x \exp \left((\mu - \frac{1}{2}\sigma^2)t + \sigma B_t \right).$$

Positive paths; proportional rather than additive shocks.

Geometric Brownian Motion

Apply Itô to $\log X_t$:

$$d \log X_t = (\mu - \frac{1}{2}\sigma^2) dt + \sigma dB_t.$$

Therefore

$$X_t = x \exp \left((\mu - \frac{1}{2}\sigma^2)t + \sigma B_t \right).$$

Positive paths; proportional rather than additive shocks.

$\mathbb{E}_x[X_t] = xe^{\mu t}$, while expected log growth is $\mu - \sigma^2/2$.

Generator:

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

A power x^q stays a power after applying $\mathcal{L}$.

Solving the Investment Problem

Guess: invest once $X_t \geq x^*$.

In the waiting region,

$$\frac{1}{2}\sigma^2 x^2 F'' + \mu x F' - rF = 0.$$

Try $F(x) = Ax^\beta$. The characteristic equation is

$$Q(\beta) := \frac{1}{2}\sigma^2 \beta(\beta - 1) + \mu\beta - r = 0.$$

Solving the Investment Problem

Guess: invest once $X_t \geq x^*$.

In the waiting region,

$$\frac{1}{2}\sigma^2 x^2 F'' + \mu x F' - rF = 0.$$

Try $F(x) = Ax^\beta$. The characteristic equation is

$$Q(\beta) := \frac{1}{2}\sigma^2\beta(\beta-1) + \mu\beta - r = 0.$$

Roots: $\beta_- < 0 < 1 < \beta_+$, since $r > \mu$.

$F(0+) = 0$ rules out x^{β_-} ; write $\beta = \beta_+$.

Value matching and smooth fit:

$$A(x^*)^\beta = x^* - I, \quad \beta A(x^*)^{\beta-1} = 1.$$

The Value of Waiting to Invest

Proposition

Let $\beta > 1$ solve $Q(\beta) = 0$. Then

$$x^* = \frac{\beta}{\beta - 1}I > I, \quad F(x) = \begin{cases} (x^* - I)(x/x^*)^\beta, & x < x^*, \\ x - I, & x \geq x^*. \end{cases}$$

An optimal rule is $\tau^* = \inf\{t \geq 0 : X_t \geq x^*\}$.

The firm may optimally reject a project with positive NPV.

Investing destroys the option to wait and reconsider.

$x^* - I$ compensates for that lost option.

Investment and the Option Value




Illustration: $\beta = 2$. Waiting remains optimal for $I < x < 2I$.

Verifying the Investment Rule

Candidate F is nonnegative, C^1 , and piecewise C^2 .

$F \geq x - I$: the power curve is convex and tangent at x^* .

In the waiting region, $(\mathcal{L} - r)F = 0$.

Verifying the Investment Rule

Candidate F is nonnegative, C^1 , and piecewise C^2 .

$F \geq x - I$: the power curve is convex and tangent at x^* .

In the waiting region, $(\mathcal{L} - r)F = 0$.

In the investment region,

$$(\mathcal{L} - r)(x - I) = (\mu - r)x + rI \leq 0.$$

At $x = x^*$ this is $-\frac{1}{2}\sigma^2\beta I < 0$; it decreases in x .

Verifying the Investment Rule

Candidate F is nonnegative, C^1 , and piecewise C^2 .

$F \geq x - I$: the power curve is convex and tangent at x^* .

In the waiting region, $(\mathcal{L} - r)F = 0$.

In the investment region,

$$(\mathcal{L} - r)(x - I) = (\mu - r)x + rI \leq 0.$$

At $x = x^*$ this is $-\frac{1}{2}\sigma^2\beta I < 0$; it decreases in x .

Thus $e^{-rt}F(X_t)$ is a nonnegative local supermartingale, hence a supermartingale: it bounds every admissible payoff.

Attainment follows from the discounted hitting-time calculation next.

Discounted Hitting Times

Let $\tau_b = \inf\{t : X_t \geq b\}$, with $x < b$.

Since $(\mathcal{L} - r)x^\beta = 0$, $e^{-rt}X_t^\beta$ is a local martingale.

Stopped at τ_b , it is bounded by b^β .

Discounted Hitting Times

Let $\tau_b = \inf\{t : X_t \geq b\}$, with $x < b$.

Since $(\mathcal{L} - r)x^\beta = 0$, $e^{-rt}X_t^\beta$ is a local martingale.

Stopped at τ_b , it is bounded by b^β .

Optional sampling at $t \wedge \tau_b$ and bounded convergence give

$$x^\beta = b^\beta \mathbb{E}_x[e^{-r\tau_b} \mathbf{1}_{\{\tau_b < \infty\}}].$$

On $\{\tau_b = \infty\}$, $e^{-rt}X_t^\beta \leq e^{-rt}b^\beta \rightarrow 0$.

Discounted Hitting Times

Let $\tau_b = \inf\{t : X_t \geq b\}$, with $x < b$.

Since $(\mathcal{L} - r)x^\beta = 0$, $e^{-rt}X_t^\beta$ is a local martingale.

Stopped at τ_b , it is bounded by b^β .

Optional sampling at $t \wedge \tau_b$ and bounded convergence give

$$x^\beta = b^\beta \mathbb{E}_x[e^{-r\tau_b} \mathbf{1}_{\{\tau_b < \infty\}}].$$

$$\text{On } \{\tau_b = \infty\}, e^{-rt}X_t^\beta \leq e^{-rt}b^\beta \rightarrow 0.$$

Therefore

$$\mathbb{E}_x[e^{-r\tau_b} \mathbf{1}_{\{\tau_b < \infty\}}] = (x/b)^\beta.$$

Set $b = x^*$ to obtain exactly $F(x)$.

Investment under Greater Uncertainty

$$x^* = I\beta/(\beta - 1); \partial x^*/\partial \beta = -I/(\beta - 1)^2 < 0.$$

At the positive root, $Q_\beta > 0$.

Implicit differentiation gives

$$\frac{\partial \beta}{\partial \sigma} = -\frac{\sigma \beta (\beta - 1)}{Q_\beta} < 0.$$

Investment under Greater Uncertainty

$$x^* = I\beta/(\beta - 1); \partial x^*/\partial \beta = -I/(\beta - 1)^2 < 0.$$

At the positive root, $Q_\beta > 0$.

Implicit differentiation gives

$$\frac{\partial \beta}{\partial \sigma} = -\frac{\sigma \beta (\beta - 1)}{Q_\beta} < 0.$$

Higher volatility raises the investment threshold.

Higher I raises x^* proportionally; higher r lowers x^* when x is the given project value.

Investment under Greater Uncertainty

$$x^* = I\beta/(\beta - 1); \partial x^*/\partial \beta = -I/(\beta - 1)^2 < 0.$$

At the positive root, $Q_\beta > 0$.

Implicit differentiation gives

$$\frac{\partial \beta}{\partial \sigma} = -\frac{\sigma \beta (\beta - 1)}{Q_\beta} < 0.$$

Higher volatility raises the investment threshold.

Higher I raises x^* proportionally; higher r lowers x^* when x is the given project value.

Do not confuse project value with a profit flow.

If the state is a flow π_t and project value is $\pi_t/(r - \mu)$, then the flow trigger is $(r - \mu)x^*$; changing r also changes capitalisation.

Overview

1. Buying a Product: Brownian Uncertainty
2. Stochastic Calculus: Valuing Continuation
3. Optimal Stopping: Boundaries and Verification
4. Investment: The Value of Waiting
5. Entry and Exit: Hysteresis
6. Strategic Stopping: Races and Contests

Entry and Exit Decisions

Setup (Dixit 1989 JPE)

Output price: $dP_t = \mu P_t dt + \sigma P_t dB_t$, $r > \max\{\mu, 0\}$.

Active firm earns flow profit $P_t - w$, $w > 0$.

Idle firm earns 0.

Entry costs $I > 0$; exit costs $E \geq 0$; re-entry is possible.

Entry and Exit Decisions

Setup (Dixit 1989 JPE)

Output price: $dP_t = \mu P_t dt + \sigma P_t dB_t$, $r > \max\{\mu, 0\}$.

Active firm earns flow profit $P_t - w$, $w > 0$.

Idle firm earns 0.

Entry costs $I > 0$; exit costs $E \geq 0$; re-entry is possible.

$V_0(p)$: value when idle; $V_1(p)$: value when active.

Switching changes the regime, not the price process.

Entry and exit are linked stopping problems.

Two Linked Stopping Equations

Idle: wait or pay I and become active.

$$\max\{V_1 - I - V_0, (\mathcal{L} - r)V_0\} = 0.$$

Active: operate or pay E and become idle.

$$\max\{V_0 - E - V_1, (\mathcal{L} - r)V_1 + p - w\} = 0.$$

Two Linked Stopping Equations

Idle: wait or pay I and become active.

$$\max\{V_1 - I - V_0, (\mathcal{L} - r)V_0\} = 0.$$

Active: operate or pay E and become idle.

$$\max\{V_0 - E - V_1, (\mathcal{L} - r)V_1 + p - w\} = 0.$$

Guess: enter at p_H , exit at $p_L < p_H$.

For $p_L < p < p_H$, action depends on whether the firm is already active.

The state is (p, regime) , not just p .

Solving for the Entry and Exit Triggers

Let $\beta_- < 0 < 1 < \beta_+$ solve the same characteristic equation as before.

In each regime's continuation region,

$$V_0(p) = Ap^{\beta_+}, \quad V_1(p) = \frac{p}{r - \mu} - \frac{w}{r} + Bp^{\beta_-}.$$

The last term is the option to exit; it vanishes as $p \rightarrow \infty$.

Solving for the Entry and Exit Triggers

Let $\beta_- < 0 < 1 < \beta_+$ solve the same characteristic equation as before.

In each regime's continuation region,

$$V_0(p) = Ap^{\beta_+}, \quad V_1(p) = \frac{p}{r - \mu} - \frac{w}{r} + Bp^{\beta_-}.$$

The last term is the option to exit; it vanishes as $p \rightarrow \infty$.

Value matching and smooth fit give four equations:

$$\begin{aligned} V_0(p_H) &= V_1(p_H) - I, & V'_0(p_H) &= V'_1(p_H), \\ V_1(p_L) &= V_0(p_L) - E, & V'_1(p_L) &= V'_0(p_L). \end{aligned}$$

Solve for A, B, p_L, p_H ; check the regime inequalities.

If exit is too costly, an interior exit trigger need not exist.

Hysteresis

Consider $w > rE$, so exit can be worthwhile at low prices.

At the entry boundary, waiting must be unattractive to the idle firm:

$$(\mathcal{L} - r)(V_1 - l) = -(p - w) + rl \leq 0 \quad \implies \quad p_H \geq w + rl.$$

Hysteresis

Consider $w > rE$, so exit can be worthwhile at low prices.

At the entry boundary, waiting must be unattractive to the idle firm:

$$(\mathcal{L} - r)(V_1 - I) = -(p - w) + rI \leq 0 \implies p_H \geq w + rI.$$

At the exit boundary, continued operation must be unattractive:

$$(\mathcal{L} - r)(V_0 - E) + p - w = rE + p - w \leq 0 \implies p_L \leq w - rE.$$

With nondegenerate uncertainty and interior triggers, the standard solution has strict inequalities.

Hysteresis

Consider $w > rE$, so exit can be worthwhile at low prices.

At the entry boundary, waiting must be unattractive to the idle firm:

$$(\mathcal{L} - r)(V_1 - I) = -(p - w) + rI \leq 0 \implies p_H \geq w + rI.$$

At the exit boundary, continued operation must be unattractive:

$$(\mathcal{L} - r)(V_0 - E) + p - w = rE + p - w \leq 0 \implies p_L \leq w - rE.$$

With nondegenerate uncertainty and interior triggers, the standard solution has strict inequalities.

Hysteresis: the same current price can sustain an active or an idle firm, depending on past entry and exit.

Entry and Exit: A Numerical Example

$$r = .05, \mu = 0, \sigma = .20, w = 1, l = 2, E = 1.$$

Solving the four boundary equations gives

$$p_L \simeq 0.651, \quad p_H \simeq 1.675.$$

Compare with $w - rE = .95$ and $w + rl = 1.10$.

	$p < p_L$	$p_L < p < p_H$	$p > p_H$
Currently idle	Wait	Wait	Enter
Currently active	Exit	Operate	Operate

Sunk costs and uncertainty create a region of inaction.

A temporary price shock can have a persistent effect on activity.

Overview

1. Buying a Product: Brownian Uncertainty
2. Stochastic Calculus: Valuing Continuation
3. Optimal Stopping: Boundaries and Verification
4. Investment: The Value of Waiting
5. Entry and Exit: Hysteresis
6. Strategic Stopping: Races and Contests

Strategic Delay in R&D

Weeds (2002 REStud)

Two firms choose when to initiate irreversible research.

Patent value follows geometric Brownian motion.

Active research succeeds at a Poisson arrival rate.

The first discovery wins the patent.

Strategic Delay in R&D

Weeds (2002 REStud)

Two firms choose when to initiate irreversible research.

Patent value follows geometric Brownian motion.

Active research succeeds at a Poisson arrival rate.

The first discovery wins the patent.

Waiting preserves flexibility; investing can secure a lead.

A rival's entry reduces the value of an active project.

Competition can induce preemption, but also sustain joint delay.

The monopoly stopping threshold alone does not determine equilibrium.

Leader and Follower Values

Compute payoffs backwards:

- Value when both firms conduct research.

- Follower's option to initiate research.

- Leader's value, anticipating follower entry.

Leader and Follower Values

Compute payoffs backwards:

- Value when both firms conduct research.

- Follower's option to initiate research.

- Leader's value, anticipating follower entry.

Compare net leader and follower values, $L(x)$ and $F(x)$.

Where $L(x) > F(x)$, each firm would prefer to move first.

Whether delaying together is sustainable also requires checking deviations.

Leader and Follower Values

Compute payoffs backwards:

- Value when both firms conduct research.

- Follower's option to initiate research.

- Leader's value, anticipating follower entry.

Compare net leader and follower values, $L(x)$ and $F(x)$.

Where $L(x) > F(x)$, each firm would prefer to move first.

Whether delaying together is sustainable also requires checking deviations.

Smooth fit applies at an optimally chosen follower boundary.

At the rival's boundary, the leader's value must match; the leader does not choose that boundary, so do not impose its smooth-fit condition there.

Research Success as a Random Arrival

Symmetric race: entry cost I , research flow cost c_R , success rate λ per active firm, independent of patent value X .

In a short interval h , one active firm's success probability is $\lambda h + o(h)$.

With two active firms, the race ends at total rate 2λ .

Research Success as a Random Arrival

Symmetric race: entry cost I , research flow cost c_R , success rate λ per active firm, independent of patent value X .

In a short interval h , one active firm's success probability is $\lambda h + o(h)$.

With two active firms, the race ends at total rate 2λ .

A firm's value after both have entered is

$$\begin{aligned} A_2(X) &= \mathbb{E}_X \int_0^\infty e^{-(r+2\lambda)t} (\lambda X_t - c_R) dt \\ &= \frac{\lambda X}{r + 2\lambda - \mu} - \frac{c_R}{r + 2\lambda}. \end{aligned}$$

Assume $r > \max\{\mu, 0\}$. The rival's possible success discounts future gains.

The Follower's Stopping Problem

While the follower waits, the leader can succeed and end the opportunity at rate λ .

Therefore, before the follower's entry threshold x_F ,

$$(\mathcal{L} - r - \lambda)F = 0.$$

Write $F(x) = Kx^\gamma$, where $\gamma > 1$ solves

$$\frac{1}{2}\sigma^2\gamma(\gamma - 1) + \mu\gamma - (r + \lambda) = 0.$$

The Follower's Stopping Problem

While the follower waits, the leader can succeed and end the opportunity at rate λ .

Therefore, before the follower's entry threshold x_F ,

$$(\mathcal{L} - r - \lambda)F = 0.$$

Write $F(x) = Kx^\gamma$, where $\gamma > 1$ solves

$$\frac{1}{2}\sigma^2\gamma(\gamma - 1) + \mu\gamma - (r + \lambda) = 0.$$

Set $F(x_F) = A_2(x_F) - I$ and $F'(x_F) = A'_2(x_F)$:

$$x_F = \frac{\gamma}{\gamma - 1} \left(I + \frac{C_R}{r + 2\lambda} \right) \frac{r + 2\lambda - \mu}{\lambda}.$$

The stopping method survives; the rival changes both the payoff and the effective discount rate.

The Leader Anticipates Entry

Let $A_1(x)$ be the active leader's value before follower entry.

Until X reaches x_F ,

$$(\mathcal{L} - r - \lambda)A_1 + \lambda x - c_R = 0.$$

A solution satisfying the natural condition near zero is

$$A_1(x) = \frac{\lambda x}{r + \lambda - \mu} - \frac{c_R}{r + \lambda} + Cx^\gamma.$$

The Leader Anticipates Entry

Let $A_1(x)$ be the active leader's value before follower entry.

Until X reaches x_F ,

$$(\mathcal{L} - r - \lambda)A_1 + \lambda x - c_R = 0.$$

A solution satisfying the natural condition near zero is

$$A_1(x) = \frac{\lambda x}{r + \lambda - \mu} - \frac{c_R}{r + \lambda} + Cx^\gamma.$$

Determine C from $A_1(x_F) = A_2(x_F)$.

Net leader value before initiating research is $L(x) = A_1(x) - I$.

These values supply the timing game; choosing equilibrium initiation requires the leader–follower incentive comparisons.

Gambling in Contests

Setup (Seel & Strack 2013 JET)

$n \geq 2$ players privately observe independent processes $X_t^i = x_0 + \mu t + \sigma B_t^i$, $x_0 > 0$.

Choose a possibly randomised stopping time τ_i , with $\mathbb{E}\tau_i < \infty$.

Must stop by bankruptcy: $\tau_i \leq \inf\{t : X_t^i = 0\}$.

Highest stopped value wins a unit prize; ties split it.

Gambling in Contests

Setup (Seel & Strack 2013 JET)

$n \geq 2$ players privately observe independent processes $X_t^i = x_0 + \mu t + \sigma B_t^i$, $x_0 > 0$.

Choose a possibly randomised stopping time τ_i , with $\mathbb{E}\tau_i < \infty$.

Must stop by bankruptcy: $\tau_i \leq \inf\{t : X_t^i = 0\}$.

Highest stopped value wins a unit prize; ties split it.

Focus on $\mu \leq 0$.

Given atomless rival cdf G , stopping at x pays $G(x)^{n-1}$.

Rank incentives change the stopping payoff.

A Driftless Contest

Let $\mu = 0$. Suppose symmetric equilibrium has upper support $\bar{x}$.

Indifference on the support requires a harmonic winning probability:

$$\frac{1}{2}\sigma^2 u'' = 0, \quad u(0) = 0, \quad u(\bar{x}) = 1.$$

Thus $u(x) = x/\bar{x}$ on $[0, \bar{x}]$.

A Driftless Contest

Let $\mu = 0$. Suppose symmetric equilibrium has upper support $\bar{x}$.

Indifference on the support requires a harmonic winning probability:

$$\frac{1}{2}\sigma^2 u'' = 0, \quad u(0) = 0, \quad u(\bar{x}) = 1.$$

Thus $u(x) = x/\bar{x}$ on $[0, \bar{x}]$.

Symmetry: $u(x_0) = 1/n$, hence $\bar{x} = nx_0$.

$$G(x) = \left(\frac{x}{nx_0} \right)^{1/(n-1)}, \quad 0 \leq x \leq nx_0.$$

A candidate distribution is not yet a stopping strategy.

Constructing an Equilibrium: Two Players

Proposition ($n = 2, \mu = 0$)

Independently draw $A_i \sim U(0, x_0)$ at time zero and stop on first exit from $(x_0 - A_i, x_0 + A_i)$. These strategies form an equilibrium.

Proof

Conditional on $A_i = a$, each endpoint is hit with probability $1/2$. Mixing over a makes $X_{\tau_i}^i$ uniform on $[0, 2x_0]$.

Against this distribution, the winning payoff is $u(x) = \min\{x/(2x_0), 1\}$ for $x \geq 0$. It is concave, so $u(X_{t \wedge \tau_0})$ is a bounded supermartingale.

Every deviation pays at most $u(x_0) = 1/2$; the proposed rule attains $1/2$.

Negative Drift and Risk Taking

For $\mu < 0$, define the scale function

$$s(x) = \frac{e^{-2\mu x/\sigma^2} - 1}{-2\mu/\sigma^2}, \quad \mathcal{L}s = 0.$$

Indifference gives $u(x) = s(x)/s(\bar{x})$ and $s(\bar{x}) = ns(x_0)$.

Hence the candidate cdf on $[0, \bar{x}]$ is

$$G(x) = \left(\frac{s(x)}{ns(x_0)} \right)^{1/(n-1)}.$$

Negative Drift and Risk Taking

For $\mu < 0$, define the scale function

$$s(x) = \frac{e^{-2\mu x/\sigma^2} - 1}{-2\mu/\sigma^2}, \quad \mathcal{L}s = 0.$$

Indifference gives $u(x) = s(x)/s(\bar{x})$ and $s(\bar{x}) = ns(x_0)$.

Hence the candidate cdf on $[0, \bar{x}]$ is

$$G(x) = \left(\frac{s(x)}{ns(x_0)} \right)^{1/(n-1)}.$$

For $n = 2$, draw $A \sim U(0, 1)$ and stop when $s(X_t)/s(x_0)$ exits $(1 - A, 1 + A)$.

Optional sampling gives equal endpoint probabilities; the construction induces G .

The winning payoff is convex on the interior support: rank incentives reward dispersion despite negative drift.

Private Incentives and Expected Losses

For $n = 2$ and $\mu < 0$, the constructed rule has bounded state support and finite expected duration.

Wald's identity gives

$$\mathbb{E}[X_\tau] = x_0 + \mu \mathbb{E}[\tau] < x_0.$$

Private Incentives and Expected Losses

For $n = 2$ and $\mu < 0$, the constructed rule has bounded state support and finite expected duration.

Wald's identity gives

$$\mathbb{E}[X_\tau] = x_0 + \mu \mathbb{E}[\tau] < x_0.$$

Maximising expected stopped value would favour stopping immediately.

Maximising the probability of winning induces continued gambling.

Zero expected drift removes expected losses, but not risk taking.

Private Incentives and Expected Losses

For $n = 2$ and $\mu < 0$, the constructed rule has bounded state support and finite expected duration.

Wald's identity gives

$$\mathbb{E}[X_\tau] = x_0 + \mu \mathbb{E}[\tau] < x_0.$$

Maximising expected stopped value would favour stopping immediately.

Maximising the probability of winning induces continued gambling.

Zero expected drift removes expected losses, but not risk taking.

The same tools now verify a best response to an endogenous payoff.

Stopping Problems: A Working Method

1. Specify the state, information, feasible stopping rules, and payoffs.
2. Write the generator and the stopping equation.
3. Propose a continuation region; solve inside it.
4. Use value matching and, when justified, smooth fit.
5. Verify the payoff bound, drift inequalities, and limiting conditions.
6. Recover behaviour: exit probabilities, duration, comparative statics.

Stopping Problems: A Working Method

1. Specify the state, information, feasible stopping rules, and payoffs.
2. Write the generator and the stopping equation.
3. Propose a continuation region; solve inside it.
4. Use value matching and, when justified, smooth fit.
5. Verify the payoff bound, drift inequalities, and limiting conditions.
6. Recover behaviour: exit probabilities, duration, comparative statics.

Next: when the drift is unknown, what state should the DM use?

Learning: derive the posterior process rather than assume its law.

References: Economic Applications

- Branco, Sun, & Villas-Boas (2012 MnSc), *Optimal Search for Product Information*.
- McDonald & Siegel (1986 QJE), *The Value of Waiting to Invest*.
- Dixit (1989 JPE), *Entry and Exit Decisions under Uncertainty*.
- Dixit & Pindyck (1994), *Investment under Uncertainty*, ch. 5.
- Weeds (2002 REStud), *Strategic Delay in a Real Options Model of R&D Competition*.
- Seel & Strack (2013 JET), *Gambling in Contests*.

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 $\mathbb{P}(X_t = Y_t \ \forall t \geq 0) = 1$.

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.

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.

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 $\{\tau < \infty\}$.

Information at a Stopping Time

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.

Thus the threshold rules used in the lecture are feasible.

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.

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

Reflection and One-Sided Hitting Times

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

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

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

Reflection and One-Sided Hitting Times

Let $\tau_a = \inf\{t : B_t = a\}$ 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.

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 $p(a) = 0, p(b) = 1$:

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

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 $(x - a)(b - x)/\sigma^2$.

The scale function converts the drifted exit problem into a linear one.

Theorem (A Sufficient Verification Result)

Let $v \geq g$ and let

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

be a supermartingale. Assume optional sampling and passage to the terminal payoff are valid for every admissible rule. Then $v \geq V$.

If a feasible τ^* satisfies $v(X_{\tau^*}) = g(X_{\tau^*})$ and $Z_{t \wedge \tau^*}$ is a uniformly integrable martingale, then $v = V$ and τ^* is optimal.

For $\tau^* = \infty$ on some paths, impose the corresponding terminal limit convention and transversality condition.

The drift inequality $(\mathcal{L} - r)v + f \leq 0$ is how Itô establishes the supermartingale property, subject to integrability.

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 .

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 .

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

For a Markov stopping problem,

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

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.

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

The stopping equation becomes

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

Finite Horizons

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

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

The stopping equation becomes

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

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 $g(x) = 1$ for $x < 0$, $g(x) = 2$ 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 g 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.