

Optimal 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.

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.

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.

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.

Overview

1. Buying a Product: Brownian Uncertainty

Overview

1. Buying a Product: Brownian Uncertainty

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.

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.

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 $\Delta \downarrow 0$.

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**.

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.

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

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.

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 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.

Why Learn 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!

Why Learn 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 .

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

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

Why Learn 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 .

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 .

Why Learn 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 .

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

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) .

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.

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

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

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

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.

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

Consumer's optimal stopping problem:

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

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

Consumer's optimal stopping problem:

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

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?

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.

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

Here $\mathbb{E}[T] = \infty$; 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.; continuous paths imply $X_{\tau_{a,b}} \in \{a, b\}$.

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

Exit probabilities: $p(x) := \mathbb{P}_x(X_{\tau_{a,b}} = b)$. Then,

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

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

Volatility changes how quickly the process gets there.