Competitive Information Design in Pandora's Box

Bolin Ding[1], Yiding Feng[2], Chien-Ju Ho[3], Wei Tang[4], Haifeng Xu[5]

[1]Alibaba Research Group

[2]Microsoft Research, New England

[3]Washington University in St. Louis

[4]Columbia University

[5]Universty of Chicago

Outline

- Background: Pandora’s Box
- Informational Properties of Pandora’s Box
- Competitive Information Design, and Equilibrium

The Pandora's Box Problem [Weitzman, Econometrica'79]

- n boxes, box i has a random reward $R_i \sim G_i$, supported on $[0,1]$
- An agent can open box at cost c_i to observe realized reward r_i
- Can claim the reward from one of the opened boxes

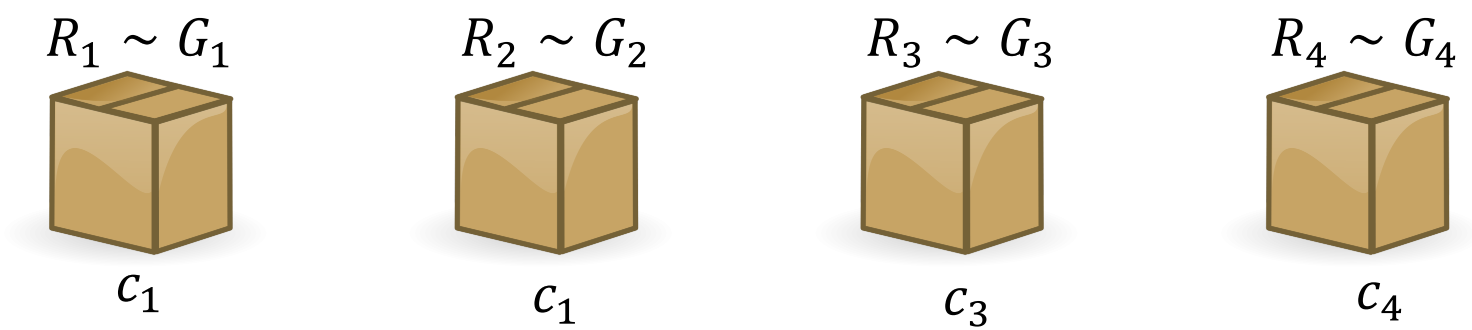

The Pandora's Box Problem [Weitzman, Econometrica'79]

- n boxes, box i has a random reward $R_i \sim G_i$, supported on $[0,1]$
- An agent can open box at cost c_i to observe realized reward r_i
- Can claim the reward from one of the opened boxes

Question: What's the utility-maximizing "dynamic search" policy?

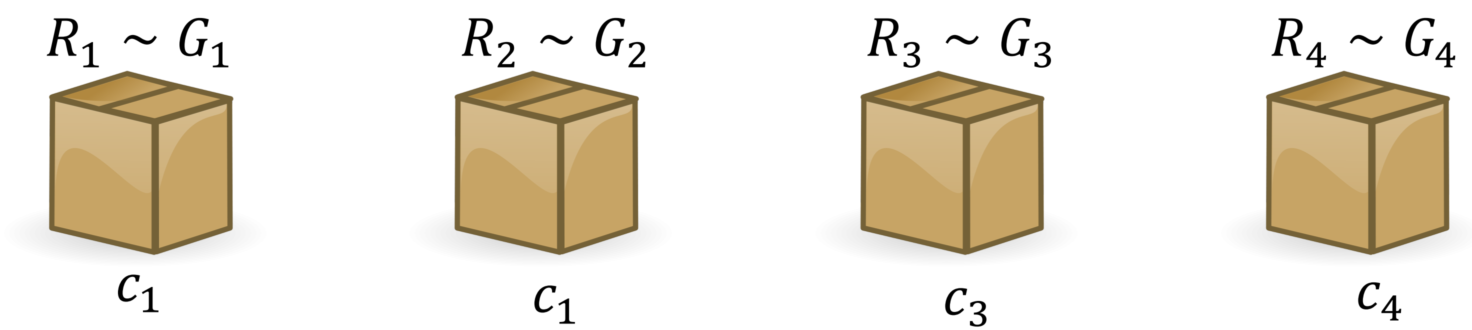

The Pandora's Box Problem [Weitzman, Econometrica'79]

- n boxes, box i has a random reward $R_i \sim G_i$, supported on $[0,1]$
- An agent can open box at cost c_i to observe realized reward r_i
- Can claim the reward from one of the opened boxes

Question: What's the utility-maximizing "dynamic search" policy?

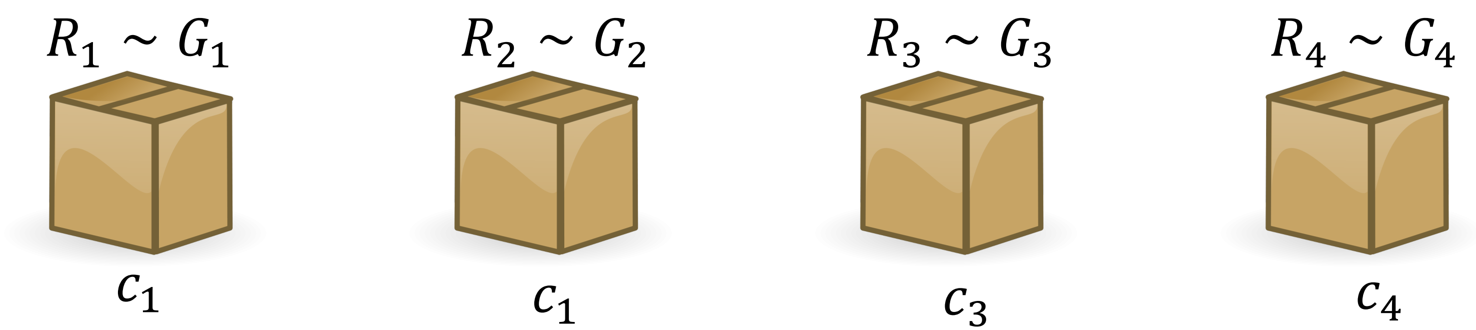

✓ Numerous applications: look for startups to fund, open house, find channels to subscribe

The Pandora's Box Problem [Weitzman, Econometrica'79]

Question: What's the utility-maximizing "dynamic search" policy?

An elegant greedy solution

- Capped reward: $\kappa_i(\sigma) = \min\{R_i, \sigma\}$
- Reservation value σ_i is the σ such that $\mathbb{E}[R_i] - \mathbb{E}[\kappa_i(\sigma)] = c_i$

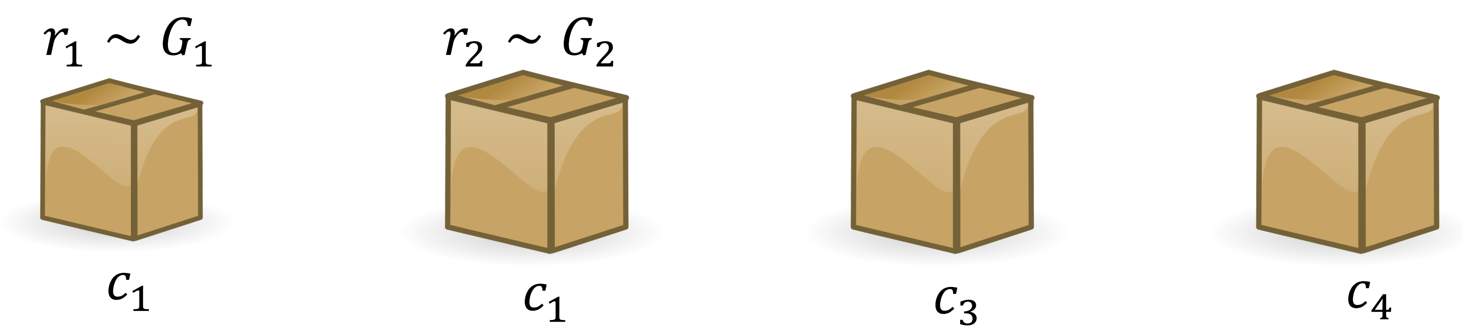

The Pandora's Box Problem [Weitzman, Econometrica'79]

Question: What's the utility-maximizing "dynamic search" policy?

An elegant greedy solution

- Capped reward: $\kappa_i(\sigma) = \min\{R_i, \sigma\}$
- Reservation value σ_i is the σ such that $\mathbb{E}[R_i] - \mathbb{E}[\kappa_i(\sigma)] = c_i$
- Optimal algorithm:
    1. Sort boxes in decreasing order of σ_i, and open sequentially

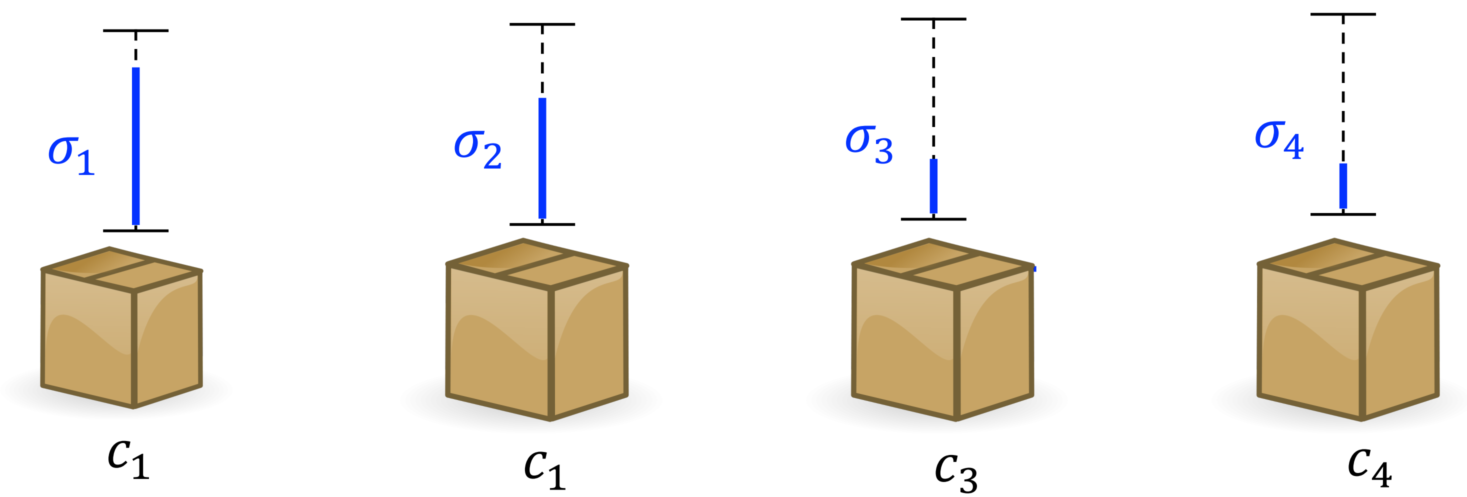

The Pandora's Box Problem [Weitzman, Econometrica'79]

Question: What's the utility-maximizing "dynamic search" policy?

An elegant greedy solution

- Capped reward: $\kappa_i(\sigma) = \min\{R_i, \sigma\}$
- Reservation value σ_i is the σ such that $\mathbb{E}[R_i] - \mathbb{E}[\kappa_i(\sigma)] = c_i$
- Optimal algorithm:
  1. Sort boxes in decreasing order of σ_i, and open sequentially
  2. Stop when largest reward so far exceeds next box's reservation value

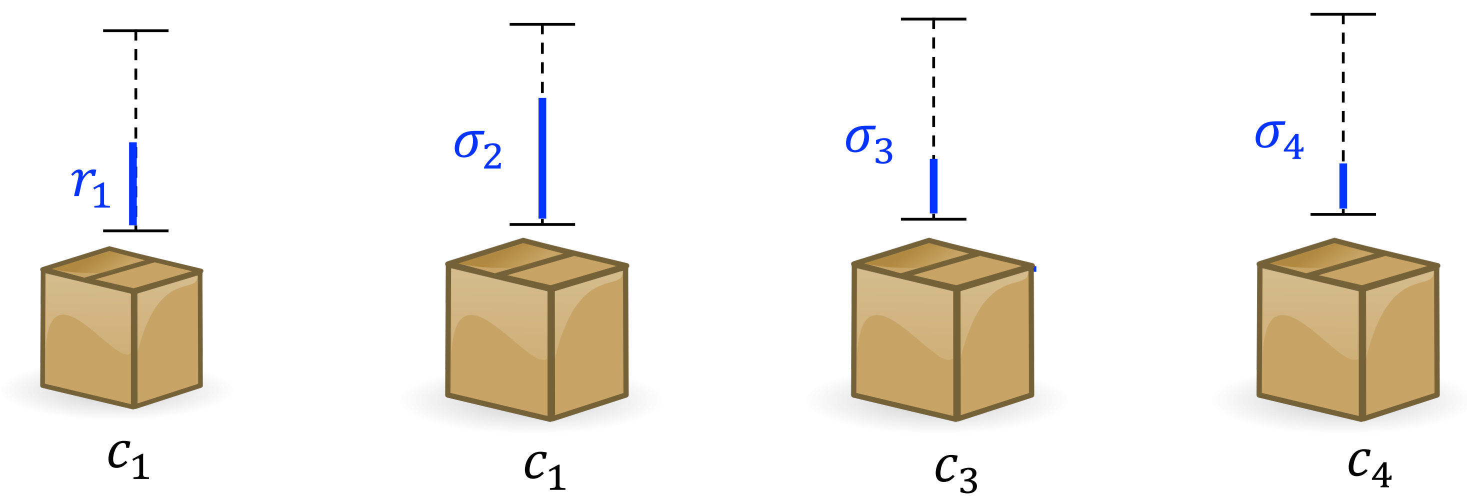

The Pandora's Box Problem [Weitzman, Econometrica'79]

Question: What's the utility-maximizing "dynamic search" policy?

An elegant greedy solution

- Capped reward: $\kappa_i(\sigma) = \min\{R_i, \sigma\}$
- Reservation value σ_i is the σ such that $\mathbb{E}[R_i] - \mathbb{E}[\kappa_i(\sigma)] = c_i$
- Optimal algorithm:
  1. Sort boxes in decreasing order of σ_i, and open sequentially
  2. Stop when largest reward so far exceeds next box's reservation value

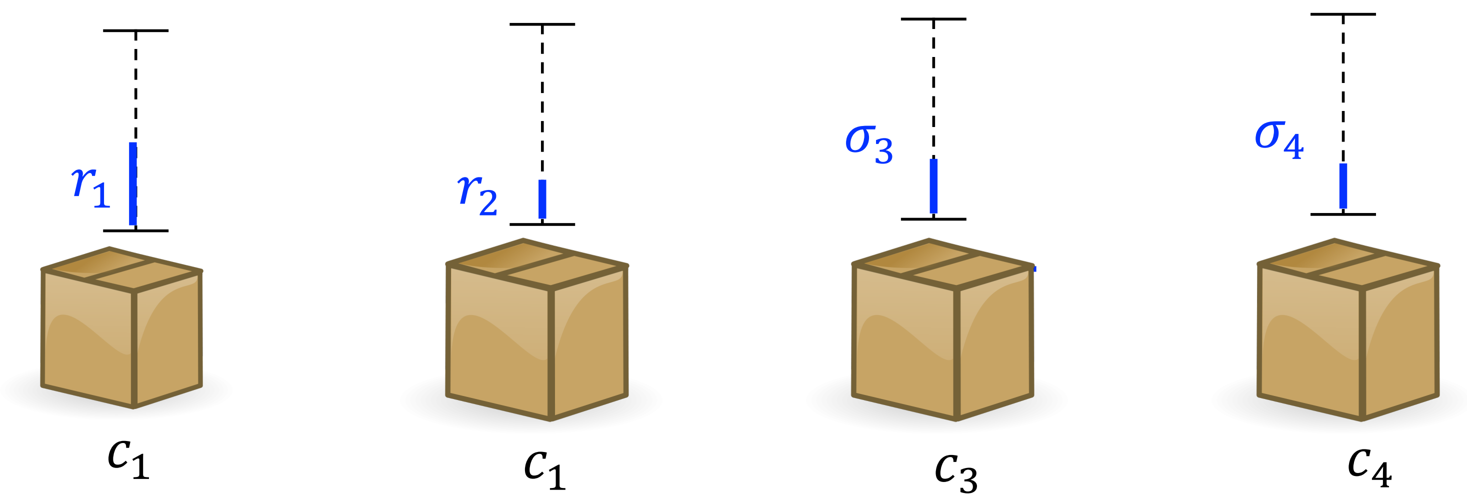

The Pandora's Box Problem [Weitzman, Econometrica'79]

Question: What's the utility-maximizing "dynamic search" policy?

An elegant greedy solution

- Capped reward: $\kappa_i(\sigma) = \min\{R_i, \sigma\}$
- Reservation value σ_i is the σ such that $\mathbb{E}[R_i] - \mathbb{E}[\kappa_i(\sigma)] = c_i$
- Optimal algorithm: [Weitzman, Econometrica'79; Kleinberg et al. EC'16]
  1. Sort boxes in decreasing order of σ_i, and open sequentially
  2. Stop when largest reward so far exceeds next box's reservation value
  3. Claim the largest reward so far

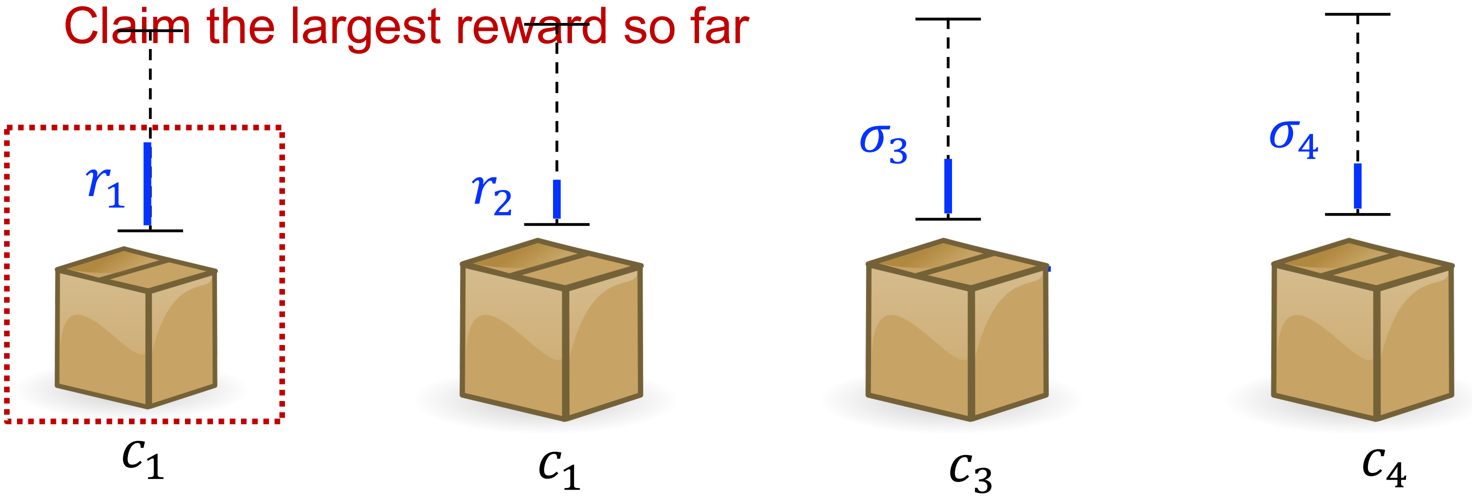

Outline

- Background: Pandora's Box
- Informational Properties of Pandora's Box
- Competitive Information Design, and Equilibrium

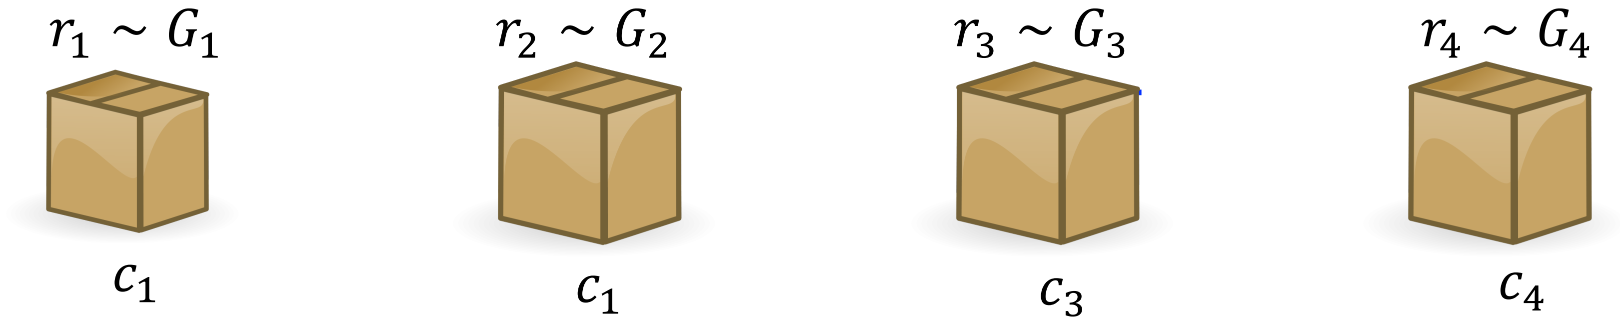

Question: how do things change if G_i becomes “less informative”?

Pandora’s box is fundamentally about tradeoff between cost and reward from information acquisition

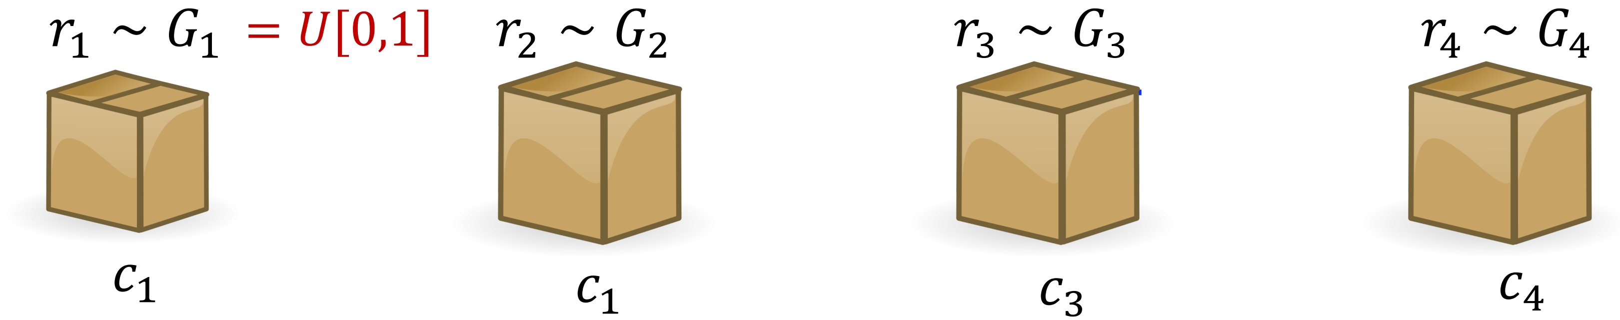

Question: how do things change if G_i becomes “less informative”?

- For example, box 1 only tells agent whether $r_1 \geq \frac{1}{2}$ or $r_1 < \frac{1}{2}$, but does not directly reveal r_i

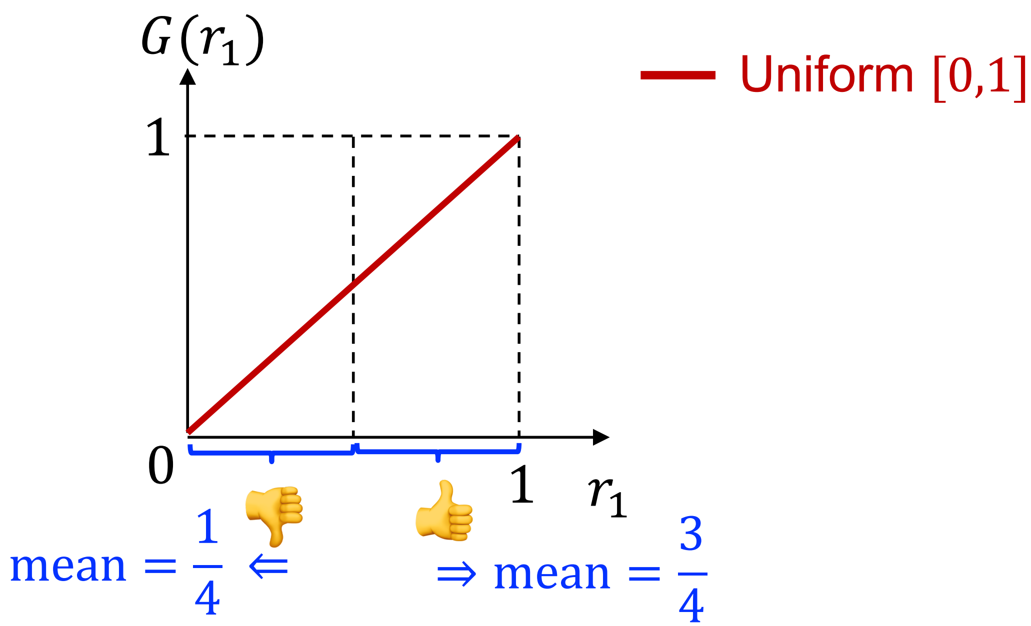

$r_1 \sim G_1 = U[0,1]$ $r_2 \sim G_2$ $r_3 \sim G_3$ $r_4 \sim G_4$

c_1 c_1 c_3 c_4

Question: how do things change if G_i becomes "less informative"?

- For example, box 1 only tells agent whether $r_1 \geq \frac{1}{2}$ or $r_1 < \frac{1}{2}$, but does not directly reveal r_i

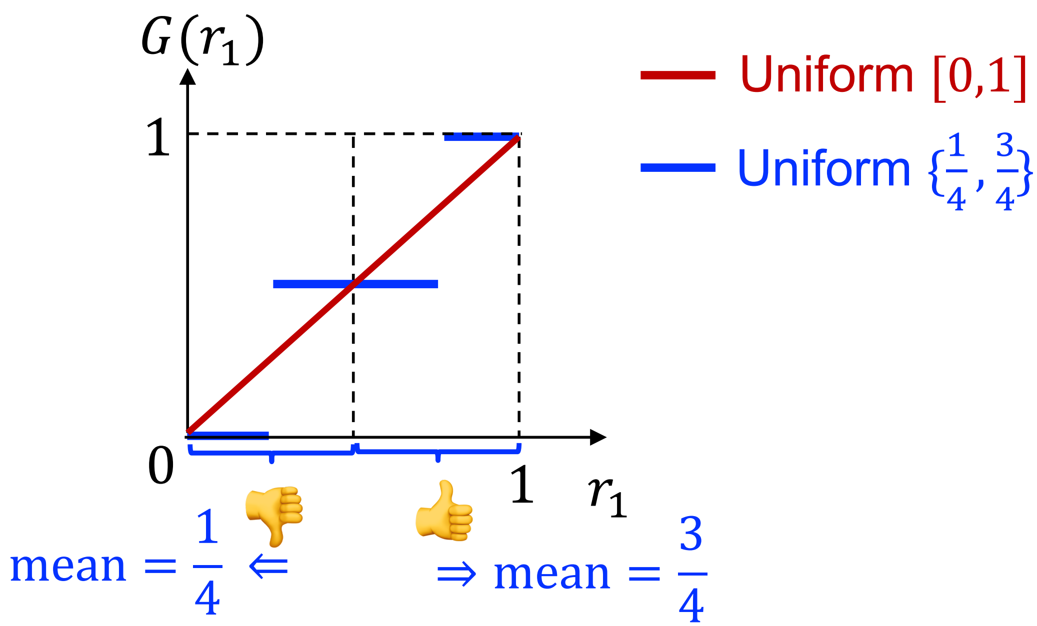

$r_1 \sim G_1 = U[0,1]$ $r_2 \sim G_2$ $r_3 \sim G_3$ $r_4 \sim G_4$

c_1 c_1 c_3 c_4

Question: how do things change if G_i becomes "less informative"?

- For example, box 1 only tells agent whether $r_1 \geq \frac{1}{2}$ or $r_1 < \frac{1}{2}$, but does not directly reveal r_i

signaling scheme

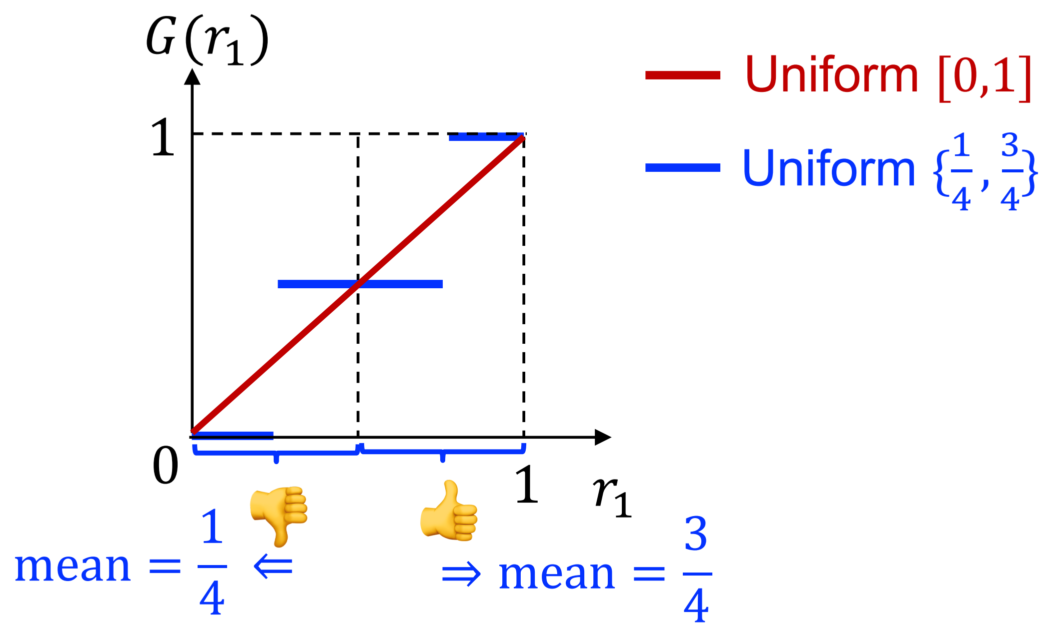

$r_1 \sim G_1 = U[0,1]$ $r_2 \sim G_2$ $r_3 \sim G_3$ $r_4 \sim G_4$

c_1 c_1 c_3 c_4

Question: how do things change if G_i becomes "less informative"?

- For example, box 1 only tells agent whether $r_1 \geq \frac{1}{2}$ or $r_1 < \frac{1}{2}$, but does not directly reveal r_i

signaling scheme

- Why?
  - Learning accurate information is costly for a box, if not impossible
  - Moreover, revealing all information is not necessarily most useful

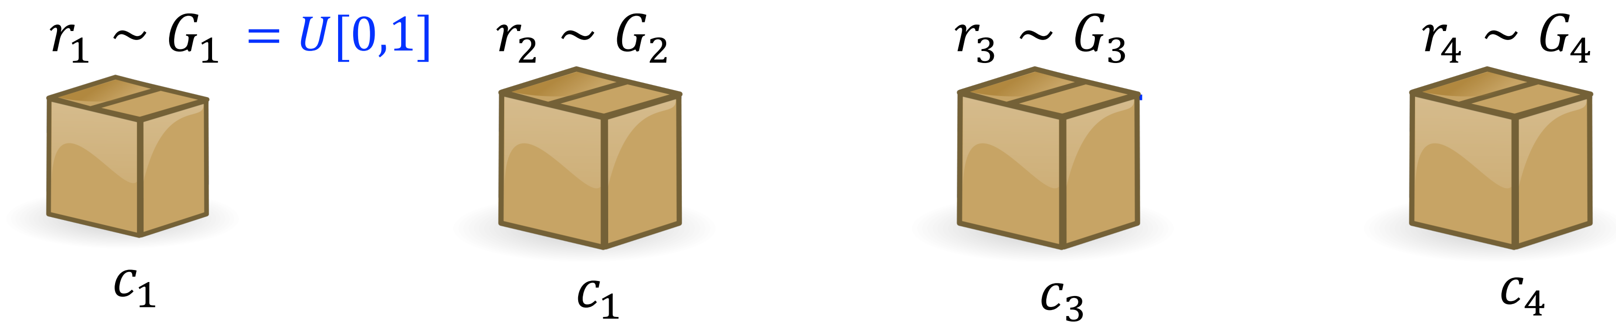

Question: how do things change if G_i becomes "less informative"?

- For example, box 1 only tells agent whether $r_1 \geq \frac{1}{2}$ or $r_1 < \frac{1}{2}$, but does not directly reveal r_i

  signaling scheme

- Order of "informativeness" of distributions has a characterization

Theorem [Blackwell'79; Aumann et al., '95]. Suppose $r \sim G$.

There exists a signaling scheme that induces some distribution H over *mean rewards*, if and only if

G is a mean-preserving spread of H.

$r_1 \sim G_1 = U[0,1]$ $r_2 \sim G_2$ $r_3 \sim G_3$ $r_4 \sim G_4$

c_1 c_1 c_3 c_4

Question: how do things change if G_i becomes "less informative"?

- For example, box 1 only tells agent whether $r_1 \geq \frac{1}{2}$ or $r_1 < \frac{1}{2}$, but does not directly reveal r_i
  signaling scheme
- Order of "informativeness" of distributions has a characterization

Theorem [Blackwell'79; Aumann et al., '95]. Suppose $r \sim G$.

There exists a signaling scheme that induces some distribution H over *mean rewards*, if and only if

G is a mean-preserving spread of H.

(i.e., $\int_0^\sigma G(x)dx \geq \int_0^\sigma H(x)dx$, with equality at $\sigma = 1$)

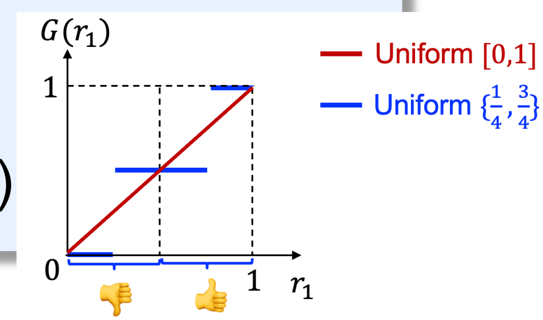

Our First Result

Pandora's Box provides an alternative characterization for the order of distributions' informativeness

Thm 1. Let $U(G_i, G_{-i})$ denote agent's optimal utility in Pandora's Box. G_i is a mean-preserving spread of H, if and only if

$$U(G_i, G_{-i}) \geq U(H, G_{-i}), \qquad \forall\, G_{-i}, \forall\, \{c_i\}_{i \in [n]}$$

G_{-i} contains all boxes' reward distributions, excluding i'th.

Our First Result

Pandora's Box provides an alternative characterization for the order of distributions' informativeness

Thm 1. Let $U(G_i, G_{-i})$ denote agent's optimal utility in Pandora's Box. G_i is a mean-preserving spread of H, if and only if

$$U(G_i, G_{-i}) \geq U(H, G_{-i}), \qquad \forall\, G_{-i}, \forall\, \{c_i\}_{i \in [n]}$$

Remarks

- More informativeness from any box induces larger agent utility
- A decision-theoretic characterization of "informativeness"

Our First Result

Pandora's Box provides an alternative characterization for the order of distributions' informativeness

Thm 1. Let $U(G_i, G_{-i})$ denote agent's optimal utility in Pandora's Box. G_i is a mean-preserving spread of H, if and only if

$$U(G_i, G_{-i}) \geq U(H, G_{-i}), \qquad \forall\, G_{-i}, \forall\, \{c_i\}_{i\in[n]}$$

Remarks

- More informativeness from any box induces larger agent utility
- A decision-theoretic characterization of "informativeness"
- Why not obvious?
  - Hiding information may lead to early stop of agent's costly search, thus reduce total search cost

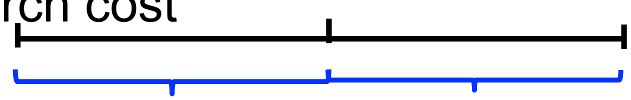

$\text{mean} = \frac{1}{4} \leftarrow r_1 < \frac{1}{2}$ $\qquad$ $r_1 \geq \frac{1}{2} \rightarrow \text{mean} = \frac{3}{4}$

Our First Result

Pandora's Box provides an alternative characterization for the order of distributions' informativeness

Thm 1. Let $U(G_i, G_{-i})$ denote agent's optimal utility in Pandora's Box. G_i is a mean-preserving spread of H, if and only if

$$U(G_i, G_{-i}) \geq U(H, G_{-i}), \qquad \forall\, G_{-i}, \forall\, \{c_i\}_{i \in [n]}$$

Remarks

- More informativeness from any box induces larger agent utility
- A decision-theoretic characterization of "informativeness"
- Why not obvious?
  - Hiding information may lead to early stop of agent's costly search, thus reduce total search cost
- Theorem shows search cost reduction never exceeds reward loss

Our First Result

Pandora's Box provides an alternative characterization for the order of distributions' informativeness

Thm 1. Let $U(G_i, G_{-i})$ denote agent's optimal utility in Pandora's Box.

G_i is a mean-preserving spread of H, if and only if

$$U(G_i, G_{-i}) \geq U(H, G_{-i}), \qquad \forall\, G_{-i}, \forall\, \{c_i\}_{i \in [n]}$$

Proof highlight

- A useful lemma we show is *less informative distribution always has smaller reservation value*

Outline

- Background: Pandora's Box
- Informational Properties of Pandora's Box
- Competitive Information Design, and Equilibrium

Pandora's Box with Strategic Boxes

$r_1 \sim G_1$ $r_2 \sim G_2$ $r_3 \sim G_3$ $r_4 \sim G_4$

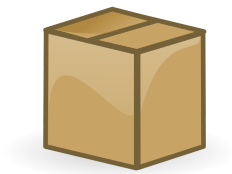
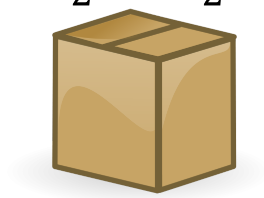
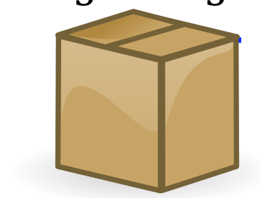
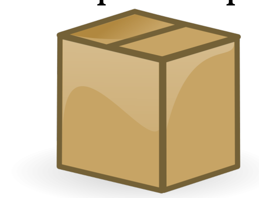

c_1 c_1 c_3 c_4

Venture capital searches for a good startup to invest

Pandora's Box with Strategic Boxes

$r_1 \sim G_1$ $r_2 \sim G_2$ $r_3 \sim G_3$ $r_4 \sim G_4$

c_1 c_1 c_3 c_4

Search for a PhD admission through their open houses

Pandora's Box with Strategic Boxes

$r_1 \sim G_1$ $r_2 \sim G_2$ $r_3 \sim G_3$ $r_4 \sim G_4$

c_1 c_1 c_3 c_4

The New York Times | FT FINANCIAL TIMES | The Economist | wp WASHINGTON POST

Search for a newsletter to subscribe to

Pandora's Box with Strategic Boxes

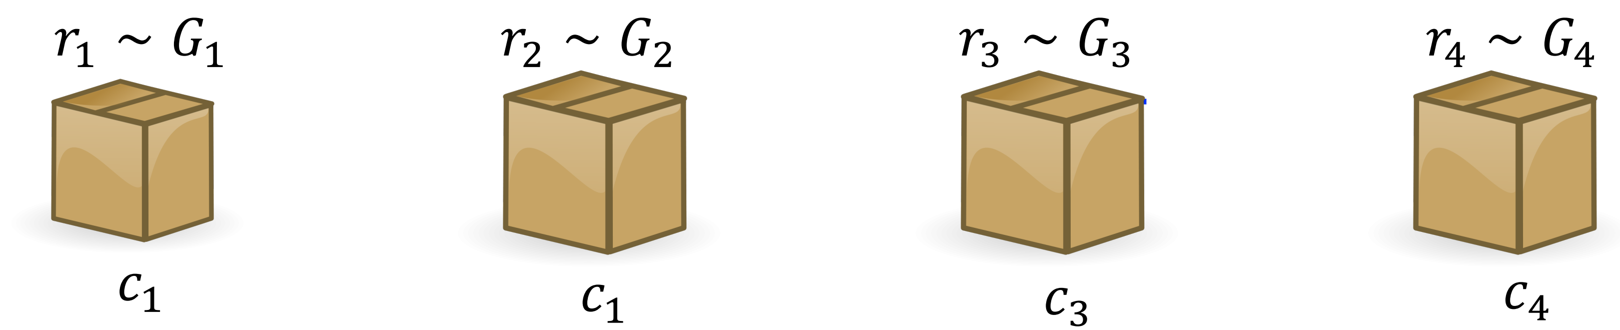

Competitive Information Design in PB

- Each box is a strategic agent
  - Maximize probability of being chosen
  - May signal partial information to increase their chance
- What is the equilibrium among boxes, assuming agent always follows with a best search?

multiple-leader-single-follower Stackelberg game

Second Main Result

Characterization of symmetric equilibrium for symmetric environment

- Each $r_i \sim G$ i.i.d.; equal cost $c_i = c$
- Box i reveals r_i partially, inducing posterior mean distribution H_i
- $(H, \cdots, H)$ is a symmetric equilibrium if no box profits from deviation

Second Main Result

Characterization of symmetric equilibrium for symmetric environment

- Each $r_i \sim G$ i.i.d.; equal cost $c_i = c$
- Box i reveals r_i partially, inducing posterior mean distribution H_i
- $(H, \cdots, H)$ is a symmetric equilibrium if no box profits from deviation

Thm 2. $(H, \cdots, H)$ is a symmetric equilibrium, if and only if

1) **Equal reservation value**: $\sigma_H = \sigma_G$
2) **Linear shape**: H is linear whenever not revealing full info
3) **No incentive to deviate** to no info or "approximately" full info

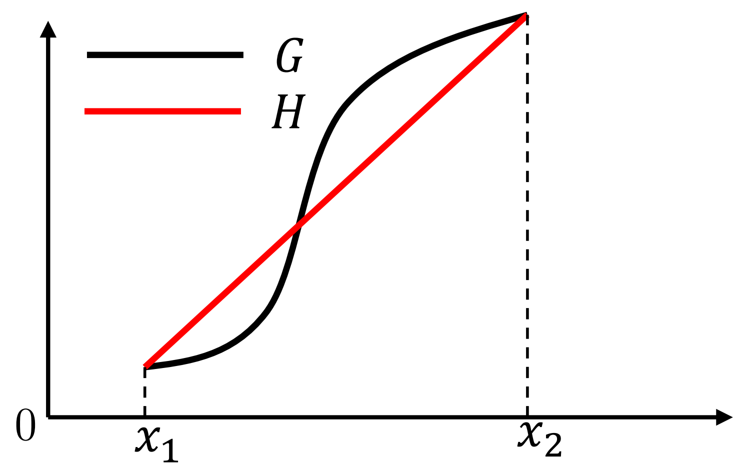

Second Main Result

Characterization of symmetric equilibrium for symmetric environment

- Each $r_i \sim G$ i.i.d.; equal cost $c_i = c$
- Box i reveals r_i partially, inducing posterior mean distribution H_i
- $(H, \cdots, H)$ is a symmetric equilibrium if no box profits from deviation

Thm 2. $(H, \cdots, H)$ is a symmetric equilibrium, if and only if

1) **Equal reservation value**: $\sigma_H = \sigma_G$
2) **Linear shape**: H is linear whenever not revealing full info
3) **No incentive to deviate** to no info or "approximately" full info

Remark.

- Strictly generalize [Au/Kawai, GEB'20; Hwang et al. 2019], which study special case with $c = 0$

Implications and Conceptual Messages

Corollary. Full information is the unique symmetric equilibrium if and only if $G^{n-1}(r)$ is convex over $[0, \sigma_G]$

Conceptual messages

- More competition → more transparency
  - $G^{n-1}(r)$ convex implies $G^{m-1}(r)$ convex for any $m \geq n$
- Larger inspection cost → more transparency
  - Larger c → smaller σ_G

Summary and Open Problem

- Many decision making problems involve costly information acquisition
  - Very often, information providers are strategic
  - Examples: secretary problem, prophet inequality, etc.
- This paper initiates an information design study for an elegant and one of the simplest problems in this space
- Open problems:
  - Equilibrium in asymmetric environment (cost and rewards)?
  - Informational aspects of many other classic sequential search problems

Thank You

Questions?

haifengxu@uchicago.edu