Optimal Stopping in Continuous Time: Part 1

Duarte Gonçalves

University College London

Topics in Economic Theory

Overview I

1. Continuous-Time Preliminaries via Optimal Pricing with Learning
   - Stochastic Processes
   - Martingale
   - Brownian Motion
   - Stochastic Calculus
   - Stopping Time
   - Properties of the Value Function
   - Optimal Stopping, Prob. Exit, and Exit Times
   - Ito's Lemma, DPP, and HJB

2. Simple Optimal Stopping Problems
   - Stopping
   - Optimal Stopping
   - Evolution of Beliefs
   - Value Function
   - Speed-Accuracy Revisited
   - A Recap of Wald's Equation/Identities

Motivation

Time as a *decision variable*

- When to sell an asset (house, car, stock)
- Deciding if/when to call a snap election
- If/when to adopt a new technology
- Response times: behavioral data that comes for free; what to make of it?

Learning takes time

- Learning about new product, new technology, market conditions, etc
- Strategic situations: preemption — e.g. improve product vs file patent first/post paper first, waiting for price to increase vs sell first, etc

Learning by doing

- Experimenting with products, interest rate adjustment, testing platform interface designs, etc

Overview I

1. Continuous-Time Preliminaries via Optimal Pricing with Learning
   - Stochastic Processes
   - Martingale
   - Brownian Motion
   - Stochastic Calculus
   - Stopping Time
   - Properties of the Value Function
   - Optimal Stopping, Prob. Exit, and Exit Times
   - Ito's Lemma, DPP, and HJB

2. Simple Optimal Stopping Problems

Application: Optimal Pricing

Setup (Branco, Sun, & Villas-Boas '12 MnSc)

(i) DM choosing whether or not to buy a product

(ii) Product has N attributes

Each attribute is good ($x_n = 1$) or bad ($x_n = -1$) with indep. equal prob

(iii) Value to DM: $u := u_0 + z \sum_{n=1}^{N} x_n$

u_0 is some baseline payoff (incl. price), $\pm z$ is payoff to good/bad attribute

(iv) It takes Δt for DM to learn about one attribute

Application: Optimal Pricing

Learning attributes

As attributes iid, order of search doesn't matter; fix order of search $1, 2, ..., N$

By time t, DM learned about $n_t := \lfloor t/\Delta t \rfloor$ attributes

Value of product at t

$$u_t := \mathbb{E}[u \mid \mathcal{F}_t] = \mathbb{E}\left[u_0 + z\sum_{n=1}^{N} x_n \middle| \mathcal{F}_t\right] = u_0 + z\sum_{n=1}^{n_t} x_n + z\sum_{n=n_t+1}^{N} \mathbb{E}[x_n \mid \mathcal{F}_t]$$

$$= u_0 + z\sum_{n=1}^{n_t} x_n$$

Value of product $\{u_t\}_t$ is stochastic process

Stochastic Processes

Stochastic Process: $X = \{X(t), t \in T\}$ S-valued r.v. defined on common prob. space $(\Omega, \mathcal{F}, P)$ mble wrt Σ; T: index set

- $X(\omega) \in S^T$: path for ω
- $\mathbb{F} = \{\mathcal{F}_t\}_{t \in T}$: filtration on prob. space $(\Omega, \mathcal{F}, P)$ family of σ-alg $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ $\forall 0 \leq s \leq t$
- Natural filtration of X $\{\mathcal{F}_t^X, t \in T\}$: smallest σ-alg wrt X_t is mble $\forall t \in T$
- X is adapted wrt $\mathbb{F}$ if X_t is $\mathcal{F}_t$-mble

Application: Optimal Pricing

Learning attributes

- As attributes iid, order of search doesn't matter; fix order of search $1, 2, ..., N$
- By time t, DM learned about $n_t := \lfloor t/\Delta t \rfloor$ attributes
  - Value of product at t

$$u_t := \mathbb{E}[u \mid \mathcal{F}_t] = \mathbb{E}\left[u_0 + z\sum_{n=1}^{N} x_n \middle| \mathcal{F}_t\right] = u_0 + z\sum_{n=1}^{n_t} x_n + z\sum_{n=n_t+1}^{N} \mathbb{E}[x_n \mid \mathcal{F}_t] = u_0 + z\sum_{n=1}^{n_t} x_n$$

- Value of product $\{u_t\}_t$ is stochastic process
- Value of product evolves as random walk with z increments: $u_{t+\Delta t} - u_t = +z$ or $-z$ with equal prob
- Value of product is martingale: $\mathbb{E}[u_s \mid \mathcal{F}_t] = u_t \; \forall s \geq t$

Martingale

Martingale: $\mathbb{E}[|X_t|] < \infty$ and $\mathbb{E}[X_t \mid \mathcal{F}_s] = X_s$ a.s. $\forall 0 \leq s \leq t$ (super $\leq$, sub $\geq$)

Application: Optimal Pricing

Simplification: continuum of attributes: taking $\Delta t \to 0$

- Many attributes $\to$ each contributes a small amount to value; $z = \sigma \Delta t^{1/2}$
- Note $|n_t - t/\Delta t| \leq \Delta t$; i.e. $n_t \approx t/\Delta t$ for small Δt
- Value at t is then

$$u_t = u_0 + u_t - u_0 = u_0 + z \sum_{n=1}^{n_t} x_n = u_0 + \sigma (\Delta t)^{1/2} \sum_{n=1}^{n_t} x_n$$

$$\approx u_0 + \sigma t^{1/2} n_t^{-1/2} \sum_{n=1}^{n_t} x_n \xrightarrow{d} N(u_0, \sigma^2 t)$$

- Increment in u from time s to t is

$$u_t - u_s = z \sum_{n=n_s}^{n_t} x_n = \sigma (\Delta t)^{1/2} \sum_{n=n_s}^{n_t} x_n \approx \sigma (t-s)^{1/2} (n_t - n_s)^{-1/2} \sum_{n=n_s}^{n_t} x_n \xrightarrow{d} N(0, (t-s)\sigma^2 t)$$

where $n_t - n_s \approx (t-s)/\Delta t$

u_t is (scaled) Brownian motion

Brownian Motion

Standard Brownian Motion

$B = \{B_t, t \in T\}$ $\mathbb{R}^d$-valued stochastic process s.t.

(i) $B_0 = 0$ (ii) $\forall 0 \leq s < t$ in T, $B_t - B_s \sim N(0, (t-s)I)$

(iii) $\forall 0 \leq s < t$ in T, $B_t - B_s$ is independent from B_u, $\forall u \leq s$

(iv) B is a.s. continuous in t

Equiv dfn: BM continuous martingale with $B_0 = 0$ and $[B]_t = t$

Quadratic Variation: $[X]_t = \lim_{\Delta \to 0} \sum_{n=1}^{N} (X_{t_n} - X_{t_{n-1}})^2$,

where $0 < t_0 < \cdots < t_N = t$ and $\Delta = \max_n t_n - t_{n-1}$

Other properties of BM

- Symmetry: $-B$ is also std BM
- Scaling: $\{B_{\lambda^2 t}/\lambda, t \in T\}$ std BM
- Invariance by translation: for all $s > 0$, $\{B_{t+s} - B_t, t \in T\}$ std BM indep from B_s
- B_t as limit process of symmetric random walk which changes by $\pm\sqrt{\Delta t}$ (w equal prob) each period Δt, when $\Delta t \to 0$

Standard Brownian Motion

Heuristically: $\mathrm{d}B_t = \sqrt{\mathrm{d}t}\,\varepsilon_t$, $\varepsilon_t \sim N(0,1)$

BM a.s. not differentiable. Intuition: $\lim_{\mathrm{d}t\to 0} \|\frac{B_{t+\mathrm{d}t}-B_t}{\mathrm{d}t}\| = |\varepsilon_t|(\mathrm{d}t)^{-1/2} = \infty$ a.s.

Making sense of $\int f\,\mathrm{d}B$...

Riemann-Stieltjes Integration:

π^n partition of $[a,b]$, such that $a = \pi_0^n < \pi_1^n < \cdots < \pi_n^n = b$

Π: set of partitions; mesh/norm $\|\pi_n\| := \sup_i |\pi_{i+1}^n - \pi_i^n|$

$$\int_a^b f\,\mathrm{d}G = \lim_{n\to\infty} \sum_{i=0}^{n} f(y_i)(G(\pi_{i+1}^n) - G(\pi_i^n))$$

where limit taken over seq of partitions π^n of $[a,b]$ and $\{y_i\}_{i=1}^n$ such that $y_i \in [\pi_i^n, \pi_{i+1}^n]$, and $\|\pi^n\| \to 0$

G with bounded variation and right semicont.

Stochastic Calculus

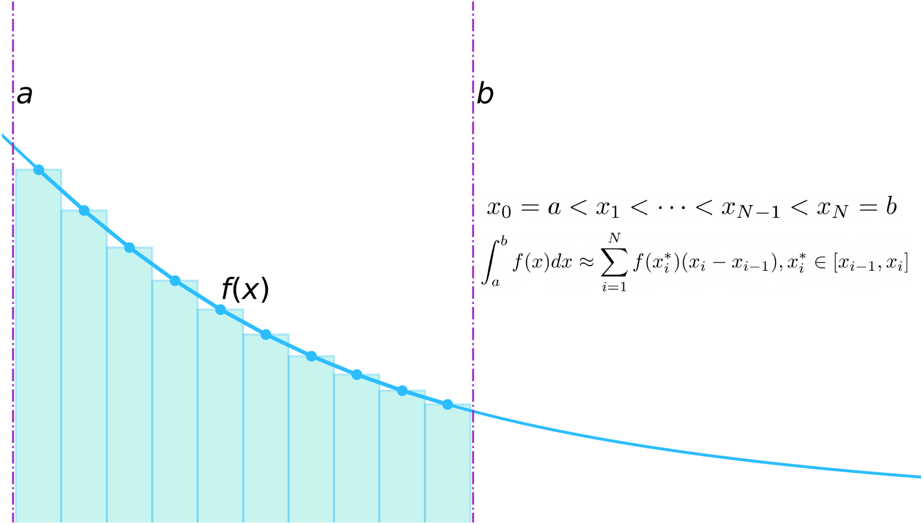

$$x_0 = a < x_1 < \cdots < x_{N-1} < x_N = b$$

$$\int_a^b f(x)dx \approx \sum_{i=1}^{N} f(x_i^*)(x_i - x_{i-1}), x_i^* \in [x_{i-1}, x_i]$$

Ito Integration:

$$\int_0^t f_s \, \mathrm{d}B_s = \lim_{n\to\infty} \sum_{i=0}^{n} f(y_i)(B(\pi_{i+1}^n) - B(\pi_i^n))$$

BM does **not** have bounded variation, but Ito integral is well-defined (in prob) if $\mathbb{E}[\int_0^t f_s^2 \, \mathrm{d}s] < \infty$ and f adapted stoch process

Idea: approximate paths of f by sequence of simple/stepwise-linear functions

Nice Properties: $I_{a,b}(f) := \int_a^b f_s \, \mathrm{d}B_s$ sat.

(1) linearity: $I_{a,b}(\alpha f + \beta g) = \alpha I_{a,b}(f) + \beta I_{a,b}(g)$

(2) additivity: $I_{a,b}(f) = I_{a,c}(f) + I_{c,b}(f)$

(3) unbiasedness: $\mathbb{E}[I_{0,t}(f)] = 0$

(4) continuous martingale: $I_t(f) \equiv I_{0,t}(f)$; $\mathbb{E}[I_{t+h}(f) \mid \mathcal{F}_t] = I_t(f)$

(5) Ito isometry: $\mathbb{E}[\int_0^t f_s g_s \, \mathrm{d}s] = \mathbb{E}[I_t(f) I_t(g)]$

Ito Processes

Ito Process: any process X (on filtered prob space) that can be written as

$X_t = X_0 + \int_0^t \mu_s \, ds + \int_0^t \sigma_s \, dB_s$

where B_t is BM, μ_t, σ_t are adapted processes s.t. $\int_0^t |\mu_s| \, ds, \int_0^t \sigma_s^2 \, ds < \infty$ for all t

Differential form: $dX_t = \mu_t \, dt + \sigma_t \, dB_t$

X is continuous and adapted process

Notation: X_t^x to denote process satisfying $dX_t = \mu_t \, dt + \sigma_t \, dB_t$ and $X_0 = x$

Application: Optimal Pricing

Stopping Problem

DM can stop at any point in time and choose whether or not to buy the product

Expected payoff if stop at t and (i) buy: u_t; in (ii) not buy: 0

Value upon stopping at t: $v_t := u_t^+ \equiv \max\{0, \mathbb{E}[u \mid \mathcal{F}_t]\}$

Note $\mathrm{d}u_t = \sigma\, \mathrm{d}B_t$: u is Ito process and martingale; v_t is a *submartingale*:

$\mathbb{E}[v_{t+h} \mid \mathcal{F}_t] = \mathbb{E}[\max\{0, u_{t+h}\} \mid \mathcal{F}_t] \geq \max\{0, \mathbb{E}[u_{t+h} \mid \mathcal{F}_t]\} = v_t$

(if f is convex, X is a martingale, and $\mathbb{E}[|f(X_t)|] < \infty$, then $Y = f(X)$ is submartingale)

DM's Stopping time τ: may depend not only on time but also on what was observed thus far

i.e. stopping time is actually a r.v.

Value given stopping time τ: v_τ

Take it that if never stop payoff is $-\infty$

Stopping Time

Stopping Time: $\tau : \Omega \to [0, \infty]$ adapted to $\mathbb{F}$ s.t. $\{\tau \leq t\} := \{\omega \in \Omega : \tau(\omega) \leq t\} \in \mathcal{F}_t \ \forall t \in T$

Stopped Process: $X_{\tau \wedge t}, X^\tau$

Result: if X is a sub/super/martingale, then so is X^τ

Uniform Integrability of X: $\lim_{k\to\infty} \sup_{t\in T} \mathbb{E}[|X_t| \mathbf{1}_{|X_t| \geq k}] = 0$

Martingale Convergence Theorem

Let X be cadlag martingale. X is UI off X_t converges a.s. and in L^1 to r.v. X_∞ and $X_t = \mathbb{E}[X_\infty \mid \mathcal{F}_t] \ \forall t$

Optional Stopping/Sampling Theorem

Let $\sigma \leq \tau$ be stopping times and X cadlag martingale. For all $t > 0$, $\mathbb{E}[X_{\tau\wedge t}] = X_0$.

If moreover

(a) σ, τ a.s. bounded, then X_σ, X_τ are integrable and $\mathbb{E}[X_\tau \mid \mathcal{F}_\sigma] = X_\sigma$ a.s.

(b) τ a.s. finite, $\mathbb{E}[|X_{\tau\wedge t}|] < \infty$, and $\lim_{t\to\infty} \mathbb{E}[|X_t| \mathbf{1}_{\tau > t}] = 0$ a.s., then $\mathbb{E}[X_\tau] = X_0$ a.s.

(super $\leq$, sub $\geq$)

Application: Optimal Pricing

Optimal Stopping

Benefit to learning: v_t is submartingale Cost to learning: time, flow cost c

DM chooses stopping time $\tau \in \mathbb{T}$ $V(u_0) := \sup_{\tau} \mathbb{E}[v_\tau - c\tau]$

Goal: get stopping policy 'out of' value function

Application: Optimal Pricing

Optimal Stopping: $$V(u_0) := \sup_{\tau} \mathbb{E}[v_\tau - c\tau]$$

Some Properties of the Value Function

(i) Value function V is convex (hence continuous)

(ii) Alexandrov thm: left/right diff everywhere V'_+, V'_-, twice diff V'' a.e.

(iii) $V \geq v_0$ (stopping immediately is always feasible)

(iv) V is increasing in u_0:

$u_t - u_0 = \sigma B_t$ is (scaled) BM

For any stopping time τ,

$$u'_0 > u_0 \implies u'_0 + \sigma B_\tau - c\tau \geq u_0 + \sigma B_\tau - c\tau$$

$$\implies (u'_0 + \sigma B_\tau)^+ - c\tau = v'_\tau - c\tau \geq v_\tau - c\tau = (u_0 + \sigma B_\tau)^+ - c\tau$$

$$\implies \mathbb{E}[v'_\tau - c\tau] = V(u'_0, \tau) \geq V(u_0, \tau) = \mathbb{E}[v_\tau - c\tau]$$

hence $V(u'_0) = \sup_\tau V(u'_0, \tau) \geq \sup_\tau V(u_0, \tau) = V(u_0)$

Application: Optimal Pricing

Fixed Boundaries: $\exists \underline{u}, \overline{u}$ s.t. the earliest optimal stopping time
$\tau := \inf\{t \geq 0 \mid V(u_t) \leq u_t^+\} = \inf\{t \geq 0 \mid u_t \notin (\underline{u}, \overline{u})\}$

Figure 1 Example of $u(t)$

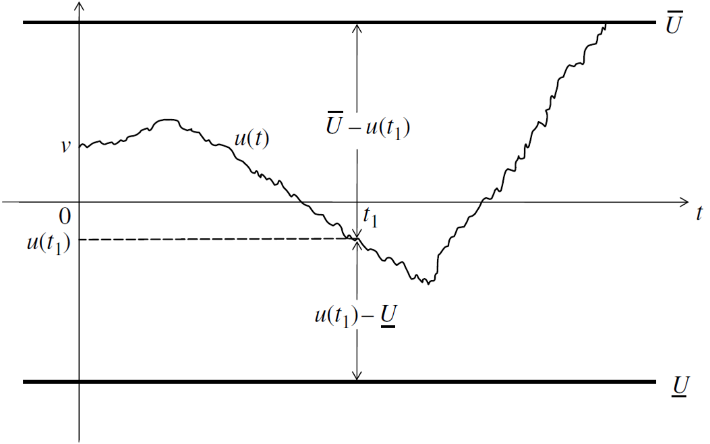

Application: Optimal Pricing

Fixed Boundaries: $\exists \underline{u}, \overline{u}$ s.t. the earliest optimal stopping time

$$\tau := \inf\{t \geq 0 \mid V(u_t) \leq u_t^+\} = \inf\{t \geq 0 \mid u_t \notin (\underline{u}, \overline{u})\}$$

Proof

Define $\underline{u} := \sup\{u \in \mathbb{R} \mid V(u) \leq 0\}$ and $\overline{u} := \inf\{u \in \mathbb{R} \mid V(u) \leq u\}$

WTS $\overline{u}$ and $\underline{u}$ are well-defined

(i) As $V(u) \geq u^+ \geq 0$, $\overline{u} \geq 0$; As $V(\underline{u}) = 0$, $\underline{u} \leq 0$

(ii) If neither $\underline{u}, \overline{u}$ were well-defined, then DM searches forever and incurs $-\infty$ exp. payoff, implying stopping immediately is strictly better

(iii) Suppose that only $\underline{u}$ is well-defined. For any $u > \underline{u}^+$, $V(u) > u^+$ which implies DM hasn't stopped at u

$\Longrightarrow$ conclude DM never chooses to buy

$\Longrightarrow$ if $u_0 > \underline{u}^+$ DM pays cost of time for u_t reach $\underline{u}$ and gets $0 - c\mathbb{E}[\tau]$ (things are actually worse...)

Application: Optimal Pricing

Fixed Boundaries: $\exists \underline{u}, \overline{u}$ s.t. the earliest optimal stopping time

$$\tau := \inf\{t \geq 0 \mid V(u_t) \leq u_t^+\} = \inf\{t \geq 0 \mid u_t \notin (\underline{u}, \overline{u})\}$$

Proof

(iv) **Fact**: if $\tau := \inf\{t \geq 0 \mid u_t = \tilde{u}\}$ and $u_0 \neq \tilde{u}$, as $u_t - u_0$ is a scaled BM, then $\mathbb{E}[\tau] = \infty$ (even though $\mathbb{P}(\tau < \infty) = 1$)

(v) Suppose then that only $\overline{u}$ is well-defined.

Then DM never stops when $u_t < 0$ and, for $u_0 < \overline{u}$, in expectation pays $-\infty$ in time costs

$\Longrightarrow$ better to stop immediately

Application: Optimal Pricing

Fixed Boundaries: $\exists \underline{u}, \overline{u}$ s.t. the earliest optimal stopping time

$\tau := \inf\{t \geq 0 \mid V(u_t) \leq u_t^+\} = \inf\{t \geq 0 \mid u_t \notin (\underline{u}, \overline{u})\}$

Proof WTS $\forall u \leq \underline{u}$ or $u \geq \overline{u}$, $V(u) = u^+$ (ie DM stops)

Suppose $u > \overline{u}$ but $V(u) > u$ **Draw a picture**

$\because V(\overline{u}) = \overline{u}$, $V(u) \geq u^+$, and V is convex, then for any $u' > u$,

$$V(u') \geq V(u) + \frac{V(u) - V(\overline{u})}{u - \overline{u}}(u' - u) = V(u) + \frac{V(u) - \overline{u}}{u - \overline{u}}(u' - u) > u + \frac{u - \overline{u}}{u - \overline{u}}(u' - u) = u'$$

This implies that if $u_0 \geq u$, then DM stops only (and buys) at $\overline{u}$, yielding an exp. payoff $\overline{u} - c\mathbb{E}[\tau] = -\infty$

Since V is nondec., $0 = u^+ \leq V(u) \leq V(\underline{u}) = 0$ for all $u \leq \underline{u}$

Application: Optimal Pricing

Purchase prob.: (assuming $u_0 \in (\underline{u}, \overline{u})$)

(1) u_t martingale, u_τ martingale, $u_\tau \in \{\underline{u}, \overline{u}\} \implies u_\tau$ bounded martingale

(2) optional stopping: $u_0 = \mathbb{E}[u_\tau] = q\overline{u} + (1-q)\underline{u} \implies \mathbb{P}(\text{buy}) = \mathbb{P}(u_\tau = \overline{u}) = q = \frac{u_0 - \underline{u}}{\overline{u} - \underline{u}}$

Expected stopping time: (assuming $u_0 \in (\underline{u}, \overline{u})$)

Recall $u_t = \sigma B_t$ and that $B_t \sim N(0,t)$, $(B_{t+h} - B_t) \sim N(0,h)$, $B_t \perp\!\!\!\perp (B_{t+h} - B_t)$

$y_t := u_t^2 - \sigma^2 t$ is a martingale:

$$\begin{aligned}\mathbb{E}[y_{t+h} \mid \mathcal{F}_t] &= \mathbb{E}[u_{t+h}^2 \mid \mathcal{F}_t] - \sigma^2(t+h) \\ &= \mathbb{E}[(u_{t+h} - u_t)^2 \mid \mathcal{F}_t] + 2u_t\mathbb{E}[u_{t+h} - u_t \mid \mathcal{F}_t] + u_t^2 - \sigma(t+h) \\ &= \sigma^2 h + u_t^2 - \sigma^2(t+h) = u_t^2 - \sigma^2 t\end{aligned}$$

y_τ is a martingale; optional stopping implies

$$\mathbb{E}[u_\tau^2 - \sigma^2\tau] = u_0^2 = q\overline{u}^2 + (1-q)\underline{u}^2 - \sigma^2\mathbb{E}[\tau] = \frac{u_0 - \underline{u}}{\overline{u} - \underline{u}}(\overline{u} - \underline{u})(\overline{u} + \underline{u}) + \underline{u}^2 - \sigma^2\mathbb{E}[\tau]$$

$$\iff \mathbb{E}[\tau] = \frac{(\overline{u} - u_0)(u_0 - \underline{u})}{\sigma^2}$$

Application: Optimal Pricing

What are exactly $\overline{u}, \underline{u}$?

A possible way forward

For *any* $\Delta > 0$, $V(u_t) \geq \mathbb{E}[V(u_{t+\Delta})] - c\Delta$

For any $\varepsilon > 0$, there is $h > 0$ such that $\forall \Delta \leq h$ s.t. $V(u_t) - \varepsilon \leq \mathbb{E}[V(u_{t+\Delta})] - c\Delta$

Suggests something to taking derivative and getting an HJB so that " dV = [some clean function]" and we can then study properties of V from this clean function

Infinitesimal Generator

Given stoch. proc. Y, ***What is*** $\boldsymbol{d}f(y)$***?***

$Y^y_{t\,t}$ stoch. proc. where $Y_0 = y$ (starting at y)

Infinitesimal Generator

$$(Af)(y) := \lim_{\mathrm{d}t \to 0} \frac{\mathbb{E}[f(Y^y_{\mathrm{d}t})] - f(y)}{\mathrm{d}t}$$

(Note: we can always think of f as a function of a single multi-dim process $Y_t = (t, X_t)$)

Exact form of Af depends on the process

Infinitesimal Generator and Ito's Lemma

Given stoch. proc. Y, ***What is*** $\boldsymbol{d}f(y)$***?***

$f \in \mathcal{C}^2 : (t, x) \mapsto f(t, x) \in \mathbb{R}$

Ito process $\mathrm{d}X_t = \mu_t \, \mathrm{d}t + \sigma_t \, \mathrm{d}B_t$

Ito's Lemma

$$\mathrm{d}f = f_t \, \mathrm{d}t + \mu_t f_x \, \mathrm{d}t + \frac{\sigma_t^2}{2} f_{xx} \, \mathrm{d}t + \sigma_t f_x \, \mathrm{d}B_t$$

Ito's lemma $\approx$ chain rule of stochastic differentiation

Extends to multidim Ito processes: $Af = (\nabla f)^T \mu + \frac{1}{2} \mathrm{tr}((\nabla^2 f)\sigma\sigma^T)$

Infinitesimal Generator and Ito's Lemma

Ito's Lemma

$$\mathrm{d}f = f_t\,\mathrm{d}t + \mu_t f_x\,\mathrm{d}t + \frac{\sigma_t^2}{2} f_{xx}\,\mathrm{d}t + \sigma_t f_x\,\mathrm{d}B_t$$

Intuition

(1) 2nd-order Taylor expansion of f at (t, x):

$$\mathrm{d}f = f_t\,\mathrm{d}t + \tfrac{1}{2} f_{tt}(\mathrm{d}t)^2 + f_x\,\mathrm{d}x + \tfrac{1}{2} f_{xx}(\mathrm{d}x)^2 + \cdots$$

(2) Sub $\mathrm{d}x$:

$$\mathrm{d}f = f_t\,\mathrm{d}t + \tfrac{1}{2} f_{tt}(\mathrm{d}t)^2 + f_x(\mu_t\,\mathrm{d}t + \sigma_t\,\mathrm{d}B_t) + \tfrac{1}{2} f_{xx}(\mu_t^2(\mathrm{d}t)^2 + \mu_t\sigma_t\,\mathrm{d}t\,\mathrm{d}B_t + \sigma_t^2(\mathrm{d}B_t)^2) + \cdots$$

As $\mathrm{d}t \to \mathbf{0}$, $(\mathrm{d}t)^2$, $\mathrm{d}t\,\mathrm{d}B_t$, and higher-order terms go to zero fast

Replace $(\mathrm{d}t)^2 = \mathrm{d}t\,\mathrm{d}B_t =$ higher order terms $= \mathbf{0}$, and $(\mathrm{d}B_t)^2 = \mathrm{d}t$ and rearrange terms

Infinitesimal Generator and Ito's Lemma

Ito's Lemma

$$df = f_t \, dt + \mu_t f_x \, dt + \frac{\sigma_t^2}{2} f_{xx} \, dt + \sigma_t f_x \, dB_t$$

Infinitesimal Generator for Ito Processes

$dY_t = \mu_t \, dt + \sigma_t \, dB_t; \ f \in \mathcal{C}^2$

$$\begin{aligned}
(Af)(y) &:= \lim_{dt \to 0} \frac{\mathbb{E}[f(Y^y_{dt})] - f(y)}{dt} \\
&= \lim_{dt \to 0} \frac{\mathbb{E}[f_t \, dt + \mu_t f_x \, dt + \frac{\sigma_t^2}{2} f_{xx} \, dt + \sigma_t f_x \, dB_t]}{dt} \\
&= \lim_{dt \to 0} \frac{f_t \, dt + \mu_t f_x \, dt + \frac{\sigma_t^2}{2} f_{xx} \, dt + \sigma_t f_x \mathbb{E}[\, dB_t]}{dt} \\
&= f_t + \mu_t f_x + \frac{\sigma_t^2}{2} f_{xx}
\end{aligned}$$

Dynkin's Formula

Dynkin's Formula

Let (1) be an Ito process, (2) $\tau \in \mathbb{T}$ stopping time with $\mathbb{E}[\tau] < \infty$, and (3) $f \in \mathcal{C}^2$ with compact support. Then

$$\mathbb{E}[f(X^x_\tau)] = f(x) + \mathbb{E}\left[\int_0^\tau (Af)(X^x_t)\,\mathrm{d}t\right]$$

Furthermore, if $\tau := \inf\{t \geq 0 \mid X_t \notin D\}$ for some bounded set $D \subset \mathbb{R}^n$, one can dispense with compact support of f

Stopping an Ito Process

Ito Process $Y_t = (t, X_t)$

$$dY_t = \mu_t \, dt + \sigma_t \, dB_t \quad \text{where}$$

$$\mu_t := \mu(Y_t) \quad \text{and} \quad \sigma_t := \sigma(Y_t)$$

Payoffs

Flow payoff $f(Y_t)$

Terminal payoff/scrap value $g(Y_t)$

Value Function given stopping time $\tau := \inf\{t \geq 0 \mid Y_t \notin D\}$ (where boundary ∂D is 'regular')

$$V(y) := \mathbb{E}\left[\int_0^\tau f(Y_t^y)\, dt + g(Y_\tau^y)\right]$$

Infinitesimal Generator

For $h \in \mathcal{C}^2 : y \mapsto h(y)$, from Ito's lemma

$$(Ah)(y) = h_t(y)\, dt + \mu(y) h_x(y)\, dt + \frac{1}{2}(\sigma(y))^2 h_{xx}(y)\, dt$$

Differential Characterization (HJB)

Theorem (11.2.1 Oksendal)

Suppose

(i) $V \in \mathcal{C}^2(D) \cap C^0(\overline{D})$ and bounded

(ii) $\mathbb{E}\left[|V(Y_\theta^y)| + \int_0^\theta |(AV)(Y_t^y)|\, dt\right] < \infty \quad \forall$ *bounded* stopp. times $\theta \leq \tau \;\; \forall y \in D$, $\forall a \in \mathcal{A}$

(iii) $\tau < \infty$ a.s. for all starting points $Y_0 = y \in D$

Then,

$$f(y) + (AV)(y) = 0 \;\forall y \in D \tag{HJB}$$

$$V(y) = g(y) \;\forall y \in \partial D \tag{BC}$$

Differential Characterization (HJB)

Intuition

BC is obvious

For HJB, note that $V(y) = \mathbb{E}\left[\int_0^\tau f(Y_t^y)\,\mathrm{d}t + g(Y_\tau^y)\right]$

For $y \in D$, since for vanishing $\mathrm{d}t$,

$$V(y) = \mathbb{E}\left[\int_0^{\mathrm{d}t} f(Y_{\mathrm{d}t}^y)\right] + \mathbb{E}[V(Y_{\mathrm{d}t}^y)] \Longleftrightarrow \mathbb{E}[V(Y_{\mathrm{d}t}^y)] - V(y) = -\mathbb{E}\left[\int_0^{\mathrm{d}t} f(Y_{\mathrm{d}t}^y)\right],$$

then $(AV)(y) = \lim_{\mathrm{d}t\to 0} \frac{\mathbb{E}[V(Y_{\mathrm{d}t}^y)] - V(y)}{\mathrm{d}t} = -f(Y_{\mathrm{d}t}^y)$

i.e. $f(y) + (AV)(y) = 0$

Application: Optimal Pricing

Back to our DM

$f(t,u) = -c; g(u) = u^+, u_t = \sigma B_t, V(u)$

HJB: $-c + \frac{1}{2}\sigma^2 V_{uu} = 0 \iff V(u) = \frac{c}{\sigma^2}u^2 + k_1 u + k_2$

Value Matching, Smooth Pasting

(BC/'Value Matching') If $V \in \mathcal{C}^0$: $V(\underline{u}) = \underline{u}^+ = 0$ and $V(\overline{u}) = \overline{u}^+ = \overline{u}$

('Smooth Pasting') If $V \in \mathcal{C}^1$: $V'(\underline{u}) = v'(\underline{u}) = 0$ and $V'(\overline{u}) = v'(\overline{u}) = 1$

$$V(\underline{u}) = V'(\underline{u}) = V(\overline{u}) - \overline{u} = V'(\overline{u}) - 1 = 0$$

$$\Longrightarrow \overline{u} = -\underline{u} = \sigma^2/(4c), \; k_1 = 1/2, \; k_2 = \sigma^2/(16c)$$

Stopping thresholds don't depend on starting point u_0 (no dependence on time)

Application: Optimal Pricing

$\overline{u} = -\underline{u} = \sigma^2/(4c)$

Purchase prob: $D(u_0) = 1/2 + 1/2 u_0/\overline{u}$ if $u_0 \in (\underline{u}, \overline{u})$

$p = 1$ if $u_0 \geq \overline{u}$, and $p = 0$ if $u_0 \leq \underline{u}$

Exp. stop. time: $\mathbb{E}[\tau] = 1_{u_0 \in (\underline{u}, \overline{u})}(\sigma^2/(16c^2) - u_0^2/\sigma^2)$

Firm's problem:

cost $k \in [0, \overline{u})$, demand $D(u_0 - p)$, profit $(p - k)D(u_0 - p)$

$u_0 - p \geq \underline{u}$ (ow no sale) and $p \geq k$ (ow no profit)

$p \in [k, u_0 - \underline{u}]$; if $k + \underline{u} \geq u_0$, no sale

Concave problem, unique solution:

if $k + \underline{u} < u_0 < k + 3\overline{u}$ price at $p^* = (u_0 + k + \overline{u})/2$

if $u_0 > k + 3\overline{u}$ price at $p^* = u_0 - \underline{u}$

Application: Optimal Pricing

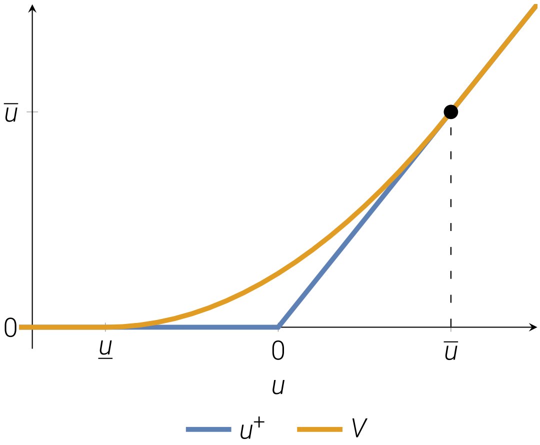

Application: Optimal Pricing

Comparative Statics

- Purchase prob: if u_0 is large enough, $\uparrow c, \downarrow \boldsymbol{\sigma}$; ow $\downarrow c, \uparrow \boldsymbol{\sigma}$
  (recall: if u_0 is very high, $u_0 - p^* = \overline{u}$ and purchase prob = 1)
- p^*: $\uparrow u_0, k$, but non-monotone (quasiconvex) in $\boldsymbol{\sigma}, c$
- Profit: $\uparrow u_0, \downarrow k$, but non-monotone (quasiconvex) in $\boldsymbol{\sigma}^2/(4c)$
  decreasing in $\boldsymbol{\sigma}^2/(4c)$ iff $u_0 - k \geq \overline{u} = \boldsymbol{\sigma}^2/(4c)$
  (intuition: at $p = k$, with $u_0 - p \geq \overline{u}$, seller breaks even and gets sale wp1, hence can make profit;
  higher $\overline{u}$ curtails seller's markup)
- $\mathbb{E}[\tau]$ and DM's EU nonmonotone in $\boldsymbol{\sigma}, c$

Moral of the story: eqm comparative statics may differ substantially from indiv DM

Overview I

1. Continuous-Time Preliminaries via Optimal Pricing with Learning

2. Simple Optimal Stopping Problems
   - Stopping
   - Optimal Stopping
     - Canonical Forms
     - Optimal Constant Thresholds
   - Evolution of Beliefs
   - Value Function
   - Speed-Accuracy Revisited
   - A Recap of Wald's Equation/Identities

Simple Stopping Problems

Binary choice problem

DM chooses $\alpha \in A = \{a, b\}$; Unknown state $\theta \in \Theta = \{a, b\}$; equal prob. $p_0 = 1/2$

Payoffs $u : A \times \Theta \to \mathbb{R}$; $u(\theta, \theta) > u(\neg\theta, \theta)$

Right/wrong answers, binary type models

Sequential sampling

$X_t = \mu s(\theta)t + \sigma B_t$; $s(\theta) := (1\{\theta = a\} - 1\{\theta = b\})$

B_t std Brownian motion; given binary drift, $\pm\mu$ wlog

Problem complexity μ/σ: rate of info flow

Application: Understand how behavior is affected by problem complexity

Important methodological points:

(1) Wald's equation

(2) Characterize evolution of beliefs (filtering)

Simple Stopping Problems

Exogenous Stopping: Suppose DM stops if $p_t = \overline{p} = 1 - \underline{p}$

Posterior Beliefs: $p_t = \dfrac{p_0 \phi((X_t - \mu t)/(\sigma\sqrt{t}))}{p_0 \phi((X_t - \mu t)/(\sigma\sqrt{t})) + (1 - p_0)\phi((X_t + \mu t)/(\sigma\sqrt{t}))}$

Posterior Log-Odds Beliefs: $\ell_t = \ln(p_t/(1 - p_t)) = 2\dfrac{\mu}{\sigma^2} X_t$

$\Longrightarrow$ DM stops when $X_t \notin (\underline{X}, \overline{X})$, $\overline{X} = \frac{\sigma^2}{2\mu}\ln(\overline{p}/(1 - \overline{p})) = -\frac{\sigma^2}{2\mu}\underline{X} = -\ln(\underline{p}/(1 - \underline{p}))$

Stopping time $\tau := \inf\{t \geq 0 \mid X_t \notin (\underline{X}, \overline{X})\}$

Simple Stopping Problems: Accuracy

Exogenous Stopping: Suppose DM stops if $p_t = \overline{p} = 1 - \underline{p}$

Choice Probability: $P_a \equiv \mathbb{P}(\text{choose } a) = \mathbb{P}(X_\tau = \overline{X})$

Given $\theta = a$ (dependence omitted),

Moment Generating Function (MGF) for Gaussian:

$M_X(\lambda) := \int \exp(\lambda x)\phi((x - \mu)/\sigma)\, dx$

$Z_t := M_X(\lambda)^{-t}\mathbb{E}[\exp(\lambda X_t)]$ is a martingale

Note: for $\lambda^* = -2\mu/\sigma^2$, $M_X(\lambda^*) = 1$

Optional stopping: $\mathbb{E}[Z_\tau] = \mathbb{E}[M_X(\lambda^*)^{-\tau}]\mathbb{E}[\exp(\lambda^* X_\tau)] = \mathbb{E}[\exp(\lambda^* X_\tau)] = Z_0 = 1$

Then, as $\overline{X} = -\underline{X}$ $\quad 1 = \mathbb{E}[\exp(\lambda^* X_\tau)] = P_a \exp(\lambda^* \overline{X}) + (1 - P_a) \exp(\lambda^* \underline{X})$

$$\iff P_a = \frac{(1 - \exp(\lambda^* \underline{X}))}{\exp(\lambda^* \overline{X}) - \exp(\lambda^* \underline{X})} = (1 + \exp(-2\mu/\sigma^2 \overline{X}))^{-1}$$

Accuracy

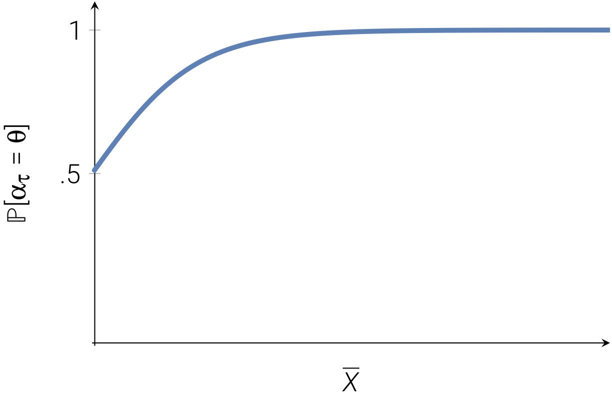

Choice Probability: $P_a \equiv \mathbb{P}(\text{choose } a) = \mathbb{P}(X_\tau = \overline{X})$ (given $\theta = a$),

$\Longrightarrow$ Logit choice!

$$P_a = (1 + \exp(-2\mu/\sigma^2 \overline{X}))^{-1}$$

$\mathbb{P}(\alpha_\tau = \theta)$ increases in $\overline{X}, -\underline{X}$

Accuracy

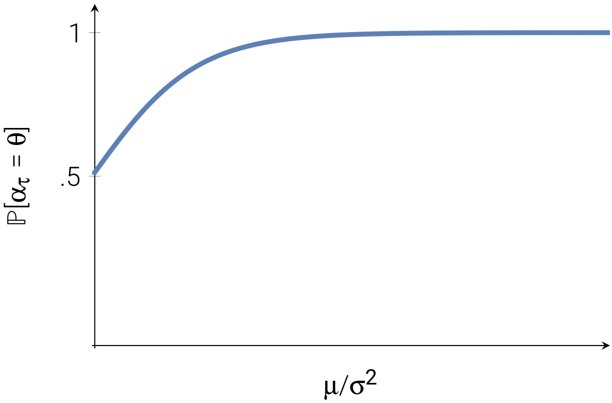

Choice Probability: $P_a \equiv \mathbb{P}(\text{choose } a) = \mathbb{P}(X_\tau = \overline{X})$ (given $\theta = a$),

$\Longrightarrow$ Logit choice!

$$P_a = (1 + \exp(-2\mu/\sigma^2 \overline{X}))^{-1}$$

$\mathbb{P}(\alpha_\tau = \theta)$ increases in $\overline{X}$, $-\underline{X}$, and increases in μ/σ^2

Simple Stopping Problems: Speed

Exogenous Stopping: Suppose DM stops if $p_t = \overline{p} = 1 - \underline{p}$

Choice Probability: $P_a \equiv \mathbb{P}(\text{choose } a) = \mathbb{P}(X_\tau = \overline{X}) = (1 + \exp(-2\mu\overline{X}))^{-1}$

Expected Stopping Time: Given $\theta = a$ (dependence omitted) and for simplicity (wlog) setting $\sigma = 1$

$\mathbb{E}[X_t] = \mu t$; optional stopping theorem yields

$\mathbb{E}[\tau]\,\mu = \mathbb{E}[X_\tau] = P_a\overline{X} + (1 - P_a)\underline{X} = (2P_a - 1)\overline{X} = \overline{X}\tanh(\overline{X}\mu)$

where $\tanh(x) = \frac{\exp(2x)-1}{\exp(2x)+1}$

Expected Stopping Time

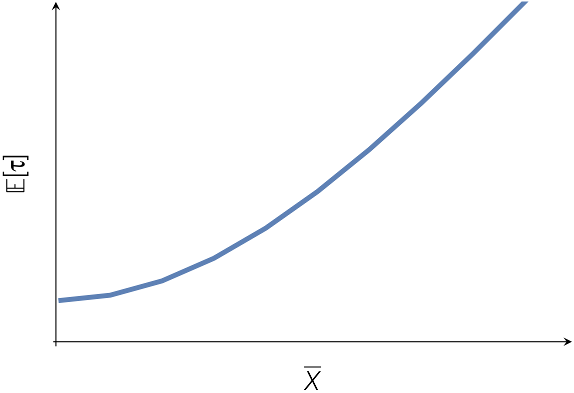

Expected Stopping Time: $\mathbb{E}[\tau] = (\overline{X})/\mu \tanh(\overline{X}\mu)$

where $\tanh(x) = \frac{\exp(2x)-1}{\exp(2x)+1}$

$\mathbb{E}[\tau]$ increases in $\overline{X}$, $-\underline{X}$

Expected Stopping Time

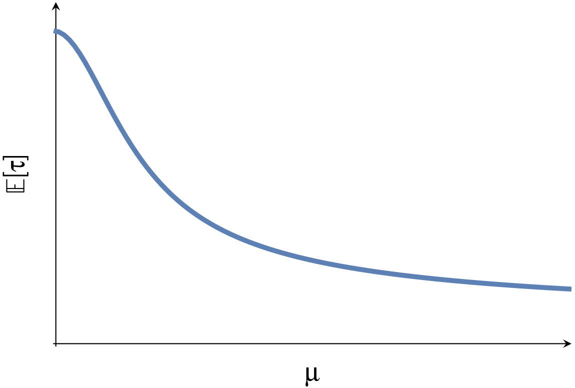

Expected Stopping Time: $\mathbb{E}[\tau] = (\overline{X})/\mu \tanh(\overline{X}\mu)$
where $\tanh(x) = \frac{\exp(2x)-1}{\exp(2x)+1}$

$\mathbb{E}[\tau]$ increases in $\overline{X}$, $-\underline{X}$, and decreases in μ

Simple Optimal Stopping Problems

Binary choice problem (Gonçalves 23 WP)

DM chooses $\alpha \in A = \{a, b\}$; Unknown state $\theta \in \Theta = \{a, b\}$; equal prob. $p_0 = 1/2$

Right/wrong answers, binary type models

Payoffs $u : A \times \Theta \to \mathbb{R}$; $u(\theta, \theta) > u(\neg\theta, \theta)$

Expected payoffs: $u(\alpha, p) := \mathbb{E}_p[u(\alpha, \theta)]$; $v(p) := \max_\alpha u(\alpha, p)$

Sequential sampling

$X_t = \mu s(\theta)t + \sigma B_t$; $s(\theta) := (1\{\theta = a\} - 1\{\theta = b\})$

B_t std Brownian motion; given binary drift, $\pm\mu$ wlog

Problem complexity μ/σ: rate of info flow

Optimal stopping

$$V(p_0) := \sup_\tau \mathbb{E}_{p_0}[v(p_\tau) - c\tau]; \qquad \text{s.t. } X_t = \mu s(\theta)t + \sigma B_t;\ X_0 = 0$$

Earliest optimal stopping time $\tau := \inf\{t \geq 0 \mid V(p_\tau) = v(p_\tau)\}$

Canonical Forms

Rewriting payoffs

Indifference $\tilde{p} : u(a, \tilde{p}) = u(b, \tilde{p})$

Stakes $\delta := \sum_\theta |u(a, \theta) - u(b, \theta)|$

Risky/Safe:

$$u_{\mathbf{a}}(\alpha, p) := 1\{\alpha = \mathbf{a}\}(p - \tilde{p}) \qquad u_{\mathbf{b}}(\alpha, p) := 1\{\alpha = \mathbf{b}\}(\tilde{p} - p)$$

$$v_{\mathbf{a}}(\alpha, p) := (p - \tilde{p})^+ \qquad v_{\mathbf{b}}(\alpha, p) := (\tilde{p} - p)^+$$

$$u(\alpha, p) = \delta u_a(\alpha, p) + u(b, p) = \delta u_b(\alpha, p) + u(a, p)$$

$$v(p) = \delta v_a(p) + u(b, p) = \delta v_b(p) + u(a, p) + u(a, p)$$

Canonical Forms

Canonical Forms

Remark

$$\arg\max_{\tau} \mathbb{E}_{p_0}[v(p_\tau) - c\tau] = \arg\max_{\tau} \mathbb{E}_{p_0}[v_a(p_\tau) - c/\delta\tau] = \arg\max_{\tau} \mathbb{E}_{p_0}[v_b(p_\tau) - c/\delta\tau]$$

Proof intuition

- For optimal τ: $\mathbb{P}_0[\tau = \infty] = 0$
- Affine transformations and optional stopping imply

$$\begin{aligned} \mathbb{E}_{p_0}[v(p_\tau) - c\tau] &= \delta\mathbb{E}_{p_0}[(p_\tau - \tilde{p})^+ - c/\delta\tau] + \mathbb{E}_{p_0}[u(b, p_\tau)] \\ &= \delta\mathbb{E}_{p_0}[v_a(p_\tau) - c/\delta\tau] + u(b, p_0) \end{aligned}$$

Focus on $V_\alpha(p_0) := \sup_\tau \mathbb{E}_{p_0}[v_\alpha(p_\tau) - c\tau]$, $\alpha = a, b$

Optimal Stopping

Notation: $\overline{\ell}, \underline{\ell}$: log-odds of $\overline{p}, \underline{p}$; $\quad \tilde{c} \equiv \dfrac{c}{2\delta}$

Proposition

Earliest optimal stopping time given by

$$\tau := \inf\{t \geq 0 \mid V(p_t) = v(p_t)\} = \inf\{t \geq 0 \mid p_t \notin (\underline{p}, \overline{p})\}$$

where optimal stopping (log-odds) beliefs $(\underline{\ell}, \overline{\ell})$ are unique solution to

$$\left(\frac{\mu}{\sigma}\right)^2 = \frac{\tilde{c}}{\tilde{p}}\left[\left(\exp(\overline{\ell}) + \overline{\ell}\right) - \left(\exp(\underline{\ell}) + \underline{\ell}\right)\right] \qquad \text{(eqn. 1)}$$

$$= \frac{\tilde{c}}{1-\tilde{p}}\left[\left(\overline{\ell} - \exp(-\overline{\ell})\right) - \left(\underline{\ell} - \exp(-\underline{\ell})\right)\right] \qquad \text{(eqn. 2)}$$

Proof: (1) constant thresholds, (2) characterize thresholds

Constant Thresholds

Remark

V_a is increasing in p and V_b is decreasing in p

Lemma

$\exists \overline{p}, \underline{p} \in [0,1]$ s.t. earliest optimal stopping time $\tau := \inf\{t \geq 0 \mid p_t \notin (\underline{p}, \overline{p})\}$

Constant Thresholds

Remark

V_a is increasing in p and V_b is decreasing in p

Lemma

$\exists \overline{p}, \underline{p} \in [0, 1]$ s.t. earliest optimal stopping time $\tau := \inf\{t \geq 0 \mid p_t \notin (\underline{p}, \overline{p})\}$

Proof intuition

- $\underline{p} := \max\{p \in [0, 1] : V_a(p) \leq 0\}$ $\overline{p} := \min\{p \in [0, 1] : V_b(p) \leq 0\}$
- $V_a(p) > 0$ iff $p > \underline{p}$ and $V_b(p) > 0$ iff $p < \overline{p}$
- $p \leq \underline{p}$: $V_a(p) = 0 = u_a(b, p) \iff V(p) = u(b, p) \implies$ stop and choose b
- $p \geq \overline{p}$: $V_b(p) = 0 = u_b(a, p) \iff V(p) = u(a, p) \implies$ stop and choose a
- $p \in (\underline{p}, \overline{p})$: $V(p) > v(p) \implies$ continue

Canonical Forms

Evolution of Beliefs

Evolution of beliefs

$$dp_t := 2\frac{\mu}{\sigma}p_t(1-p_t)\,d\tilde{B}_t$$

Intuition

Posterior Beliefs: $p_t = \dfrac{p_0\phi((X_t-\mu t)/(\sigma\sqrt{t}))}{p_0\phi((X_t-\mu t)/(\sigma\sqrt{t}))+(1-p_0)\phi((X_t+\mu t)/(\sigma\sqrt{t}))}$

Posterior Log-Odds Beliefs: $\ell_t = \ln(p_t/(1-p_t)) = 2\dfrac{\mu}{\sigma^2}X_t$

Scaling: $dX_t = \mu_t\,dt + \sigma_t\,dB_t$ is observationally equivalent to $d\tilde{X}_t = \frac{\mu_t}{\sigma_t^2}\,dt + d\tilde{B}_t$

- $\sigma B_t = B_{t\sigma^2} \longrightarrow X_t = \mu t + B_{t\sigma^2}$
- Change time: $t\sigma^2 \mapsto s \longrightarrow \tilde{X}_s = \frac{\mu}{\sigma^2}s + B_s$
- If $\tilde{X}_t = X_t$, posterior beliefs are the same; wlog normalize $\sigma = 1$

Evolution of Beliefs

Evolution of beliefs

$$dp_t := 2\frac{\mu}{\sigma} p_t(1-p_t)\, d\tilde{B}_t$$

Intuition

Posterior Beliefs: $p_t = \dfrac{p_0\phi((X_t-\mu t)/(\sigma\sqrt{t}))}{p_0\phi((X_t-\mu t)/(\sigma\sqrt{t})) + (1-p_0)\phi((X_t+\mu t)/(\sigma\sqrt{t}))}$

Posterior Log-Odds Beliefs: $\ell_t = \ln(p_t/(1-p_t)) = 2\dfrac{\mu}{\sigma^2} X_t$

Scaling: wlog normalize $\sigma = 1$

Let $P(x) := (1+\exp(-2\mu x - \ell_0))^{-1}$; then $p_t = P(X_t)$

Note: $P'(x) = 2\mu P(x)(1-P(x))$ and $P''(x) = (2\mu)^2 P(x)(1-P(x))(1-2P(x))$

Taylor expand $P(X_{t+dt})$ and use that $(dB_t)^2 = dt$, $dB_t\, dt = 0 = (dt)^k = 0$ for $k \geq 3/2$

$$\begin{aligned} dp_t \text{"="} P(X_{t+dt}) - P(X_t) &\approx P'(X_t) dX_t + \frac{1}{2} P''(X_t)(dX_t)^2 \\ &= 2\mu P(X_t)(1-P(X_t))[dX_t - \mu(1-2P(X_t))\, dt] \end{aligned}$$

Evolution of Beliefs

Evolution of beliefs

$$dp_t := 2\frac{\mu}{\sigma}p_t(1-p_t)\,d\tilde{B}_t$$

Intuition

$$dp_t = 2\mu P(X_t)(1-P(X_t))[\,dX_t - \mu(1-2P(X_t))\,dt]$$

From DM's perspective, given information at t,

$$dX_t = (2p_t - 1)\mu\,dt + d\tilde{B}_t,$$

where $\tilde{B}_t$ is some std BM

Then *from DM's perspective*, $dp_t = 2\mu P(X_t)(1-P(X_t))\,d\tilde{B}_t$

Characterize Thresholds

Evolution of beliefs

$$dp_t := 2\frac{\mu}{\sigma} p_t(1 - p_t) d\tilde{B}_t$$

Lemma

Value function V given by unique (viscosity, later) solution to free-boundary problem:

$$0 = -c + 2(\mu/\sigma)^2 \left(p(1-p)\right)^2 V''(p) \text{ on } p \in (\underline{p}, \overline{p}) \qquad \text{(HJB)}$$

$$0 = V(p) - v(p) \text{ on } p \notin (\underline{p}, \overline{p}) \qquad \text{(BC/‘value matching’)}$$

implying

$$\lim_{p \uparrow \overline{p}} V'(p) = v'(\overline{p}) \text{ and } \lim_{p \downarrow \underline{p}} V'(p) = v'(\underline{p}) \qquad \text{(‘smooth pasting’)}$$

Optimal Stopping

Proposition

Earliest optimal stopping time given by

$$\tau := \inf\{t \geq 0 \mid V(p_t) = v(p_t)\} = \inf\{t \geq 0 \mid p_t \notin (\underline{p}, \overline{p})\}$$

where optimal stopping (log-odds) beliefs $(\underline{\ell}, \overline{\ell})$ are unique solution to

$$\left(\frac{\mu}{\sigma}\right)^2 = \frac{\tilde{c}}{\tilde{p}}\left[\left(\exp(\overline{\ell}) + \overline{\ell}\right) - \left(\exp(\underline{\ell}) + \underline{\ell}\right)\right] \qquad \text{(eqn. 1)}$$

$$= \frac{\tilde{c}}{1-\tilde{p}}\left[\left(\overline{\ell} - \exp(-\overline{\ell})\right) - \left(\underline{\ell} - \exp(-\underline{\ell})\right)\right] \qquad \text{(eqn. 2)}$$

Accuracy

Lemma

Under optimal stopping, accuracy $\mathbb{P}[\alpha_\tau = \theta]$ is increasing in μ/σ

Proof intuition

(1) **Log-odds beliefs $\Leftrightarrow$ process:** $\ell_t = \ln\left(\frac{\phi((X_t - \mu t)/\sigma\sqrt{t})}{\phi((X_t + \mu t)/\sigma\sqrt{t})}\right) = \frac{2\mu}{\sigma^2}X_t$

(2) **Optional stopping:** (bounded martingale) given θ, $\exp(-(2\mu s(\theta)/\sigma^2)X_\tau)$ is a martingale

$\implies \mathbb{E}[\exp(-s(\theta)\ell_\tau) \mid \theta] = \mathbb{E}[\exp(-(2\mu s(\theta)/\sigma^2)X_\tau) \mid \theta] = \mathbb{E}[\exp(-(2\mu s(\theta)/\sigma^2)X_0) \mid \theta] = 1$

(3) **Choice probabilities:** $\mathbb{P}[\alpha_\tau = a \mid \theta] = \mathbb{P}[\ell_\tau = \overline{\ell} \mid \theta]$ and

$1 = \mathbb{E}[\exp(-s(\theta)\ell_\tau) \mid \theta] = \mathbb{P}[\ell_\tau = \overline{\ell} \mid \theta]\exp(-s(\theta)\overline{\ell}) + \mathbb{P}[\ell_\tau = \underline{\ell} \mid \theta]\exp(-s(\theta)\underline{\ell})$

Closed form for $\mathbb{P}[\alpha_\tau = \theta]$: increases in $\overline{\ell}$ and decreases in $\underline{\ell}$ for $\theta = a, b$

Accuracy

Speed

Proposition

Optimal $\mathbb{E}[\tau]$ is non-monotone and single-peaked in μ/σ

Proof intuition

(1) $\mathbb{P}[\alpha_\tau = a] = \mathbb{P}[p_\tau = \bar{p}] = \mathbb{P}[X_\tau = \frac{\sigma^2}{2\mu}\bar{\ell}]$

(2) **Monotonicity of thresholds:** $(\bar{\ell}, \underline{\ell})$ monotone in μ/σ; $\lim_{\mu/\sigma\to\infty} \bar{\ell} = \infty, \lim_{\mu/\sigma\to\infty} \underline{\ell} = -\infty$

(3) **Complexity $\Leftrightarrow$ Thresholds:**

$$(\mu/\sigma)^2 = \frac{\tilde{c}}{\tilde{p}}\left((\exp(\bar{\ell}) + \bar{\ell}) - (\exp(\underline{\ell}) + \underline{\ell})\right) = \frac{c/\delta}{2(1-\tilde{p})}\left((\bar{\ell} - \exp(-\bar{\ell})) - (\underline{\ell} - \exp(-\underline{\ell}))\right)$$

(4) **Wald's identity:** (tbc)

$$s(\theta)\mu\mathbb{E}[\tau \mid \theta] = \mathbb{E}[X_\tau \mid \theta] = \mathbb{P}[X_\tau = \bar{X} \mid \theta]\bar{X} + \mathbb{P}[X_\tau = \underline{X} \mid \theta]\underline{X}$$

$$= \frac{1}{2\mu/\sigma^2}\left(\mathbb{P}[\alpha_\tau = a \mid \theta]\bar{\ell} + \mathbb{P}[\alpha_\tau = b \mid \theta]\underline{\ell}\right)$$

$$\mathbb{E}[\tau] =: T(\bar{\ell}, \underline{\ell})$$

Speed

Proposition

Optimal $\mathbb{E}[\tau]$ is non-monotone and single-peaked in μ/σ

Proof intuition

(4) **Nonmonotonicity of RT:**

$\overline{\ell}, -\underline{\ell} \uparrow \infty \implies T(\overline{\ell}, \underline{\ell}) \to 0$

$\overline{\ell}, -\underline{\ell} \downarrow \tilde{\ell} \equiv \text{log-odds}(\tilde{p}) \implies T(\overline{\ell}, \underline{\ell}) \to 0$

(5) **Logconcavity:** $T(\overline{\ell}, \underline{\ell})$ logconcave in $(\overline{\ell}, \underline{\ell})$

Implicit function: $\underline{\ell}(\overline{\ell}) \implies T(\overline{\ell}, \underline{\ell}(\overline{\ell}))$ logconcave

(6) **Single-peakedness:** $\overline{\ell}$ increasing in $\mu/\sigma \implies T(\mu/\sigma)$ nonmonotone and qconcave

Speed

(a) $\tilde{p} = 1/2$

(b) $\tilde{p} = 3/5$

Speed, Accuracy, and Complexity

Recap: accuracy increases in μ/σ; speed single-peaked in μ/σ

Corollary: inverse *U*-shaped relation speed-accuracy

Useful Results: Wald's Equation (Discrete Time)

Wald's Equation

Let $\{X_t\}_{t\in\mathbb{N}}$ be a real-valued stochastic process and τ a stopping time defined on $(\Omega, \mathcal{B}, \mathbb{F}, P)$.

If **(i)** $\mathbb{E}[X_t]$ is finite, **(ii)** $\mathbb{E}[X_t \mathbf{1}\{\tau \geq t\}] = \mathbb{E}[X_t]\mathbb{P}(\tau \geq t)$ for every t, and **(iii)** $\sum_{t\in\mathbb{N}} \mathbb{E}[|X_t|\mathbf{1}\{\tau \geq t\}] < \infty$,

then $S_\tau := \sum_{t=1}^{\tau} X_t$ and $T_\tau := \sum_{t=1}^{\tau} \mathbb{E}[X_t]$ have finite expectation and satisfy $\mathbb{E}[S_\tau] = \mathbb{E}[T_\tau]$.

If, moreover, $\mathbb{E}[X_t] = m \ \forall t$ and $\mathbb{E}[\tau] < \infty$,

then $\mathbb{E}[S_\tau] = \mathbb{E}[\tau]m$.

Note: the theorem does *not* require X_t to be iid.

For simplicity, we'll prove the theorem assuming X_t are in fact iid.

Proof

Let $Y_t := S_t - tm$.

Note that $\mathbb{E}[Y_{t+1} \mid \mathcal{F}_t] = S_t - tm + \mathbb{E}[X_{t+1} \mid \mathcal{F}_t] - m = Y_t$, therefore, Y_t is a martingale.

Since $\mathbb{E}[|Y_{t+1} - Y_t| \mid \mathcal{F}_t] = \mathbb{E}[|X_t - m|] < \infty$, we can apply the optional stopping theorem;

we then find that $0 = \mathbb{E}[Y_1] = \mathbb{E}[Y_\tau] = \mathbb{E}[S_\tau] - \mathbb{E}[\tau]m$

Useful Results: Wald's Identities

There are different versions of Wald's equation/identity, typically called **Wald's 1st/2nd/3rd identity** which follow from the optional stopping theorem

$\{X_t\}_{t\in\mathbb{N}}$: stochastic process s.t. X_t independent; $\mathbb{F}$ its natural filtration; τ adapted stopping time with $\mathbb{E}[\tau] < \infty$

(i) $S_t := \sum_{\ell=1}^{t} X_t$; (ii) $m_t := \sum_{\ell=1}^{t} \mathbb{E}[X_\ell]$; (iii) $v_t := \sum_{\ell=1}^{t} \mathbb{V}(X_\ell)$; (iv) $\phi(\theta) := \mathbb{E}[\exp(\theta X_1)]$ (v) $M_t^1 := S_t - m_t$; (vi) $M_t^2 := (S_t - m_t)^2 - v_t$; and (vii) $M_t^3 := \phi(\theta)^{-t}\exp(\theta S_t)$

1. If $\sup_t \mathbb{E}[|X_t|] < \infty$, then M_t^1 is a martingale and $\mathbb{E}[M_\tau^1] = \mathbb{E}[M_1^1] = 0$.
   In particular, if X_t are iid with mean $\mathbb{E}[X_t] = m$, then $\mathbb{E}[S_\tau] = m\mathbb{E}[\tau]$.
2. If $\sup_t \mathbb{E}[X_t^2] < \infty$, then M_t^2 is a martingale and $\mathbb{E}[M_\tau^2] = \mathbb{E}[M_1^2] = 0$.
   In particular, if X_t are iid with variance $\mathbb{V}(X_t) = \sigma^2$, then $\mathbb{E}[(S_\tau - m_\tau)^2] = \sigma^2\mathbb{E}[\tau]$.
3. If X_t are iid, the moment generating function $\phi(\theta) < \infty$, and τ is a.s. bounded or $M_t^3 \mathbf{1}\{\tau \geq t\} \leq \delta < \infty$ for all t, then M_t^3 is a martingale and $\mathbb{E}[\phi(\theta)^{-\tau}\exp(\theta S_\tau)] = 1$.

Back to Continuous Time

Wald's identity with BM:

$\mathbb{E}[\tau] < \infty$ (we know this)

Given θ, $X_t - \mu s(\theta) t = \sigma B_t$

$\mathbb{E}[B_\tau] = 0$ (optional stopping; details omitted)

$\mathbb{E}[X_\tau] = \mu s(\theta) \mathbb{E}[\tau]$

(May also be of use to recall another martingale: $B_\tau^2 - \tau$)