Searching

Duarte Gonçalves
University College London

Topics in Economic Theory

Overview

1. Pandora’s Problem
2. Martingales and etc.
3. Gittins Index
4. Web Search
5. Pricing with Pandora Consumers
6. Summary

Overview

Pandora

Directed Search/Info Acquisition

Before: when to stop.

Now: which info to acquire.

Pandora

Directed Search/Info Acquisition

Before: when to stop.

Now: which info to acquire.

Setup (Weitzman 1979 Ecta)

$T < \infty$ alternatives, $t \in [T] = \{\mathbf{1}, ..., T\}$ and an outside option $\mathbf{0}$;

Payoffs $X_t \sim F_t$, independent;

DM can pay cost c_t to learn X_t;

Recall: DM can stop at any time and pick best alternative explored thus far.

Pandora

Directed Search/Info Acquisition

- Before: when to stop.
- Now: which info to acquire.

Setup (Weitzman 1979 Ecta)

- $T < \infty$ alternatives, $t \in [T] = \{\mathbf{1}, ..., T\}$ and an outside option $\mathbf{0}$;
- Payoffs $X_t \sim F_t$, independent;
- DM can pay cost c_t to learn X_t;
- Recall: DM can stop at any time and pick best alternative explored thus far.

Motivation:

- Firm interviewing applicants;
- Consumer searching for a product.

Pandora

What changes?

Given an order of search, problem is *exactly* the same as before.

Main difficulty: deciding what to learn about next.

Simplifying assumption: independence of payoffs across alternatives
$\implies$ focus on *history-independent search orders* WLOO (why?).

Ordering Alternatives

Feasible orderings Π:

$$\Pi := \{\pi \in \mathbb{N}_0^{\mathbb{N}_0} \mid \pi \text{ is bijective and } \pi(t) = t \ \forall t \notin [T]\}.$$

π: permutation of elements in $[T] = \{1, ..., T\}$; defines an order of search.

π makes use only of time t.

$\pi(t)$: alternative DM searches/learns about if they haven't yet stopped by time t.

Ordering Alternatives

Fixing order π, problem as before: i.e. optimal stopping and choosing:

- $X_{\pi(t)} \equiv X_t^\pi$ as the (gross) payoff associated to alternative $\pi(t) = n$;
- $c_{\pi(t)}$ associated search cost;
- $M_t^\pi := \max_{s \leq t} X_s^\pi$: highest (gross) payoff thus far;
- $Y_t^\pi := M_t^\pi - \sum_{s \leq t} c_{\pi(t)}, t \leq T$; $\quad X_t^\pi = Y_t^\pi = -\infty$ for $t > T$; $X_0 = Y_0$ given;
- $\mathbb{T}^\pi$: stopping times taking values in $\mathbb{N}_0$ adapted to natural filtration given $\{X_t^\pi\}$.

Ordering Alternatives

Fixing order π, problem as before: i.e. optimal stopping and choosing:

- $X_{\pi(t)} \equiv X_t^\pi$ as the (gross) payoff associated to alternative $\pi(t) = n$;
- $c_{\pi(t)}$ associated search cost;
- $M_t^\pi := \max_{s \leq t} X_s^\pi$: highest (gross) payoff thus far;
- $Y_t^\pi := M_t^\pi - \sum_{s \leq t} c_{\pi(t)}, t \leq T$; $\quad X_t^\pi = Y_t^\pi = -\infty$ for $t > T$; $X_0 = Y_0$ given;
- $\mathbb{T}^\pi$: stopping times taking values in $\mathbb{N}_0$ adapted to natural filtration given $\{X_t^\pi\}$.

Pandora's Problem

$$\max_{\pi \in \Pi} \max_{\tau \in \mathbb{T}^\pi} \mathbb{E}[Y_\tau^\pi] \tag{W}$$

Optimal Stopping

Leveraging what we know:

$\forall \pi$, there is optimal stopping time; Π finite $\implies$ there is solution.

WLOO focus on earliest optimal stopping time for each order.

Solving Pandora's Problem

Start with simple stopping rules:

If problem were monotone, it'd suffice to consider 1-sla/threshold policy

$$\bar{x}_t := \inf\{x \in \mathbb{R} \mid x \geq \mathbb{E}[(X_t \vee x)] - c_t\}.$$

Whether problem monotone or not depends on order of search π.

Solving Pandora's Problem

Start with simple stopping rules:

- If problem were monotone, it'd suffice to consider 1-sla/threshold policy
$$\bar{x}_t := \inf\{x \in \mathbb{R} \mid x \geq \mathbb{E}[(X_t \vee x)] - c_t\}.$$
- Whether problem monotone or not depends on order of search π.
- Given π, define following stopping time:
$$\tau^\pi := \min\left\{t \geq 0 \mid M_t^\pi \geq \max_{s>t} \bar{x}_{\pi(s)}\right\}.$$

Solving Pandora's Problem

If DM stops at t it better be that

$$Y_t^\pi \geq \mathbb{E}[Y_{t+1}^\pi \mid \mathcal{F}_t^\pi] \iff M_t^\pi \geq \mathbb{E}[M_t^\pi \vee X_{\pi(t+1)} \mid \mathcal{F}_t^\pi] - c_{\pi(t+1)},$$

(i.e. not profitable to continue)

Solving Pandora's Problem

If DM stops at t it better be that

$$Y_t^\pi \geq \mathbb{E}[Y_{t+1}^\pi \mid \mathcal{F}_t^\pi] \Longleftrightarrow M_t^\pi \geq \mathbb{E}[M_t^\pi \vee X_{\pi(t+1)} \mid \mathcal{F}_t^\pi] - c_{\pi(t+1)},$$

(i.e. not profitable to continue)

More, since DM can *choose* continuation order, we must also have

$$M_t^\pi \geq \mathbb{E}[M_t^\pi \vee X_{\pi(t+h)} \mid \mathcal{F}_t^\pi] - c_{\pi(t+h)}, \forall h > 0,$$

Solving Pandora's Problem

If DM stops at t it better be that

$$Y_t^\pi \geq \mathbb{E}[Y_{t+1}^\pi \mid \mathcal{F}_t^\pi] \Longleftrightarrow M_t^\pi \geq \mathbb{E}[M_t^\pi \vee X_{\pi(t+1)} \mid \mathcal{F}_t^\pi] - c_{\pi(t+1)},$$

(i.e. not profitable to continue)

More, since DM can *choose* continuation order, we must also have

$$M_t^\pi \geq \mathbb{E}[M_t^\pi \vee X_{\pi(t+h)} \mid \mathcal{F}_t^\pi] - c_{\pi(t+h)}, \forall h > 0,$$

i.e. DM can choose to search not only $X_{\pi(t+1)}$ but any of remaining unsearched alternatives, $X_{\pi(t+h)}, h > 0$.
If they're optimally stopping, it must not be profitable to continue and try out any of the remaining alternatives.

Solving Pandora's Problem

Under optimal π and (earliest) optimal stopping time τ, on $\{\tau = t\}$ (stopping at t)

$$M_t^\pi - \sum_{s \le t} c_{\pi(s)} \ge \mathbb{E}[M_t^\pi \vee X_{t+h}^\pi] - \sum_{s \le t} c_{\pi(s)} - c_{\pi(t+h)}, \qquad \forall h > 0$$

$$\iff M_t^\pi \ge \mathbb{E}[M_t^\pi \vee X_{t+h}^\pi] - c_{\pi(t+h)}, \qquad \forall h > 0$$

$$\iff M_t^\pi \ge \bar{x}_{\pi(t+h)}, \qquad \forall h > 0$$

$$\iff M_t^\pi \ge \max_{s>t} \bar{x}_{\pi(s)}.$$

Solving Pandora's Problem

Under optimal π and (earliest) optimal stopping time τ, on $\{\tau = t\}$ (stopping at t)

$$M_t^\pi - \sum_{s \le t} c_{\pi(s)} \ge \mathbb{E}[M_t^\pi \vee X_{t+h}^\pi] - \sum_{s \le t} c_{\pi(s)} - c_{\pi(t+h)}, \qquad \forall h > 0$$

$$\iff M_t^\pi \ge \mathbb{E}[M_t^\pi \vee X_{t+h}^\pi] - c_{\pi(t+h)}, \qquad \forall h > 0$$

$$\iff M_t^\pi \ge \bar{x}_{\pi(t+h)}, \qquad \forall h > 0$$

$$\iff M_t^\pi \ge \max_{s>t} \bar{x}_{\pi(s)}.$$

Compare to $\tau^\pi := \min\left\{t \ge 0 \mid M_t^\pi \ge \max_{s>t} \bar{x}_{\pi(s)}\right\}$.

- Implies $\tau \ge \tau^\pi$.
- Natural conjecture to check: τ^π is optimal.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y_\tau^\pi] = \mathbb{E}[Y_{\tau^\pi}^\pi]$, τ^π is regular, and $\tau^\pi \leq \tau$.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y^{\pi}_{\tau}] = \mathbb{E}[Y^{\pi}_{\tau^{\pi}}]$, τ^{π} is regular, and $\tau^{\pi} \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y_\tau^\pi] = \mathbb{E}[Y_{\tau^\pi}^\pi]$, τ^π is regular, and $\tau^\pi \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \leq t\} \subseteq \{\tau^\pi \leq t\} \subseteq \{\tau^\pi \leq t + 1\}$.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y_\tau^\pi] = \mathbb{E}[Y_{\tau^\pi}^\pi]$, τ^π is regular, and $\tau^\pi \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \leq t\} \subseteq \{\tau^\pi \leq t\} \subseteq \{\tau^\pi \leq t + 1\}$.

WTS $\{\tau^\pi \leq t\} \subseteq \{\tau \leq t\} \implies \{\tau^\pi \leq t\} = \{\tau \leq t\}$.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y^\pi_\tau] = \mathbb{E}[Y^\pi_{\tau^\pi}]$, τ^π is regular, and $\tau^\pi \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \leq t\} \subseteq \{\tau^\pi \leq t\} \subseteq \{\tau^\pi \leq t+1\}$.

WTS $\{\tau^\pi \leq t\} \subseteq \{\tau \leq t\} \implies \{\tau^\pi \leq t\} = \{\tau \leq t\}$.

Step 1. $\tau^\pi \leq T-1 \implies Y^\pi_{T-1} \geq \mathbb{E}[Y^\pi_T] \implies \tau \leq T-1$; therefore $\{\tau^\pi \leq T-1\} = \{\tau \leq T-1\}$.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y^{\pi}_{\tau}] = \mathbb{E}[Y^{\pi}_{\tau^{\pi}}]$, τ^{π} is regular, and $\tau^{\pi} \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \leq t\} \subseteq \{\tau^{\pi} \leq t\} \subseteq \{\tau^{\pi} \leq t+1\}$.

WTS $\{\tau^{\pi} \leq t\} \subseteq \{\tau \leq t\} \implies \{\tau^{\pi} \leq t\} = \{\tau \leq t\}$.

Step 1. $\tau^{\pi} \leq T-1 \implies Y^{\pi}_{T-1} \geq \mathbb{E}[Y^{\pi}_{T}] \implies \tau \leq T-1$; therefore $\{\tau^{\pi} \leq T-1\} = \{\tau \leq T-1\}$.

Step 2. Induction: assume $s \geq t+1$, $\{\tau^{\pi} \leq s\} = \{\tau \leq s\} \implies \{\tau^{\pi} = s+1\} = \{\tau = s+1\}$.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y_\tau^\pi] = \mathbb{E}[Y_{\tau^\pi}^\pi]$, τ^π is regular, and $\tau^\pi \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \leq t\} \subseteq \{\tau^\pi \leq t\} \subseteq \{\tau^\pi \leq t+1\}$.

WTS $\{\tau^\pi \leq t\} \subseteq \{\tau \leq t\} \implies \{\tau^\pi \leq t\} = \{\tau \leq t\}$.

Step 1. $\tau^\pi \leq T-1 \implies Y_{T-1}^\pi \geq \mathbb{E}[Y_T^\pi] \implies \tau \leq T-1$; therefore $\{\tau^\pi \leq T-1\} = \{\tau \leq T-1\}$.

Step 2. Induction: assume $s \geq t+1$, $\{\tau^\pi \leq s\} = \{\tau \leq s\} \implies \{\tau^\pi = s+1\} = \{\tau = s+1\}$.

Since $\{t \geq \tau^\pi\} \cap \{\tau > t\} \subseteq \{t+1 \geq \tau^\pi\} \cap \{\tau > t\} = \{\tau = t+1\}$, on $\{t \geq \tau^\pi\} \cap \{\tau > t\}$,

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y^{\pi}_{\tau}] = \mathbb{E}[Y^{\pi}_{\tau^{\pi}}]$, τ^{π} is regular, and $\tau^{\pi} \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \leq t\} \subseteq \{\tau^{\pi} \leq t\} \subseteq \{\tau^{\pi} \leq t+1\}$.

WTS $\{\tau^{\pi} \leq t\} \subseteq \{\tau \leq t\} \implies \{\tau^{\pi} \leq t\} = \{\tau \leq t\}$.

Step 1. $\tau^{\pi} \leq T-1 \implies Y^{\pi}_{T-1} \geq \mathbb{E}[Y^{\pi}_{T}] \implies \tau \leq T-1$; therefore $\{\tau^{\pi} \leq T-1\} = \{\tau \leq T-1\}$.

Step 2. Induction: assume $s \geq t+1$, $\{\tau^{\pi} \leq s\} = \{\tau \leq s\} \implies \{\tau^{\pi} = s+1\} = \{\tau = s+1\}$.

Since $\{t \geq \tau^{\pi}\} \cap \{\tau > t\} \subseteq \{t+1 \geq \tau^{\pi}\} \cap \{\tau > t\} = \{\tau = t+1\}$, on $\{t \geq \tau^{\pi}\} \cap \{\tau > t\}$,

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y^\pi_\tau] = \mathbb{E}[Y^\pi_{\tau^\pi}]$, τ^π is regular, and $\tau^\pi \le \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \le t\} \subseteq \{\tau^\pi \le t\} \subseteq \{\tau^\pi \le t+1\}$.

WTS $\{\tau^\pi \le t\} \subseteq \{\tau \le t\} \implies \{\tau^\pi \le t\} = \{\tau \le t\}$.

Step 1. $\tau^\pi \le T-1 \implies Y^\pi_{T-1} \ge \mathbb{E}[Y^\pi_T] \implies \tau \le T-1$; therefore $\{\tau^\pi \le T-1\} = \{\tau \le T-1\}$.

Step 2. Induction: assume $s \ge t+1$, $\{\tau^\pi \le s\} = \{\tau \le s\} \implies \{\tau^\pi = s+1\} = \{\tau = s+1\}$.

Since $\{t \ge \tau^\pi\} \cap \{\tau > t\} \subseteq \{t+1 \ge \tau^\pi\} \cap \{\tau > t\} = \{\tau = t+1\}$, on $\{t \ge \tau^\pi\} \cap \{\tau > t\}$,
then $\mathbb{E}[Y^\pi_\tau \mid \tau > t \ge \tau^\pi] = \mathbb{E}[Y^\pi_{t+1} \mid \tau > t \ge \tau^\pi] \le Y^\pi_t \implies \tau \le t$.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y^\pi_\tau] = \mathbb{E}[Y^\pi_{\tau^\pi}]$, τ^π is regular, and $\tau^\pi \le \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \le t\} \subseteq \{\tau^\pi \le t\} \subseteq \{\tau^\pi \le t+1\}$.

WTS $\{\tau^\pi \le t\} \subseteq \{\tau \le t\} \implies \{\tau^\pi \le t\} = \{\tau \le t\}$.

Step 1. $\tau^\pi \le T-1 \implies Y^\pi_{T-1} \ge \mathbb{E}[Y^\pi_T] \implies \tau \le T-1$; therefore $\{\tau^\pi \le T-1\} = \{\tau \le T-1\}$.

Step 2. Induction: assume $s \ge t+1$, $\{\tau^\pi \le s\} = \{\tau \le s\} \implies \{\tau^\pi = s+1\} = \{\tau = s+1\}$.

Since $\{t \ge \tau^\pi\} \cap \{\tau > t\} \subseteq \{t+1 \ge \tau^\pi\} \cap \{\tau > t\} = \{\tau = t+1\}$, on $\{t \ge \tau^\pi\} \cap \{\tau > t\}$,
then $\mathbb{E}[Y^\pi_\tau \mid \tau > t \ge \tau^\pi] = \mathbb{E}[Y^\pi_{t+1} \mid \tau > t \ge \tau^\pi] \le Y^\pi_t \implies \tau \le t$.
Contradicts regularity of τ. Hence, $\{t \ge \tau^\pi\} \cap \{\tau > t\} = \emptyset$ and $\{\tau^\pi \le t\} = \{\tau \le t\}$.

Conclude: $\tau^\pi = \tau$. More: $\forall$ optimal stopping time τ', $\mathbb{E}[Y^\pi_{\tau'}] = \mathbb{E}[Y^\pi_\tau] = \mathbb{E}[Y^\pi_{\tau^\pi}]$ and $\tau' \ge \tau^\pi$.

Pandora's Optimal Stopping Time

Proposition

Suppose (π, τ) solve Pandora's problem. Then, $\mathbb{E}[Y^\pi_\tau] = \mathbb{E}[Y^\pi_{\tau^\pi}]$, τ^π is regular, and $\tau^\pi \leq \tau$.

Proof

WL take τ regular and earliest optimal stopping time.

By the above, $\{\tau \leq t\} \subseteq \{\tau^\pi \leq t\} \subseteq \{\tau^\pi \leq t+1\}$.

WTS $\{\tau^\pi \leq t\} \subseteq \{\tau \leq t\} \implies \{\tau^\pi \leq t\} = \{\tau \leq t\}$.

Step 1. $\tau^\pi \leq T-1 \implies Y^\pi_{T-1} \geq \mathbb{E}[Y^\pi_T] \implies \tau \leq T-1$; therefore $\{\tau^\pi \leq T-1\} = \{\tau \leq T-1\}$.

Step 2. Induction: assume $s \geq t+1$, $\{\tau^\pi \leq s\} = \{\tau \leq s\} \implies \{\tau^\pi = s+1\} = \{\tau = s+1\}$.

Since $\{t \geq \tau^\pi\} \cap \{\tau > t\} \subseteq \{t+1 \geq \tau^\pi\} \cap \{\tau > t\} = \{\tau = t+1\}$, on $\{t \geq \tau^\pi\} \cap \{\tau > t\}$,
then $\mathbb{E}[Y^\pi_\tau \mid \tau > t \geq \tau^\pi] = \mathbb{E}[Y^\pi_{t+1} \mid \tau > t \geq \tau^\pi] \leq Y^\pi_t \implies \tau \leq t$.
Contradicts regularity of τ. Hence, $\{t \geq \tau^\pi\} \cap \{\tau > t\} = \emptyset$ and $\{\tau^\pi \leq t\} = \{\tau \leq t\}$.

Conclude: $\tau^\pi = \tau$. More: $\forall$ optimal stopping time τ', $\mathbb{E}[Y^\pi_{\tau'}] = \mathbb{E}[Y^\pi_\tau] = \mathbb{E}[Y^\pi_{\tau^\pi}]$ and $\tau' \geq \tau^\pi$.

Regularity of τ^π follows from definition

Pandora's Optimal Search Order

We found out the optimal stopping time for the optimal order π

But we still don't know π

Pandora's Optimal Search Order

We found out the optimal stopping time for the optimal order π

But we still don't know π

If for all t, $\bar{x}_{\pi(t)} \geq \bar{x}_{\pi(t+1)}$ problem would be monotone, since

$$M_t^\pi \geq \max_{s>t} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+1)} \implies Y_t^\pi \geq \mathbb{E}[Y_{t+1}^\pi \mid \mathcal{F}_t^\pi]$$

Then τ^π would simply correspond to the 1-sla rule!

Pandora's Optimal Search Order

Claim: There's beauty and order in the universe. Should check for simple solutions.

Conjecture: if π were *not* inducing a decreasing sequence of $\bar{x}_{\pi(t)}$, rearranging π so that $\{\bar{x}_{\pi(t)}\}_t$ is decreasing sequence improves payoffs.

Pandora's Optimal Search Order

Claim: There's beauty and order in the universe. Should check for simple solutions.

Conjecture: if π were *not* inducing a decreasing sequence of $\overline{x}_{\pi(t)}$, rearranging π so that $\{\overline{x}_{\pi(t)}\}_t$ is decreasing sequence improves payoffs.

Turns out it's true:

Pandora's Optimal Search Order

Claim: There's beauty and order in the universe. Should check for simple solutions.

Conjecture: if π were *not* inducing a decreasing sequence of $\overline{x}_{\pi(t)}$, rearranging π so that $\{\overline{x}_{\pi(t)}\}_t$ is decreasing sequence improves payoffs.

Turns out it's true:

Pandora's Optimal Search Order

Claim: There's beauty and order in the universe. Should check for simple solutions.

Conjecture: if π were *not* inducing a decreasing sequence of $\bar{x}_{\pi(t)}$, rearranging π so that $\{\bar{x}_{\pi(t)}\}_t$ is decreasing sequence improves payoffs.

Turns out it's true:

Proposition

If (π, τ^{π}) solve Pandora's problem, then $\{\bar{x}_{\pi(s)}\}_{s \leq t}$ is nonincreasing for all $t : \mathbb{P}(\tau^{\pi} \geq t) > 0$

Pandora's Optimal Search Order

Proof

Suppose $\bar{x}_{\pi(s)}$ nonincreasing for $s \geq t + 1$, but $\bar{x}_{\pi(t)} < \bar{x}_{\pi(t+1)}$.

Pandora's Optimal Search Order

Proof

Suppose $\bar{x}_{\pi(s)}$ nonincreasing for $s \geq t+1$, but $\bar{x}_{\pi(t)} < \bar{x}_{\pi(t+1)}$.

Assume that $\tau^{\pi} \geq t$ occurs with positive prob., as ow it is WL to rearrange the order following t (no impact on payoffs).

Pandora's Optimal Search Order

Proof

Suppose $\overline{x}_{\pi(s)}$ nonincreasing for $s \geq t + 1$, but $\overline{x}_{\pi(t)} < \overline{x}_{\pi(t+1)}$.

Assume that $\tau^{\pi} \geq t$ occurs with positive prob., as ow it is WL to rearrange the order following t (no impact on payoffs).

Define new order δ: same as π except it swaps t-th and $(t+1)$-th alternatives.

Pandora's Optimal Search Order

Proof

Suppose $\overline{x}_{\pi(s)}$ nonincreasing for $s \geq t+1$, but $\overline{x}_{\pi(t)} < \overline{x}_{\pi(t+1)}$.

Assume that $\tau^{\pi} \geq t$ occurs with positive prob., as ow it is WL to rearrange the order following t (no impact on payoffs).

Define new order δ: same as π except it swaps t-th and $(t+1)$-th alternatives.

Define stopping time $\tau := \min\{s \geq 0 \mid M_s^{\delta} \geq \max_{n>s} \overline{x}_{\pi(s)}\}$

Pandora's Optimal Search Order

Proof

Suppose $\overline{x}_{\pi(s)}$ nonincreasing for $s \geq t+1$, but $\overline{x}_{\pi(t)} < \overline{x}_{\pi(t+1)}$.

Assume that $\tau^{\pi} \geq t$ occurs with positive prob., as ow it is WL to rearrange the order following t (no impact on payoffs).

Define new order δ: same as π except it swaps t-th and $(t+1)$-th alternatives.

Define stopping time $\tau := \min\{s \geq 0 \mid M_s^{\delta} \geq \max_{n>s} \overline{x}_{\pi(s)}\}$

Then $\{\tau \leq s\} = \{\tau^{\pi} \leq s\}, \forall s \neq t$; i.e. only difference between τ^{π} and τ is if they disagree in stopping *at* t.

Pandora's Optimal Search Order

Proof

Suppose $\overline{x}_{\pi(s)}$ nonincreasing for $s \geq t+1$, but $\overline{x}_{\pi(t)} < \overline{x}_{\pi(t+1)}$.

Assume that $\tau^{\pi} \geq t$ occurs with positive prob., as ow it is WL to rearrange the order following t (no impact on payoffs).

Define new order δ: same as π except it swaps t-th and $(t+1)$-th alternatives.

Define stopping time $\tau := \min\{s \geq 0 \mid M_s^{\delta} \geq \max_{n>s} \overline{x}_{\pi(s)}\}$

Then $\{\tau \leq s\} = \{\tau^{\pi} \leq s\}, \forall s \neq t$; i.e. only difference between τ^{π} and τ is if they disagree in stopping *at* t.

A long derivation (here) reveals that $\mathbb{E}[Y_{\tau^{\pi}}^{\pi} - Y_{\tau}^{\delta}] < 0$, contradicting optimality of π (*not* contradicting optimality of τ^{π} — we've seen that if π were an optimal order, τ^{π} is an optimal stopping time).

Pandora's Optimal Search Order

Proof

Suppose $\bar{x}_{\pi(s)}$ nonincreasing for $s \geq t+1$, but $\bar{x}_{\pi(t)} < \bar{x}_{\pi(t+1)}$.

Assume that $\tau^\pi \geq t$ occurs with positive prob., as ow it is WL to rearrange the order following t (no impact on payoffs).

Define new order δ: same as π except it swaps t-th and $(t+1)$-th alternatives.

Define stopping time $\tau := \min\{s \geq 0 \mid M_s^\delta \geq \max_{n>s} \bar{x}_{\pi(s)}\}$

Then $\{\tau \leq s\} = \{\tau^\pi \leq s\}, \forall s \neq t$; i.e. only difference between τ^π and τ is if they disagree in stopping *at* t.

A long derivation (here) reveals that $\mathbb{E}[Y^\pi_{\tau^\pi} - Y^\delta_\tau] < 0$, contradicting optimality of π (*not* contradicting optimality of τ^π — we've seen that if π were an optimal order, τ^π is an optimal stopping time).

Iterating argument optimal order π sat. $\{\bar{x}_{\pi(s)}\}_{s \leq t}$ nonincreasing.

Pandora's Optimal Search Order

Proposition

If (π, τ^{π}) solve Pandora's problem, then $\{\bar{x}_{\pi(s)}\}_{s \leq t}$ is nonincreasing for all $t : \mathbb{P}(\tau^{\pi} \geq t) > 0$

We find out that the solution to Pandora's problem is s.t. (1) optimal order is simply to order alternatives according to $\bar{x}_{\pi(t)}$; (2) use 1-sla stopping rule!

Variations to Pandora's Problem

Pandora's rule doesn't generalise easily.

Variations to Pandora's Problem

Pandora's rule doesn't generalise easily.

Breaking Independence

- All hell breaks loose: now optimal order may depend on history of observed payoffs.
- Learning X_n changes beliefs about X_m!

Variations to Pandora's Problem

Pandora's rule doesn't generalise easily.

Breaking Independence

All hell breaks loose: now optimal order may depend on history of observed payoffs.

Learning X_n changes beliefs about X_m!

Choosing without Search

Things go a bit awry, but Doval (2018 JET) provides sufficient conditions for Pandora's rule to remain optimal except for last alternative.

Also shows how for simple setting (binary outcomes) breaking independence leads to spectacular failure of Pandora's rule to be optimal.

Variations to Pandora's Problem

Pandora's rule doesn't generalise easily.

Payoffs depend on values of all $\{X_n\}$

- E.g. alternatives as different components of research project.
  - DM decides if to work on them.
  - Total payoff depends on what is ultimately included in paper.
- Olszewski & Weber (2015 JET): sufficient conditions for Pandora's rule to remain optimal.

Variations to Pandora's Problem

Pandora's rule doesn't generalise easily.

Payoffs depend on values of all $\{X_n\}$

- E.g. alternatives as different components of research project.
  DM decides if to work on them.
  Total payoff depends on what is ultimately included in paper.
- Olszewski & Weber (2015 JET): sufficient conditions for Pandora's rule to remain optimal.

Flexible Learning

- I haven't checked thoroughly, but I think an open problem is to retain independence and characterise the optimal solution when allowing for more flexible learning.
  Hope for something nice/tractable that would do well in applications.
  E.g., when assessing candidates, scan some CVs, then dig into some good ones and interview them; if turn out not great, go back to pile.

Application: R&D and Project Selection

T different research projects

Project n:

- Ex-ante prob succsess $p_n \in (0, 1)$ and payoff $r_n > 0$; ow payoff 0.
- Associated cost c_n.
- DM only considers projects with positive exp. value: $p_n r_n - c_n > 0$.
- Outside option yields 0.

Application: R&D and Project Selection

T different research projects

Project n:

- Ex-ante prob succsess $p_n \in (0, 1)$ and payoff $r_n > 0$; ow payoff 0.
- Associated cost c_n.
- DM only considers projects with positive exp. value: $p_n r_n - c_n > 0$.
- Outside option yields 0.

Solution

$$\bar{x}_n := \inf\{x \in \mathbb{R} \mid x \geq \mathbb{E}[r_n \vee x] - c_n\} = r_n - \frac{c_n}{p_n},$$

DM trades off expected net reward and prob success:

- Explore first projects with high potential reward net of cost scaled by prob success.
- Stop myopically (1-sla).

Overview

Martingales and etc.

We'll do a quick detour to recall some fundamental results about martingales.

Martingales and etc.

We'll do a quick detour to recall some fundamental results about martingales.

Martingales and etc.

Let $X = \{X_t\}_{t\in[T]}$ be an adapted process.

- X is a **supermatingale** if $\mathbb{E}[X_t^-] < \infty \; \forall t \in T$ and $\mathbb{E}[X_t \mid \mathcal{F}_s] \leq X_s$ a.s. for all $s \leq t$, $s, t \in [T]$.
- X is a **submartingale** if $-X$ is a supermatingale.
- X is a **martingale** if it is both a super- and submartingale. Specifically, a martingale satisfies $\mathbb{E}[X_t \mid \mathcal{F}_s] = X_s$ for $s \leq t$.
- X is a **predictable process** if X_t is $\mathcal{F}_{t-1}$-measurable.

Doob's Decomposition Theorem

Doob's Decomposition Theorem

If X is an $\mathbb{F}$-adapted process satisfying $\mathbb{E}[\sup_{s\geq t}|X_s|] < \infty$, then $\exists$ martingale $Z : Z_0 = \mathbf{0}$ and an integrable predictable process $A : A_0 = \mathbf{0}$ s.t. $X_t = X_0 + Z_t + A_t \quad \forall t$, with decomposition being unique a.s.

Implication: super/martingale = martingale + a.s. de/increasing predictable process.

Doob's Martingale Convergence Theorem

Doob's Martingale Convergence Theorem

If X is supermartingale s.t. $\sup_{t>0} \mathbb{E}[X_t^-] < \infty$, then X_t converges pointwise to X_∞, $X_t(\omega) \to X_\infty(\omega)$, and $X_\infty < \infty$ a.s.
Moreover, if X is martingale, it is uniformly integrable if and only if X_t converges a.s. and in L^1 to X_∞ satisfying $X_t = \mathbb{E}[X_\infty \mid \mathcal{F}_t]$ for all t.

Doob's Martingale Convergence Theorem

Doob's Martingale Convergence Theorem

If X is supermartingale s.t. $\sup_{t>0} \mathbb{E}[X_t^-] < \infty$, then X_t converges pointwise to X_∞, $X_t(\omega) \to X_\infty(\omega)$, and $X_\infty < \infty$ a.s.
Moreover, if X is martingale, it is uniformly integrable if and only if X_t converges a.s. and in L^1 to X_∞ satisfying $X_t = \mathbb{E}[X_\infty \mid \mathcal{F}_t]$ for all t.

Recall, (i) X is uniformly integrable if $\lim_{a\to\infty} \sup_t \mathbb{E}[|X_t| \mathbf{1}_{\{|X_t|>a\}}] = 0$; (ii)

$$X_t \xrightarrow{L^1} X_\infty \iff \mathbb{E}[|X_t - X_\infty|] \to 0.$$

Doob's Martingale Convergence Theorem

Doob's Martingale Convergence Theorem

If X is supermartingale s.t. $\sup_{t>0} \mathbb{E}[X_t^-] < \infty$, then X_t converges pointwise to X_∞, $X_t(\omega) \to X_\infty(\omega)$, and $X_\infty < \infty$ a.s.
Moreover, if X is martingale, it is uniformly integrable if and only if X_t converges a.s. and in L^1 to X_∞ satisfying $X_t = \mathbb{E}[X_\infty \mid \mathcal{F}_t]$ for all t.

Recall, (i) X is uniformly integrable if $\lim_{a\to\infty} \sup_t \mathbb{E}[|X_t| \mathbf{1}_{\{|X_t|>a\}}] = 0$; (ii)

$$X_t \xrightarrow{L^1} X_\infty \iff \mathbb{E}[|X_t - X_\infty|] \to 0.$$

Doob's Optional Stopping Theorem

Let X be supermartingale (resp. submartingale) and τ a stopping time. Suppose one of the following holds:

(i) $\exists c < \infty : \tau \leq c$ a.s.

(ii) $\mathbb{E}[\tau] < \infty$ and $\exists c < \infty : \forall t,\ \mathbb{E}[|X_{t+1} - X_t| \mid \mathcal{F}_t] \leq c$ on $\{\tau > t\}$.

(iii) $\exists c < \infty : \forall t,\ |X_{t\wedge\tau}| \leq c$ a.s.

Then, X_τ is a.s. well-defined r.v. and $\mathbb{E}[X_\tau] \leq$ (resp. $\geq$)$\mathbb{E}[X_0]$.

Doob's Martingale Convergence Theorem

Counterexample

Wealth after sequence of iid fair bets, $S_t = \sum_{s \leq t} X_s$, where $X_s = \pm\mathbf{1}$ wp 1/2 and $X_0 = \mathbf{0}$.
$\tau := \inf\{t \geq \mathbf{0} \mid S_t = \mathbf{1}\}$. Then S_τ is martingale.
$\mathbb{E}[\tau] = \infty$ and optional stopping theorem doesn't apply: $\mathbb{E}[S_\tau] = \mathbf{1} > \mathbf{0} = \mathbb{E}[S_t]$.
S_t doesn't converge in mean. Also, $\{S_t\}$ not u.i.

Wald's Equations

Wald's Equation

Let X be s.t. (i) $\mathbb{E}[X_t] < \infty$, (ii) $\forall t \ \mathbb{E}[X_t \mathbf{1}\{\tau \geq t\}] = \mathbb{E}[X_t]\mathbb{P}(\tau \geq t)$, and (iii) $\sum_{t \in \mathbb{N}} \mathbb{E}[|X_t| \mathbf{1}\{\tau \geq t\}] < \infty$. Define $S_\tau := \sum_{t=1}^{\tau} X_t$ and $T_\tau := \sum_{t=1}^{\tau} \mathbb{E}[X_t]$.
Then $\mathbb{E}[S_\tau] = \mathbb{E}[T_\tau] < \infty$. If, moreover, $\mathbb{E}[X_t] = m \ \forall t$ and $\mathbb{E}[\tau] < \infty$, then $\mathbb{E}[S_\tau] = \mathbb{E}[\tau]m$.

Wald's Equations

There are a number of versions of this result typically labeled Wald's First/ Second/ Third Identity which follow from the optional stopping theorem:

Wald's First/Second/Third Identities

Let $\{X_t\}_{t\in\mathbb{N}}$ be a stochastic process such that X_t are independent, with $\mathbb{F}$ denoting its natural filtration, and τ be an adapted stopping time with $\mathbb{E}[\tau] < \infty$. Define (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 $M_t^3 \mathbb{1}\{\tau \geq t\} \leq \delta < \infty$ for all t, then M_t^3 is a martingale and $\mathbb{E}[\phi(\theta)^{-\tau}\exp(\theta S_\tau)] = 1$.

Overview

1. Pandora's Problem
2. Martingales and etc.
3. Gittins Index
4. Web Search
5. Pricing with Pandora Consumers
6. Summary

Multi-Armed Bandit Problems

Setup:

Each period $t = 0, 1, ...,$

DM ('gambler') chooses one action ('arm') $a_t \in A := \{1, ..., K\}$

receives a random payoff $x_t^{a_t}$

whose distribution $F^{a_t}(\cdot; s_t^k)$ depends on the state $s_t = (s_t^1, ..., s_t^K)$,

and the state evolves according to a Markov chain s.t. $s_{t+1}^k = \phi_k(x_t^k, s_t^k)$ if $a_t = k$ and $s_{t+1}^k = s_t^k$ if $a_t \neq k$. ('restless' bandit when allowing $s_{t+1}^k \neq s_t^k$ even if $a_t \neq k$)

Payoffs are discounted by $\delta \in [0, 1)$.

Multi-Armed Bandit Problems

Setup:

Each period $t = 0, 1, \ldots,$

DM ('gambler') chooses one action ('arm') $a_t \in A := \{1, \ldots, K\}$

receives a random payoff $x_t^{a_t}$

whose distribution $F^{a_t}(\cdot; s_t^k)$ depends on the state $s_t = (s_t^1, \ldots, s_t^K)$,

and the state evolves according to a Markov chain s.t. $s_{t+1}^k = \phi_k(x_t^k, s_t^k)$ if $a_t = k$ and $s_{t+1}^k = s_t^k$ if $a_t \neq k$. ('restless' bandit when allowing $s_{t+1}^k \neq s_t^k$ even if $a_t \neq k$)

Payoffs are discounted by $\delta \in [0, 1)$.

Motivation: learning-by-doing + choosing what to do.

Choosing and switching streaming platform subscription.

Learning about job match value and decide to switch to another job.

Learning about effects of enacted policy.

Joining a queue and switching to another another.

Draw-down retirement savings by depleting different assets.

Seller learning demand by experimenting with prices.

Scheduling of experiments (e.g., Pandora's boxes!)

Multi-Armed Bandit Problems

Setup:

Each period $t = 0, 1, ...,$

DM ('gambler') chooses one action ('arm') $a_t \in A := \{1, ..., K\}$

receives a random payoff $x_t^{a_t}$

whose distribution $F^{a_t}(\cdot; s_t^k)$ depends on the state $s_t = (s_t^1, ..., s_t^K)$,

and the state evolves according to a Markov chain s.t. $s_{t+1}^k = \phi_k(x_t^k, s_t^k)$ if $a_t = k$ and $s_{t+1}^k = s_t^k$ if $a_t \neq k$. ('restless' bandit when allowing $s_{t+1}^k \neq s_t^k$ even if $a_t \neq k$)

Payoffs are discounted by $\delta \in [0, 1)$.

Example of Learning Framing:

Unknown parameters: $\theta^k \in \Theta^k$, $\theta^k \sim \mu_0^k \in \Delta(\Theta^k)$. $\theta = (\theta^1, ..., \theta^k)$ and $\mu_0 = \times_k \mu_0^k$ (independence; product measure).

Objective payoff distributions: $X_t^k \sim G^k(\cdot \mid \theta^k)$ iid.

Posterior beliefs: $\mu_{t+1}^k = \mu_t^k \mid X_t^k$ if $a_t = k$ and $\mu_{t+1}^k = \mu_t^k$ if $a_t \neq k$ (our Markovian 'state').

Subjective payoff distributions: $F^k(\cdot \mid \mu_t^k) := \mathbb{E}_{\mu_t^k}[G^k(\cdot \mid \theta^k)]$.

Actions entail payoffs *and* learning about payoff distribution.

Multi-Armed Bandit Problems

Setup:

Histories: $h^t := (s_0, s_1, ..., s_t) \in H_t$ and $\mathcal{H} := \cup_{t=0}^{\infty} H_t$.

Strategies: $\alpha \in A^{\mathcal{H}}$.

Payoffs $\mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Goal: $V(s_0) := \sup_{\alpha} \mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Multi-Armed Bandit Problems

Setup:

Histories: $h^t := (s_0, s_1, ..., s_t) \in H_t$ and $\mathcal{H} := \cup_{t=0}^{\infty} H_t$.

Strategies: $\alpha \in A^{\mathcal{H}}$.

Payoffs $\mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Goal: $V(s_0) := \sup_\alpha \mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Denote $X^k(s_t^k) \sim F^k(\cdot; s_t^k)$.

Theorem (Gittins & Jones, 1974)

The optimal policy satisfies $\alpha(s_t) \in \arg\max_{a \in A} m^a(s_t^a)$, where

$$(1-\delta) m^k(s_0^k) := \sup_\tau \frac{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k X^k(s_t^k)]}{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k]}.$$

Multi-Armed Bandit Problems

Setup:

Histories: $h^t := (s_0, s_1, ..., s_t) \in H_t$ and $\mathcal{H} := \cup_{t=0}^{\infty} H_t$.

Strategies: $\alpha \in A^{\mathcal{H}}$.

Payoffs $\mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Goal: $V(s_0) := \sup_\alpha \mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Denote $X^k(s_t^k) \sim F^k(\cdot; s_t^k)$.

Theorem (Gittins & Jones, 1974)

The optimal policy satisfies $\alpha(s_t) \in \arg\max_{a \in A} m^a(s_t^a)$, where

$$(1 - \delta) m^k(s_0^k) := \sup_\tau \frac{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k X^k(s_t^k)]}{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k]}.$$

Comments:

Decomposes K-dimensional problem to solve K 1-dimensional problems.

Formulation very general: with transition can capture use-costs of arm, countable number of arms, etc.

Multi-Armed Bandit Problems

Setup:

Histories: $h^t := (s_0, s_1, ..., s_t) \in H_t$ and $\mathcal{H} := \cup_{t=0}^{\infty} H_t$.

Strategies: $\alpha \in A^{\mathcal{H}}$.

Payoffs $\mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Goal: $V(s_0) := \sup_{\alpha} \mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Multi-Armed Bandit Problems

Setup:

Histories: $h^t := (s_0, s_1, ..., s_t) \in H_t$ and $\mathcal{H} := \cup_{t=0}^{\infty} H_t$.

Strategies: $\alpha \in A^{\mathcal{H}}$.

Payoffs $\mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Goal: $V(s_0) := \sup_{\alpha} \mathbb{E}[\sum_{t=0}^{\infty} \delta^t X_t^{\alpha_t}]$.

Denote $X^k(s_t^k) \sim F^k(\cdot; s_t^k)$.

Theorem (Gittins & Jones, 1974)

The optimal policy satisfies $\alpha(s_t) \in \arg\max_{a \in A} m^a(s_t^a)$, where

$$(1 - \delta) m^k(s_0^k) := \sup_{\tau} \frac{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k X^k(s_t^k)]}{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k]}.$$

Seminal papers: Gittins & Jones (1974, 1979), Gittins (1979), Weber (1992).
Also Gittins (1989), Karatzas (1984; Brownian bandits), Banks & Sundaram (1994; switching costs).

Applications: Pricing and learning demand: Rothschild (1974); McLennan (1984); Rustichini & Wolinsky (1995); Keller & Rady (1999); Bergemann & Välimäki (2006),

Multi-Armed Bandit Problems

Where does the Gittins-Jones index come from?

Multi-Armed Bandit Problems

Where does the Gittins-Jones index come from?

Consider two actions: pull arm k or get lump-sum reward M.

Lump-sum reward optimal $\implies s_{t+1}^k = s_t^k \implies$ Lump-sum reward remains optimal.

Multi-Armed Bandit Problems

Where does the Gittins-Jones index come from?

Consider two actions: pull arm k or get lump-sum reward M.

Lump-sum reward optimal $\implies s_{t+1}^k = s_t^k \implies$ Lump-sum reward remains optimal.

Effectively optimal stopping problem: when to stop pulling arm k.

$$V(s_0^k, M) := \sup_{\tau} \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau M\right] \qquad \text{(SP)}$$

$$= \max\{M, X^k(s_0^k) + \delta \mathbb{E}[V(s_1^k, M) \mid s_0^k]\} \qquad \text{(DP)}$$

Where does the Gittins-Jones index come from?

Consider two actions: pull arm k or get lump-sum reward M.

Lump-sum reward optimal $\implies s^k_{t+1} = s^k_t \implies$ Lump-sum reward remains optimal.

Effectively optimal stopping problem: when to stop pulling arm k.

$$V(s^k_0, M) := \sup_{\tau} \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s^k_t) + \delta^\tau M\right] \qquad \text{(SP)}$$

$$= \max\{M, X^k(s^k_0) + \delta \mathbb{E}[V(s^k_1, M) \mid s^k_0]\} \qquad \text{(DP)}$$

Note: (i) V increasing in M,

(ii) convex in M (maximising over linear functions of M), and

(iii) $\exists M' : V(s^k_0, M) = M \iff M \geq M'$ (both stopping and continuation region are intervals).

Where does the Gittins-Jones index come from?

$$V(s_0^k, M) := \sup_{\tau} \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau M\right] = \max\{M, X^k(s_0^k) + \delta \mathbb{E}[V(s_1^k, M) \mid s_0^k]\} \qquad \text{(DP)}$$

$m^k(s^k) := \inf\{M \mid V(s^k, M) = M\}$ is smallest lump-sum prize that DM will take to stop pulling arm k in state s^k.

Where does the Gittins-Jones index come from?

$$V(s_0^k, M) := \sup_{\tau} \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau M\right] = \max\{M, X^k(s_0^k) + \delta \mathbb{E}[V(s_1^k, M) \mid s_0^k]\} \qquad \text{(DP)}$$

$m^k(s^k) := \inf\{M \mid V(s^k, M) = M\}$ is smallest lump-sum prize that DM will take to stop pulling arm k in state s^k.

Experimentation suboptimal if $M = m^k(s^k)$. Hence, $\forall \tau$,

$$m^k(s_0^k) \geq \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau m^k(s_0^k)\right] \iff m^k(s_0^k)(1 - \mathbb{E}[\delta^\tau]) \geq \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k)\right].$$

Where does the Gittins-Jones index come from?

$$V(s_0^k, M) := \sup_{\tau} \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau M\right] = \max\{M, X^k(s_0^k) + \delta \mathbb{E}[V(s_1^k, M) \mid s_0^k]\} \qquad \text{(DP)}$$

$m^k(s^k) := \inf\{M \mid V(s^k, M) = M\}$ is smallest lump-sum prize that DM will take to stop pulling arm k in state s^k.

Experimentation suboptimal if $M = m^k(s^k)$. Hence, $\forall \tau$,

$$m^k(s_0^k) \geq \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau m^k(s_0^k)\right] \iff m^k(s_0^k)(1 - \mathbb{E}[\delta^\tau]) \geq \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k)\right].$$

Since $\sum_{s=0}^{t-1} \delta^s = \frac{1-\delta^t}{1-\delta}$, $m^k(s_0^k)(1 - \mathbb{E}[\delta^\tau]) = (1-\delta) m^k(s_0^k) \mathbb{E}[\sum_{t=0}^{\tau-1} \delta^t]$. Then,

Multi-Armed Bandit Problems

Where does the Gittins-Jones index come from?

$$V(s_0^k, M) := \sup_\tau \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau M\right] = \max\{M, X^k(s_0^k) + \delta \mathbb{E}[V(s_1^k, M) \mid s_0^k]\} \qquad \text{(DP)}$$

$m^k(s^k) := \inf\{M \mid V(s^k, M) = M\}$ is smallest lump-sum prize that DM will take to stop pulling arm k in state s^k.

Experimentation suboptimal if $M = m^k(s^k)$. Hence, $\forall \tau$,

$$m^k(s_0^k) \geq \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k) + \delta^\tau m^k(s_0^k)\right] \iff m^k(s_0^k)(1 - \mathbb{E}[\delta^\tau]) \geq \mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k)\right].$$

Since $\sum_{s=0}^{t-1} \delta^s = \frac{1-\delta^t}{1-\delta}$, $m^k(s_0^k)(1 - \mathbb{E}[\delta^\tau]) = (1-\delta) m^k(s_0^k) \mathbb{E}[\sum_{t=0}^{\tau-1} \delta^t]$. Then,

$$(1-\delta) m^k(s_0^k) \geq \frac{\mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k)\right]}{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^t]}, \qquad \forall \tau,$$

with equality at optimal stopping:

$$(1-\delta) m^k(s_0^k) = \sup_\tau \frac{\mathbb{E}\left[\sum_{t=0}^{\tau-1} \delta^t X^k(s_t^k)\right]}{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^t]}.$$

Multi-Armed Bandit Problems

Why does the Gittins-Jones index solve the optimisation problem?

Intuition from Weber (1992):

- Suppose seller rents arms out to competitive market of operators, all risk neutral with common discount factor.

Multi-Armed Bandit Problems

Why does the Gittins-Jones index solve the optimisation problem?

Intuition from Weber (1992):

- Suppose seller rents arms out to competitive market of operators, all risk neutral with common discount factor.
- Optimal to obtain high rental incomes in early periods.

Multi-Armed Bandit Problems

Why does the Gittins-Jones index solve the optimisation problem?

Intuition from Weber (1992):

- Suppose seller rents arms out to competitive market of operators, all risk neutral with common discount factor.
- Optimal to obtain high rental incomes in early periods.
- Rental market operated as descending price auction: fee for operating an arbitrary arm is lowered until an operator accepts price.
- At accepted price, operator can operate arm while they want.

Multi-Armed Bandit Problems

Why does the Gittins-Jones index solve the optimisation problem?

Intuition from Weber (1992):

- Suppose seller rents arms out to competitive market of operators, all risk neutral with common discount factor.
- Optimal to obtain high rental incomes in early periods.
- Rental market operated as descending price auction: fee for operating an arbitrary arm is lowered until an operator accepts price.
- At accepted price, operator can operate arm while they want.
- Since market for operators is competitive, price is s.t. under optimal stopping rule, operator breaks even. Hence, highest acceptable price for arm k is its Gittins-Jones index $m^k(s_t^k)$.

Why does the Gittins-Jones index solve the optimisation problem?

Intuition from Weber (1992):

- Suppose seller rents arms out to competitive market of operators, all risk neutral with common discount factor.
- Optimal to obtain high rental incomes in early periods.
- Rental market operated as descending price auction: fee for operating an arbitrary arm is lowered until an operator accepts price.
- At accepted price, operator can operate arm while they want.
- Since market for operators is competitive, price is s.t. under optimal stopping rule, operator breaks even. Hence, highest acceptable price for arm k is its Gittins-Jones index $m^k(s_t^k)$.
- Operator runs arm until Gittins-Jones index falls below price, i.e. its original Gittins-Jones Index.
- Once arm is abandoned, restart process of lowering the price offer.

Multi-Armed Bandit Problems

Why does the Gittins-Jones index solve the optimisation problem?

Intuition from Weber (1992):

- Suppose seller rents arms out to competitive market of operators, all risk neutral with common discount factor.
- Optimal to obtain high rental incomes in early periods.
- Rental market operated as descending price auction: fee for operating an arbitrary arm is lowered until an operator accepts price.
- At accepted price, operator can operate arm while they want.
- Since market for operators is competitive, price is s.t. under optimal stopping rule, operator breaks even. Hence, highest acceptable price for arm k is its Gittins-Jones index $m^k(s_t^k)$.
- Operator runs arm until Gittins-Jones index falls below price, i.e. its original Gittins-Jones Index.
- Once arm is abandoned, restart process of lowering the price offer.
- Since operators get zero surplus and they are operating under optimal rules, this method of allocating arms results in maximal surplus to owner, solving the original MAB problem.

Multi-Armed Bandit Problems

Theorem (Gittins & Jones, 1974)

The optimal policy satisfies $\alpha(s_t) \in \arg\max_{a \in A} m^a(s_t^a)$, where

$$(1-\delta)m^k(s_0^k) := \sup_{\tau} \frac{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k X^k(s_t^k)]}{\mathbb{E}[\sum_{t=0}^{\tau-1} \delta^k]}.$$

Very nice result. Hard to compute index in closed-form other than in special cases.

Fragility: result breaks with minor variations.

- non-geometric discounting, e.g., fixed time horizon;
- arms with correlated priors;
- actions affecting more than one arm at a time;
- payoffs depending on state of 2 or more arms;
- delayed feedback;
- ‘restless’ arms that change state without being pulled;
- switching costs; etc.

Overview

1. Pandora's Problem
2. Martingales and etc.
3. Gittins Index
4. Web Search
5. Pricing with Pandora Consumers
6. Summary

Choi and Smith WP

Overview

1. Pandora's Problem
2. Martingales and etc.
3. Gittins Index
4. Web Search
5. Pricing with Pandora Consumers
   - Setup
   - Consumer Behaviour
   - Market Equilibrium
   - Comparative Statics
6. Summary

Pricing with Pandora Consumers

Pricing with search (Choi, Dai, & Kim, 2018 Ecta)

Sellers

- N sellers; each supplies product n at price $p_n \geq 0$; marginal cost c_n.
- p: vector of prices; p_{-n}: prices of n's competitors.
- Demand for n given prices: $D_n(p)$.
- Sellers maximise profit: $\pi_n(p) := D_n(p)(p_n - c_n)$.

Pricing with Pandora Consumers

Consumers

- Consumer i's valuation of n: $X_n = V_n + W_n - p_n$.
  - V_n is known; W_n is idiosyncratic-value component revealed only when consumer learns about n.
- Cost to acquire info and learn W_n: $k_n > 0$.
- Outside option X_0; search with recall.
- V_n and W_n are independently drawn from F_n and G_n, both smooth.
- Independence allows use of Pandora's rule.
- If consumer stops after searching sellers $A \subseteq [N]$, the consumer accrues a payoff $\max_{n \in A \cup \{0\}} X_{i,n} - \sum_{m \in A} k_m$, where $p_0 = k_0 = 0$.
- Consumers know p.

Pandora's Rule

Define $\overline{w}_n$: $k_n = \mathbb{E}[(W_n - \overline{w}_n)^+]$

Given V_n define $\overline{x}_n := V_n + \overline{w}_n - p_n$

Given order $\pi \in \Pi$, let $M_t^\pi := \max_{s \leq t} X_{\pi(t)} = \max_{s \leq t} V_{\pi(t)} + W_{\pi(t)} - p_{\pi(t)}$

Consumer Behaviour

Pandora's Rule

Define $\overline{w}_n$: $k_n = \mathbb{E}[(W_n - \overline{w}_n)^+]$

Given V_n define $\overline{x}_n := V_n + \overline{w}_n - p_n$

Given order $\pi \in \Pi$, let $M_t^\pi := \max_{s \leq t} X_{\pi(t)} = \max_{s \leq t} V_{\pi(t)} + W_{\pi(t)} - p_{\pi(t)}$

Pandora's Rule

Define $\overline{w}_n$: $k_n = \mathbb{E}[(W_n - \overline{w}_n)^+]$

Given V_n define $\bar{x}_n := V_n + \overline{w}_n - p_n$

Given order $\boldsymbol{\pi} \in \boldsymbol{\Pi}$, let $M_t^{\pi} := \max_{s \leq t} X_{\pi(t)} = \max_{s \leq t} V_{\pi(t)} + W_{\pi(t)} - p_{\pi(t)}$

Proposition

Given a price vector p and realization V, it is optimal for the consumer
(1) to learn about sellers in decreasing order of $\bar{x}_n$, with the optimal order of search being given by $\boldsymbol{\pi} \in \boldsymbol{\Pi}$ such that $\bar{x}_{\pi}(t) \geq \bar{x}_{\pi}(t+1)$; and
(2) to stop whenever $M_t^{\pi} \geq \bar{x}_{\pi(t+1)}$, with the earliest optimal stopping time being given by $\boldsymbol{\tau} := \min\{t \geq 0 \mid M_t^{\pi} \geq \bar{x}_{\pi(t+1)}\}$

Pandora's Rule

Define $\overline{w}_n$: $k_n = \mathbb{E}[(W_n - \overline{w}_n)^+]$

Given V_n define $\overline{x}_n := V_n + \overline{w}_n - p_n$

Given order $\pi \in \Pi$, let $M_t^\pi := \max_{s \leq t} X_{\pi(t)} = \max_{s \leq t} V_{\pi(t)} + W_{\pi(t)} - p_{\pi(t)}$

Proposition

Given a price vector p and realization V, it is optimal for the consumer
(1) to learn about sellers in decreasing order of $\overline{x}_n$, with the optimal order of search being given by $\pi \in \Pi$ such that $\overline{x}_\pi(t) \geq \overline{x}_\pi(t+1)$; and
(2) to stop whenever $M_t^\pi \geq \overline{x}_{\pi(t+1)}$, with the earliest optimal stopping time being given by $\tau := \min\{t \geq 0 \mid M_t^\pi \geq \overline{x}_{\pi(t+1)}\}$

Proof

This is just Pandora's rule.

Consumer Behaviour

Theorem 1

Given a price vector p and realizations V, W, consumer i chooses product n if $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$ and only if $X_n \wedge \bar{x}_n \geq \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$

We've already proved a version of this for the satisficing setup.

(A proof for this specific setup is here.)

Write consumers' expected payoff given V as $\mathbb{E}[\max_n(X_n \wedge \bar{x}_n) \vee X_0 \mid V]$!

Intuition

Let $\overline{M} := \max_{n \in [T]} X_n \wedge \bar{x}_n \vee X_0$:

(1) n is chosen if $X_n \wedge \bar{x}_n > \overline{M}$, and

(2) the consumer learns about n whenever $\bar{x}_n > \overline{M}$, incurring in cost $k_n = \mathbb{E}[(X_n - \bar{x}_n)^+]$.

Market Equilibrium

Let (1) H_n be s.t. $X_n \wedge \bar{x}_n \sim H_n$;

(2) $\overline{H}_n$ be s.t. $\max_{m \neq n} X_m \wedge \bar{x}_m \vee X_{\mathbf{0}} \sim \overline{H}_n$.

Market Equilibrium

Let (1) H_n be s.t. $X_n \wedge \bar{x}_n \sim H_n$;
(2) $\overline{H}_n$ be s.t. $\max_{m \neq n} X_m \wedge \bar{x}_m \vee X_0 \sim \overline{H}_n$.

Demand for *n*:

$$D_n(p) = \mathbb{P}(\max_{m \neq n} X_m \wedge \bar{x}_m \vee X_0 \leq x < X_n \wedge \bar{x}_n) = \int (1 - H_n(x_n)) d\overline{H}_n(x_n).$$

p is **eqm price** if $\forall n$,

$$p_n \in \arg\max_{p'_n} D_n(p'_n, p_{-n})(p'_n - c_n).$$

FOC:

$$\frac{1}{p_n - c_n} = -\frac{\mathrm{d}D_n(p)/\,\mathrm{d}p_n}{D_n(p)}.$$

Market Equilibrium

Let (1) H_n be s.t. $X_n \wedge \bar{x}_n \sim H_n$;
(2) $\overline{H}_n$ be s.t. $\max_{m \neq n} X_m \wedge \bar{x}_m \vee X_0 \sim \overline{H}_n$.

Demand for *n*:

$$D_n(p) = \mathbb{P}(\max_{m \neq n} X_m \wedge \bar{x}_m \vee X_0 \leq x < X_n \wedge \bar{x}_n) = \int (1 - H_n(x_n)) d\overline{H}_n(x_n).$$

p is **eqm price** if $\forall n$,

$$p_n \in \arg\max_{p_n'} D_n(p_n', p_{-n})(p_n' - c_n).$$

FOC:

$$\frac{1}{p_n - c_n} = -\frac{\mathrm{d}D_n(p)/\,\mathrm{d}p_n}{D_n(p)}.$$

Assumption 1: H_n and $1 - H_n$ are log-concave $\forall n$.

Assumption 2: supp H_n has no upper bound $\forall n$.

Market Equilibrium

Theorem

Under Assumption 1, $D_n(p)$ is log-concave in p_n and $\log D_n(p)$ has strictly increasing differences in (p_n, p_m).

Under Assumptions 1 and 2, there is a unique Nash equilibrium of the pricing game, and this equilibrium is in pure strategies.

Market Equilibrium

Theorem

Under Assumption 1, $D_n(p)$ is log-concave in p_n and $\log D_n(p)$ has strictly increasing differences in (p_n, p_m).

Under Assumptions 1 and 2, there is a unique Nash equilibrium of the pricing game, and this equilibrium is in pure strategies.

Intuition

Existence of PSNE from game being supermodular (assumed away).

Uniqueness is due to particular structure of quasilinear preferences.

Beautiful result; ties-in:
(i) learning-based discrete choice with (ii) imperfect competition
within an eqm pricing model.

Market Equilibrium

Theorem

Under Assumption 1, $D_n(p)$ is log-concave in p_n and $\log D_n(p)$ has strictly increasing differences in (p_n, p_m).

Under Assumptions 1 and 2, there is a unique Nash equilibrium of the pricing game, and this equilibrium is in pure strategies.

Intuition

Existence of PSNE from game being supermodular (assumed away).

Uniqueness is due to particular structure of quasilinear preferences.

Beautiful result; ties-in:
(i) learning-based discrete choice with (ii) imperfect competition
within an eqm pricing model.

Issue: H_n is endogenous!

Market Equilibrium

Theorem 2

Under Assumption 1, $D_n(p)$ is log-concave in p_n and $\log D_n(p)$ has strictly increasing differences in (p_n, p_m).

Under Assumptions 1 and 2, there is a unique Nash equilibrium of the pricing game, and this equilibrium is in pure strategies.

Sufficient conditions for uniqueness

(i) f_n and g_n log-concave and $\sup \operatorname{supp} F_n = +\infty \implies \mathbf{1} - H_n$ log-concave.

(ii) F_n entails sufficiently high variance and ($f_n(\underline{v}_n) = \mathbf{0}$ or $\inf \operatorname{supp} F_n = -\infty$) $\implies H_n$ log-concave.

(iii) $f_n'(X_\mathbf{0} + c_n - \overline{w}_n) \leq \mathbf{0} \implies H_n$ log-concave on relevant part of its support.

Horizontal Differentiation

Imperfect Price Competition

Always exist *some* consumers who value n much more than others.

Typically this is *assumed*, e.g. loyal buyers vs shoppers, horizontal differentiation.

Here: degree of horizontal differentiation depend on V (and W).

Horizontal Differentiation

Imperfect Price Competition

- Always exist *some* consumers who value n much more than others.
- Typically this is *assumed*, e.g. loyal buyers vs shoppers, horizontal differentiation.
- Here: degree of horizontal differentiation depend on V (and W).

How to speak of more or less horizontal differentiation in this setting?

Horizontal Differentiation

Imperfect Price Competition

How to speak of more or less horizontal differentiation in this setting?

H_2 is **more dispersed** than H_1 if more dispersed for any w.

Higher values more dispersed.

Horizontal Differentiation

Imperfect Price Competition

How to speak of more or less horizontal differentiation in this setting?

Distribution H_2 is **more dispersed** than H_1 **above** w if $H_2^{-1}(b) - H_2^{-1}(a) \geq H_1^{-1}(b) - H_1^{-1}(a)$ for any $H_1(w) < a \leq b < 1$.

H_2 is **more dispersed** than H_1 if more dispersed for any w.

Higher values more dispersed.

Horizontal Differentiation

Imperfect Price Competition

How to speak of more or less horizontal differentiation in this setting?

Distribution H_2 is **more dispersed** than H_1 **above** w if $H_2^{-1}(b) - H_2^{-1}(a) \geq H_1^{-1}(b) - H_1^{-1}(a)$ for any $H_1(w) < a \leq b < 1$.

H_2 is **more dispersed** than H_1 if more dispersed for any w.

Higher values more dispersed.

Simplifying assumption: **symmetric environment**; i.e. for all $n, m \in [T]$ $F \equiv F_n = F_m$, $G \equiv G_n = G_m$, and $k \equiv k_n = k_m$ (implying $H \equiv H_n = H_m$), and $c \equiv c_n = c_m$.

Comparative Statics

Proposition 3

In symmetric environments, eqm price increases as H becomes more dispersed above $X_0 + c$ and $H(X_0 + c)$ weakly decreases.

Comparative Statics

Proposition 3

In symmetric environments, eqm price increases as H becomes more dispersed above $X_0 + c$ and $H(X_0 + c)$ weakly decreases.

Intuition

With higher preference heterogeneity (horizontal differentiation), sellers can charge higher mark-ups.

Comparative Statics

Proposition 3

In symmetric environments, eqm price increases as H becomes more dispersed above $X_0 + c$ and $H(X_0 + c)$ weakly decreases.

Intuition

With higher preference heterogeneity (horizontal differentiation), sellers can charge higher mark-ups.

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$:
(1) G (match values) becomes more dispersed and f is log-concave, or
(2) F (prior/common values) becomes more dispersed, g is log-concave, f is decreasing above $X_0 + c - \overline{w}$, and $\inf \operatorname{supp} F \leq X_0 + c - \overline{w}$.

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$:
(1) G (match values) becomes more dispersed and f is log-concave, (...)

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$:
(1) G (match values) becomes more dispersed and f is log-concave, (...)

Intuition for (1):

+ preference heterogeneity (+ differentiation) $\Longrightarrow$ sellers can charge higher mark-ups.

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$:
(1) G (match values) becomes more dispersed and f is log-concave, (...)

Intuition for (1):

- + preference heterogeneity (+ differentiation) $\Longrightarrow$ sellers can charge higher mark-ups.
- *Not* sufficient to have higher 'ex-post' preference heterogeneity: heterogeneity depends on how much consumers learn about their valuations. 'Effective' (or post-learning) preference heterogeneity.

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$:
(1) G (match values) becomes more dispersed and f is log-concave, (...)

Intuition for (1):

+ preference heterogeneity (+ differentiation) $\Longrightarrow$ sellers can charge higher mark-ups.

Not sufficient to have higher 'ex-post' preference heterogeneity: heterogeneity depends on how much consumers learn about their valuations.
'Effective' (or post-learning) preference heterogeneity.

+ dispersion in G $\Longrightarrow$ + dispersion 'ex-post' valuations; *and*
+ dispersion in G $\Longrightarrow$ + value to learning.
Combined, both reinforce + effective heterogeneity.

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$: (...)
(2) F (prior/common values) becomes more dispersed, g is log-concave, f is decreasing above $X_0 + c - \overline{w}$, and $\inf \operatorname{supp} F \leq X_0 + c - \overline{w}$.

Intuition for (2):

F + dispersed $\equiv$ + ex-ante heterogeneity.

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$: (...)
(2) F (prior/common values) becomes more dispersed, g is log-concave, f is decreasing above $X_0 + c - \overline{w}$, and $\inf \operatorname{supp} F \leq X_0 + c - \overline{w}$.

Intuition for (2):

F + dispersed $\equiv$ + ex-ante heterogeneity.

Implies higher ex-post heterogeneity, *but* may also preclude learning.

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$: (...)
(2) F (prior/common values) becomes more dispersed, g is log-concave, f is decreasing above $X_0 + c - \overline{w}$, and $\inf \operatorname{supp} F \leq X_0 + c - \overline{w}$.

Intuition for (2):

F + dispersed $\equiv$ + ex-ante heterogeneity.

Implies higher ex-post heterogeneity, *but* may also preclude learning.

Learning is source of effective heterogeneity, no conflicting forces.

Comparative Statics

Proposition 4

In symmetric environments, either condition is sufficient for H to become more dispersed above $X_0 + c$: (...)
(2) F (prior/common values) becomes more dispersed, g is log-concave, f is decreasing above $X_0 + c - \overline{w}$, and $\inf \operatorname{supp} F \leq X_0 + c - \overline{w}$.

Intuition for (2):

F + dispersed $\equiv$ + ex-ante heterogeneity.

Implies higher ex-post heterogeneity, *but* may also preclude learning.

Learning is source of effective heterogeneity, no conflicting forces.

Conditions in (2) ensure stronger incentives to learning, so that we have both forces going in same direction.

Comparative Statics

Proposition 5

In symmetric environments, if f is log-concave, both eqm price and profit decrease in cost to learning k.

Comparative Statics

Proposition 5

In symmetric environments, if f is log-concave, both eqm price and profit decrease in cost to learning k.

***Counter*-Intuitive**!

Expect higher learning costs to be exploited by sellers! (e.g. Diamond paradox)

Comparative Statics

Proposition 5

In symmetric environments, if f is log-concave, both eqm price and profit decrease in cost to learning k.

***Counter*-Intuitive**!

- Expect higher learning costs to be exploited by sellers! (e.g. Diamond paradox)
- Why do we get this then?
  - Intuition relies on presumption that prices not known to consumers prior to learning.
  - Here: prices known; it's valuations that are not known.

Comparative Statics

Proposition 5

In symmetric environments, if f is log-concave, both eqm price and profit decrease in cost to learning k.

Intuition: Prices known in advance.

Comparative Statics

Proposition 5

In symmetric environments, if f is log-concave, both eqm price and profit decrease in cost to learning k.

Intuition: Prices known in advance.

If seller lowers price, + consumers learning about seller's product, + demand.

Comparative Statics

Proposition 5

In symmetric environments, if f is log-concave, both eqm price and profit decrease in cost to learning k.

Intuition: Prices known in advance.

- If seller lowers price, + consumers learning about seller's product, + demand.
- When learning costs are higher, consumers less willing to learn and demand + sensitive to prices.

Comparative Statics

Proposition 5

In symmetric environments, if f is log-concave, both eqm price and profit decrease in cost to learning k.

Intuition: Prices known in advance.

- If seller lowers price, + consumers learning about seller's product, + demand.
- When learning costs are higher, consumers less willing to learn and demand + sensitive to prices.
- Thus: $\uparrow$ learning costs $\implies$ + price competition $\implies$ $\downarrow$ eqm price $\implies$ $\downarrow$ profits. (this last bit is not immediate: outside option)

Summary

Search as flexible framework for directed info acquisition

- a lot to be done...
  - Directed job search from first principles;
  - Market entry with learning;
  - Structural estimation of consumer demand;
  - Manipulation/obfuscation via search costs;
  - Uncertainty about prices and valuations;
  - Dynamic pricing;
  - Inference based on decision times (see Choi & Smith for comparative statics);
  - Optimal R&D funding design.

Pandora's Optimal Search Order

Proof (Back)

Suppose $\bar{x}_{\pi(s)}$ nonincreasing for $s \geq t+1$, but $\bar{x}_{\pi(t)} < \bar{x}_{\pi(t+1)}$.

Assume that $\tau^{\pi} \geq t$ occurs with positive prob., as ow it is WL to rearrange the order following t (no impact on payoffs).

Define new order δ: same as π except it swaps t-th and $(t+1)$-th alternatives.

Define stopping time $\tau := \min\{s \geq 0 \mid M_s^{\delta} \geq \max_{n>s} \bar{x}_{\pi(s)}\}$

Then $\{\tau \leq s\} = \{\tau^{\pi} \leq s\}, \forall s \neq t$; i.e. only difference between τ^{π} and τ is if they disagree in stopping *at* t.

$\{\tau \in \{t, t+1\}\} = \{\tau^{\pi} \in \{t, t+1\}\} = \{M_{t-1}^{\pi} < \bar{x}_{\pi(t+1)}$ and $M_{t+1}^{\pi} \geq \bar{x}_{\pi(t+2)}\}$.

$\{\tau^{\pi} = t\} = \{M_{t-1}^{\pi} < \bar{x}_{\pi(t+1)} = \max_{s>t-1} \bar{x}_{\pi(s)}$ and $X_t = M_t \geq \bar{x}_{\pi(t+1)} = \max_{s>t-1} \bar{x}_{\pi(s)}\}$, whereas $\{\tau = t\} = \{M_{t-1}^{\pi} < \bar{x}_{\pi(t+1)} = \max_{s>t-1} \bar{x}_{\pi(s)}$ and $X_{t+1} = M_t^{\delta} \geq \bar{x}_{\pi(t+1)} = \max_{s>t-1} \bar{x}_{\pi(s)}\}$.

$\mathbb{E}[Y_{\tau^{\pi}}^{\pi} - Y_{\tau}^{\delta}] = \mathbb{E}[\mathbf{1}_{\{\tau \in \{t,t+1\}\}}(Y_{\tau^{\pi}}^{\pi} - Y_{\tau}^{\delta})]$

Pandora's Optimal Search Order

Proof **(Back)**

$\{\tau \in \{t, t+1\}\} = \{\tau^\pi \in \{t, t+1\}\} = \{M^\pi_{t-1} < \bar{x}_{\pi(t+1)}\} \cap \{M^\pi_{t+1} \geq \bar{x}_{\pi(t+2)}\}$. Furthermore,

$\{\tau^\pi = t\} = \{M^\pi_{t-1} < \max_{s>t-1} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+1)}\} \cap \{X^\pi_t = M^\pi_t \geq \max_{s>t-1} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+1)}\}$;

$\{\tau = t\} = \{M^\delta_{t-1} = M^\pi_{t-1} < \max_{s>t-1} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+1)}\} \cap \{X^\pi_{t+1} = M^\delta_t \geq \max_{s>t-1} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+1)}\}$.

$\{\tau^\pi = t+1\} = \{X^\pi_t \leq M^\pi_t < \max_{s>t} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+1)}\} \cap \{M^\pi_{t+1} = M^\delta_{t+1} \geq \max_{s>t+1} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+2)}\}$;

$\{\tau = t+1\} = \{X^\pi_{t+1} \leq M^\delta_t < \max_{s>t} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+1)}\} \cap \{M^\pi_{t+1} = M^\delta_{t+1} \geq \max_{s>t} \bar{x}_{\pi(s)} = \bar{x}_{\pi(t+2)}\}$.

Pandora's Optimal Search Order

Proof **(Back)**

Note that

$$\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_t^\pi\wedge X_{t+1}^\pi\geq\bar{x}_{\pi(t+1)}\}}(Y_{\tau^\pi}^\pi - Y_\tau^\delta)$$
$$= \mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_t^\pi\wedge X_{t+1}^\pi\geq\bar{x}_{\pi(t+1)}\}}((X_t^\pi - X_{t+1}^\pi) - (c_{\pi(t)} - c_{\pi(t+1)}))$$

$$\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_t^\pi\geq\bar{x}_{\pi(t+1)}>X_{t+1}^\pi\}}(Y_{\tau^\pi}^\pi - Y_\tau^\delta)$$
$$= \mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_t^\pi\geq\bar{x}_{\pi(t+1)}>X_{t+1}^\pi\}}((X_t^\pi - X_t^\pi) - (c_{\pi(t)} - c_{\pi(t+1)} - c_{\pi(t)}))$$
$$= \mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_t^\pi\geq\bar{x}_{\pi(t+1)}>X_{t+1}^\pi\}}c_{\pi(t+1)}$$

$$\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_{t+1}^\pi\geq\bar{x}_{\pi(t+1)}>X_t^\pi\}}(Y_{\tau^\pi}^\pi - Y_\tau^\delta)$$
$$= \mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_{t+1}^\pi\geq\bar{x}_{\pi(t+1)}>X_t^\pi\}}((X_{t+1}^\pi - X_{t+1}^\pi) - (c_{\pi(t)} + c_{\pi(t+1)} - c_{\pi(t+1)}))$$
$$= -\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{X_{t+1}^\pi\geq\bar{x}_{\pi(t+1)}>X_t^\pi\}}c_{\pi(t)}$$

$$\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{\bar{x}_{\pi(t+1)}>X_t^\pi\vee X_{t+1}^\pi\}}(Y_{\tau^\pi}^\pi - Y_\tau^\delta)$$
$$= \mathbf{1}_{\{\tau\in\{t,t+1\}\}}\mathbf{1}_{\{\bar{x}_{\pi(t+1)}>X_t^\pi\vee X_{t+1}^\pi\}}((M_{t-1}^\pi - M_{t-1}^\pi) - (c_{\pi(t)} + c_{\pi(t+1)} - c_{\pi(t+1)} - c_{\pi(t)})) = 0$$

Pandora's Optimal Search Order

Proof **(Back)**

Since $\mathbf{1} = \mathbf{1}_{\{X_t^\pi \wedge X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}} + \mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)} > X_{t+1}^\pi\}} + \mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)} > X_t^\pi\}} + \mathbf{1}_{\{X_t^\pi \vee X_{t+1}^\pi < \bar{x}_{\pi(t+1)}\}}$ we get:

$$\mathbb{E}[Y_{\rho^\pi}^\pi - Y_\tau^\delta]$$

$$= \mathbb{E}\left[\mathbf{1}_{\{\tau \in \{t,t+1\}\}}\left(\mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}}((X_t^\pi - X_{t+1}^\pi) - (c_{\pi(t)} - c_{\pi(t+1)})) + \mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)} > X_{t+1}^\pi\}}c_{\pi(t+1)} - \mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)} > X_t^\pi\}}c_{\pi(t)}\right)\right]$$

$$= \mathbb{E}\left[\mathbf{1}_{\{\tau \in \{t,t+1\}\}}\left(\mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}}(X_t^\pi - X_{t+1}^\pi) + \mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)}\}}c_{\pi(t+1)} - \mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}}c_{\pi(t)}\right)\right]$$

Pandora's Optimal Search Order

Proof **(Back)**

Recalling that $c_{\pi(t)} := \mathbb{E}[(X_t^\pi - \bar{x}_{\pi(t)})^+]$ and $\bar{x}_{\pi(t)} < \bar{x}_{\pi(t+1)}$, we get

$$\mathbb{E}[Y^\pi_{\rho^\pi} - Y^\delta_\tau]$$

$$= \mathbb{E}\Bigg[\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\Bigg(\mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}}(X_t^\pi - X_{t+1}^\pi)$$

$$+ \mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbb{E}[(X_{t+1}^\pi - \bar{x}_{\pi(t+1)})^+] - \mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbb{E}[(X_t^\pi - \bar{x}_{\pi(t)})^+]\Bigg)\Bigg]$$

$$< \mathbb{E}\Bigg[\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\Bigg(\mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}}(X_t^\pi - X_{t+1}^\pi)$$

$$+ \mathbf{1}_{\{X_t^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbb{E}[(X_{t+1}^\pi - \bar{x}_{\pi(t+1)})^+] - \mathbf{1}_{\{X_{t+1}^\pi \geq \bar{x}_{\pi(t+1)}\}}\mathbb{E}[(X_t^\pi - \bar{x}_{\pi(t+1)})^+]\Bigg)\Bigg]$$

Pandora's Optimal Search Order

Proof **(Back)**

$$\mathbb{E}[Y^{\pi}_{\rho^{\pi}} - Y^{\delta}_{\tau}] < \mathbb{E}\Bigg[\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\Bigg(\mathbf{1}_{\{X^{\pi}_t \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X^{\pi}_{t+1} \geq \bar{x}_{\pi(t+1)}\}}(X^{\pi}_t - X^{\pi}_{t+1})$$

$$+ \mathbf{1}_{\{X^{\pi}_t \geq \bar{x}_{\pi(t+1)}\}}\mathbb{E}[(X^{\pi}_{t+1} - \bar{x}_{\pi(t+1)})^+] - \mathbf{1}_{\{X^{\pi}_{t+1} \geq \bar{x}_{\pi(t+1)}\}}\mathbb{E}[(X^{\pi}_t - \bar{x}_{\pi(t+1)})^+]\Bigg)\Bigg]$$

$$= \mathbb{E}\Bigg[\mathbf{1}_{\{\tau\in\{t,t+1\}\}}\Bigg(\mathbf{1}_{\{X^{\pi}_t \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X^{\pi}_{t+1} \geq \bar{x}_{\pi(t+1)}\}}(X^{\pi}_t - X^{\pi}_{t+1})$$

$$+ \mathbf{1}_{\{X^{\pi}_t \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X^{\pi}_{t+1} \geq \bar{x}_{\pi(t+1)}\}}(X^{\pi}_{t+1} - \bar{x}_{\pi(t+1)}) - \mathbf{1}_{\{X^{\pi}_t \geq \bar{x}_{\pi(t+1)}\}}\mathbf{1}_{\{X^{\pi}_{t+1} \geq \bar{x}_{\pi(t+1)}\}}(X^{\pi}_t - \bar{x}_{\pi(t+1)})\Bigg)\Bigg]$$

$$= 0$$

Consumer Behaviour

Theorem

Given a price vector p and realizations V, W, consumer i chooses product n if $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$ and only if $X_n \wedge \bar{x}_n \geq \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

Theorem

Given a price vector p and realizations V, W, consumer i chooses product n if $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$ and only if $X_n \wedge \bar{x}_n \geq \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

Proof **(Back)**

Only if:

- Case 1: $X_n \wedge \bar{x}_n < \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.
  - Case 1a: $X_n \wedge \bar{x}_n = X_n$.

    If $X_n < X_0$, consumer will never purchase n.

    If $X_n < X_m \wedge \bar{x}_m$ and consumer stops after $\pi^{-1}(m)$, they will never choose n.

    If $X_n < X_m \wedge \bar{x}_m$ and stops at $t < \pi^{-1}(m)$, then $M_t^{\pi} > \bar{x}_m \geq X_m \wedge \bar{x}_m > X_n$ and so they will not choose n.

Consumer Behaviour

Theorem

Given a price vector p and realizations V, W, consumer i chooses product n if $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$ and only if $X_n \wedge \bar{x}_n \geq \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

Proof **(Back)**

Only if:

- Case 1: $X_n \wedge \bar{x}_n < \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

– Case 1b: $X_n \wedge \bar{x}_n = \bar{x}_n < X_n$.

If $\pi^{-1}(m) < \pi^{-1}(n)$, then $M^{\pi}_{\pi^{-1}(n)-1} \geq (X_m \wedge \bar{x}_m) \vee X_0 > \bar{x}_n$ and so τ calls for stopping and consumer never learns about product n and thus never buys it.

If $\pi^{-1}(m) > \pi^{-1}(n)$ and $(X_m \wedge \bar{x}_m) \geq X_0$, then $X_n > \bar{x}_n \geq \bar{x}_m \geq X_m \wedge \bar{x}_m > X_n \wedge \bar{x}_n = \bar{x}_n$, a contradiction.

If $X_n \wedge \bar{x}_n = \max_m(X_m \wedge \bar{x}_m) < X_0$, then if stop at t (but not earlier) $X_0 \leq M^{\pi}_{s-1} < \bar{x}_{\pi(s)}$.

This implies $X_0 \geq X_{\pi(s)}$, for all $s = 1, ..., t$, therefore $M^{\pi}_t = X_0$.

By assumption, $\bar{x}_n < X_n$, and so if $\pi^{-1}(n) \leq t$, $X_0 > X_n$ and n cannot be chosen.

Theorem

Given a price vector p and realizations V, W, consumer i chooses product n if $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$ and only if $X_n \wedge \bar{x}_n \geq \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

Proof **(Back)**

If:

- Case 2: $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

– Case 2a: $X_n \geq \bar{x}_n$.

Then $M_{\pi^{-1}(n)} \geq X_n \geq \bar{x}_n \geq \bar{x}_{\pi^{-1}(n)+1}$ and therefore $\tau \leq \pi^{-1}(n)$.

As $\bar{x}_n > X_m \wedge \bar{x}_m$, for all $m : \pi^{-1}(m) < \pi^{-1}(n)$,

$\implies \bar{x}_m \geq \bar{x}_n$ and $X_n \geq \bar{x}_n > X_m$

$\implies n$ must be chosen.

Theorem

Given a price vector p and realizations V, W, consumer i chooses product n if $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$ and only if $X_n \wedge \bar{x}_n \geq \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

Proof **(Back)**

If:

- Case 2: $X_n \wedge \bar{x}_n > \max_{m \neq n}(X_m \wedge \bar{x}_m) \vee X_0$.

– Case 2b: $X_n < \bar{x}_n$.

Then $\bar{x}_m \geq \bar{x}_n > X_n > X_m \wedge x_m$ for all $m : \pi^{-1}(m) < \pi^{-1}(n)$ and $\tau \geq \pi^{-1}(n)$.

If the consumer stops at $t > \pi^{-1}(n)$, then

$\bar{x}_s > X_n > X_s \wedge \bar{x}_s$ for all $s \leq t$ such that $s \neq \pi^{-1}(n)$ and $M_t^{\pi} = X_n$

$\implies$ n is chosen.