EC'23 Tutorials on Information Design

Part 2: (Some) New Frontiers

Joint Tutorial with Konstantin Zabarnyi

Haifeng Xu

Department of Computer Science
and Data Science Institute

University of Chicago

SIGMA LAB
Strategic IntelliGence in Machine Agents

Outline

- Info Design with Monetary Transfers – Pricing of Information
- Info Design in Optimal Stopping
- Info Design in Principal-Agent Problems
- Info Design without Commitment – Cheap Talk

> - Will focus more on problems/results, less on techniques
> - All relevant papers are listed at the bottom of each slide

Outline

- Info Design with Monetary Transfers – Pricing of Information
- Info Design in Optimal Stopping
- Info Design in Principal-Agent Problems
- Info Design without Commitment – Cheap Talk

- Will focus more on problems/results, less on techniques
- All relevant papers are listed at the bottom of each slide

Motivations of information pricing – how to sell ML predictions?

Google Cloud Vertex AI

Amazon SageMaker

Restaurant type and location →  → Demand

At high level: sender “persuades” a receiver to pay him, as opposed to take certain actions?

Monopoly Pricing of Information: A Basic Model

- One seller [sender], one buyer [receiver]
- Buyer is a decision maker who faces a binary choice: an active action 1 and a passive action 0
  - Active action: open a restaurant, approve loan, invest stock X, etc.
- Payoff of passive action $\equiv 0$

Open a restaurant here?

Google Cloud Vertex AI

(seller)

(buyer)

Monopoly Pricing of Information: A Basic Model

- One seller [sender], one buyer [receiver]
- Buyer is a decision maker who faces a binary choice: an active action 1 and a passive action 0
  - Active action: open a restaurant, approve loan, invest stock X, etc.
- Payoff of passive action $\equiv 0$
- Payoff of active action $= v(q, t)$
  - q is a *state of nature*, t is buyer type
- $t \sim F(t)$ is privy to buyer

$$v(q,t) = \underbrace{q}_{\text{demand}} \times \underbrace{t}_{\text{Profit per person}} - \underbrace{2}_{\text{operation cost}}$$

Google Cloud Vertex AI

(seller)

Open a restaurant here?

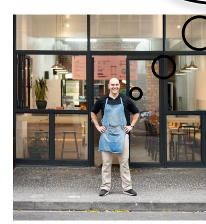

(buyer)

Monopoly Pricing of Information: A Basic Model

- One seller [sender], one buyer [receiver]
- Buyer is a decision maker who faces a binary choice: an active action 1 and a passive action 0
  - Active action: open a restaurant, approve loan, invest stock X, etc.
- Payoff of passive action $\equiv 0$
- Payoff of active action $= v(q, t)$
  - q is a *state of nature*, t is buyer type
- $\mathrm{t} \sim F(t)$ is privy to buyer
- $q \sim G(q)$ and seller reveals information about q

Open a restaurant here?

Google
Cloud
Vertex
AI

(seller)

Reveal learned information about state

(buyer)

Monopoly Pricing of Information: A Basic Model

- One seller [sender], one buyer [receiver]
- Buyer is a decision maker who faces a binary choice: an active action 1 and a passive action 0
  - Active action: open a restaurant, approve loan, invest stock X, etc.
- Payoff of passive action $\equiv 0$
- Payoff of active action $= v(q, t)$
  - q is a *state of nature*, t is buyer type
- $t \sim F(t)$ is privy to buyer
- $q \sim G(q)$ and seller reveals information about q

Mechanism design question: like persuasion, seller designs a signaling scheme about q, but to maximize charges from buyer ?

Selling Threshold Experiments Turns Out to Suffice

Recall buyer value $v(q, t)$

Def. A personalized threshold experiments (i.e., signaling scheme) is determined by some threshold function $\theta(t)$ – it simply predicts $q \geq \theta(t)$ or not for each buyer type t

Characterizing the Optimal Mechanism

Def. Lower virtual value function: $\underline{\phi}(t) = t - \frac{1-F(t)}{f(t)}$ [Myerson'81]

Upper virtual value function: $\bar{\phi}(t) = t + \frac{F(t)}{f(t)}$

Mixed virtual value function: $\phi_c(t) = c\underline{\phi}(t) + (1-c)\bar{\phi}(t)$

Characterizing the Optimal Mechanism

Def. Lower virtual value function: $\underline{\phi}(t) = t - \frac{1-F(t)}{f(t)}$ [Myerson'81]

Upper virtual value function: $\bar{\phi}(t) = t + \frac{F(t)}{f(t)}$

Mixed virtual value function: $\phi_c(t) = c\underline{\phi}(t) + (1-c)\bar{\phi}(t)$

Theorem (informal, [LSX', EC21]).

The mechanism with threshold predictions $\theta^*(t) = -\phi_c(t)$ and following payment function represents an optimal mechanism to previous problem:

$$p^*(t) = \int_{q\in Q} \pi^*(q,t)g(q)v(q,t)\mathrm{d}q - \int_{t_1}^{t}\int_{q\in Q} \pi^*(q,x)g(q)v_1(q)\mathrm{d}q\ \mathrm{d}x$$

where $\phi_c(t)$ is the "mixed virtual value" function and c is determined by certain equation.

Characterizing the Optimal Mechanism

Economic insights:

- ✓ Optimal mechanism sells processed information, not the q itself
- ✓ Personalized threshold for different user types (like private persuasion [AB, JET'19])
- ✓ Consequently, much more power for price discrimination

Theorem (informal, [LSX', EC21]).

The mechanism with threshold predictions $\theta^*(t) = -\phi_c(t)$ and following payment function represents an optimal mechanism to previous problem:

$$p^*(t) = \int_{q\in Q} \pi^*(q,t)g(q)v(q,t)\mathrm{d}q - \int_{t_1}^{t}\int_{q\in Q} \pi^*(q,x)g(q)v_1(q)\mathrm{d}q \ \mathrm{d}x$$

where $\phi_c(t)$ is the "mixed virtual value" function and c is determined by certain equation.

Characterizing the Optimal Mechanism

Deriving closed-form optimal mechanisms for more general models appears much less tractable…

Theorem (informal, [LSX', EC21]).

The mechanism with threshold predictions $\theta^*(t) = -\phi_c(t)$ and following payment function represents an optimal mechanism to previous problem:

$$p^*(t) = \int_{q \in Q} \pi^*(q,t) g(q) v(q,t) \mathrm{d}q - \int_{t_1}^{t} \int_{q \in Q} \pi^*(q,x) g(q) v_1(q) \mathrm{d}q \ \mathrm{d}x$$

where $\phi_c(t)$ is the "mixed virtual value" function and c is determined by certain equation.

SIGMA LAB
Strategic IntelliGence in Machine Agents

A General Model of Information Pricing

- Buyer takes one of n action $i \in [n] = \{1, \cdots, n\}$
- Buyer has an arbitrary utility function $u(i, \theta; t)$ where
  - $\theta \sim dist. \mu$ is a random state of nature
  - $t \sim dist. F$ captures

✓ [BKP, EC'12] developed revelation principle, and a complex polynomial-time algorithm

Optimal Mechanisms for Selling Information. Moshe Babaioff, Robert Kleinberg and Renato Paes Leme, EC'12

A General Model of Information Pricing

- Buyer takes one of n action $i \in [n] = \{1, \cdots, n\}$
- Buyer has an arbitrary utility function $u(i, \theta; t)$ where
  - $\theta \sim dist. \mu$ is a random state of nature
  - $t \sim dist. F$ captures

✓ [BKP, EC'12] developed revelation principle, and a complex polynomial-time algorithm

- Computation is extremely complex (despite poly time) and has to go through every vertex of buyer's posterior polytope
- Optimal mechanisms may be unrealistic
  - E.g., there are examples for which buyer value is in [0, 5]
  - Optimal mechanism asks buyer to deposit $25004 first
  - Then return either 0 or 50000, yielding optimal revenue < 2

Optimal Mechanisms for Selling Information. Moshe Babaioff, Robert Kleinberg and Renato Paes Leme, EC'12

A General Model of Information Pricing

- Buyer takes one of n action $i \in [n] = \{1, \cdots, n\}$
- Buyer has an arbitrary utility function $u(i, \theta; t)$ where
  - $\theta \sim dist. \mu$ is a random state of nature
  - $t \sim dist. F$ captures

✓ [BKP, EC'12] developed revelation principle, and a complex polynomial-time algorithm

✓ [CXZ, SODA'20] significantly simplifies their algorithm to a single convex program, and allows payment constraints

Key idea: a simplified revelation principle that shows the existence of an optimal mechanism in succinct format

The Optimal Mechanism Proceeds Like Consulting

The Consulting Mechanism [CXZ, SODA'20]

1. Elicit buyer type t
2. Charge buyer x_t
3. Implement signaling scheme π_t for buyer t – recommend action i to the buyer with prob $\pi_t(\sigma_i, \theta)$ on state θ

- A consulting mechanism is described by a (payment, signaling scheme) menu $\{x_t, \pi_t\}_t$

The Optimal Mechanism Proceeds Like Consulting

The Consulting Mechanism [CXZ, SODA'20]

1. Elicit buyer type t
2. Charge buyer x_t
3. Implement signaling scheme π_t for buyer t – recommend action i to the buyer with prob $\pi_t(\sigma_i, \theta)$ on state θ

- A consulting mechanism is described by a (payment, signaling scheme) menu $\{x_t, \pi_t\}_t$
- Will be incentive compatible – reporting true t is optimal
- The recommended action is obedient – guaranteed to be the optimal action for buyer t given his information

Thm. The optimal consulting mechanism with menu $\{x_t, \pi_t\}_t$ computed by the following convex program is an optimal mechanism.

Computing the Optimal Mechanism

The convex program for computing optimal $\{x_t, \pi_t\}_t$

- Variables: $\pi_t(\sigma_i, \theta)$ = prob of sending σ_i conditioned on θ for each t
- Variable x_t is the payment from buyer type t

Expected revenue

$$
\begin{aligned}
\max \quad & \boxed{\textstyle\sum_t f(t) \cdot x_t} & \\
\text{s.t.} \quad & \textstyle\sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t & \\
& \quad \geq \textstyle\sum_i \max_j \left[\sum_\theta \mu(\theta) \pi_{t'}(\sigma_i, \theta) u(j, \theta; t) \right] - x_{t'}, & \text{for } t' \neq t \\
& \textstyle\sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t \geq \max_i \sum_\theta \mu(\theta) u(i, \theta; t), & \text{for } t \\
& \textstyle\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \geq \sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(j, \theta; t), & \text{for } i \neq j, t \\
& \textstyle\sum_i \pi_t(\sigma_i, \theta) = 1, & \text{for } \theta, t \\
& \pi_t(\sigma_i, \theta) \geq 0, & \text{for } t, \sigma_i, \theta
\end{aligned}
$$

Computing the Optimal Mechanism

The convex program for computing optimal $\{x_t, \pi_t\}_t$

- Variables: $\pi_t(\sigma_i, \theta)$ = prob of sending σ_i conditioned on θ for each t
- Variable x_t is the payment from buyer type t

Truthfully reporting true t is optimal

$$
\begin{array}{lll}
\max & \sum_t f(t) \cdot x_t & \\
\text{s.t.} & \boxed{\begin{array}{l} \sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t \\ \qquad \ge \sum_i \max_j \left[\sum_\theta \mu(\theta) \pi_{t'}(\sigma_i, \theta) u(j, \theta; t) \right] - x_{t'}, \end{array}} & \text{for } t' \neq t \\
& \sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t \ge \max_i \sum_\theta \mu(\theta) u(i, \theta; t), & \text{for } t \\
& \sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \ge \sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(j, \theta; t), & \text{for } i \neq j, t \\
& \sum_i \pi_t(\sigma_i, \theta) = 1, & \text{for } \theta, t \\
& \pi_t(\sigma_i, \theta) \ge 0, & \text{for } t, \sigma_i, \theta
\end{array}
$$

Computing the Optimal Mechanism

The convex program for computing optimal $\{x_t, \pi_t\}_t$

- Variables: $\pi_t(\sigma_i, \theta)$ = prob of sending σ_i conditioned on θ for each t
- Variable x_t is the payment from buyer type t

Participation is no worse than not

$$
\begin{array}{lll}
\max & \sum_t f(t) \cdot x_t & \\
\text{s.t.} & \sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t & \\
& \quad \geq \sum_i \max_j \left[\sum_\theta \mu(\theta) \pi_{t'}(\sigma_i, \theta) u(j, \theta; t) \right] - x_{t'}, & \text{for } t' \neq t \\
& \boxed{\sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t \geq \max_i \sum_\theta \mu(\theta) u(i, \theta; t),} & \text{for } t \\
& \sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \geq \sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(j, \theta; t), & \text{for } i \neq j, t \\
& \sum_i \pi_t(\sigma_i, \theta) = 1, & \text{for } \theta, t \\
& \pi_t(\sigma_i, \theta) \geq 0, & \text{for } t, \sigma_i, \theta
\end{array}
$$

Computing the Optimal Mechanism

The convex program for computing optimal $\{x_t, \pi_t\}_t$

- Variables: $\pi_t(\sigma_i, \theta)$ = prob of sending σ_i conditioned on θ for each t
- Variable x_t is the payment from buyer type t

$$
\begin{array}{lll}
\max & \sum_t f(t) \cdot x_t & \\
\text{s.t.} & \sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t & \\
& \quad \geq \sum_i \max_j \left[\sum_\theta \mu(\theta) \pi_{t'}(\sigma_i, \theta) u(j, \theta; t) \right] - x_{t'}, & \text{for } t' \neq t \\
& \sum_i \left[\sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \right] - x_t \geq \max_i \sum_\theta \mu(\theta) u(i, \theta; t), & \text{for } t \\
& \sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(i, \theta; t) \geq \sum_\theta \mu(\theta) \pi_t(\sigma_i, \theta) u(j, \theta; t), & \text{for } i \neq j, t \\
& \sum_i \pi_t(\sigma_i, \theta) = 1, & \text{for } \theta, t \\
& \pi_t(\sigma_i, \theta) \geq 0, & \text{for } t, \sigma_i, \theta
\end{array}
$$

Similar to constraints in persuasion

(Some) Other Models of Information Pricing

- ✓ [BBS, AER'18] considers analogous monopoly pricing problem, but buyer's type is about his belief of the state, not his utility uncertainty

The Design and Price of Information. Dirk Bergemann, Alessandro Bonatti, Alex Smolin, AER'18.

(Some) Other Models of Information Pricing

- ✓ [BBS, AER’18] considers analogous monopoly pricing problem, but buyer’s type is about his belief of the state, not his utility uncertainty
- ✓ [CV, ITCS’21] develops algorithms for the above model

How to sell information optimally: An algorithmic study. Yang Cai, Grigoris Velegkas. ITCS’21

(Some) Other Models of Information Pricing

- ✓ [BBS, AER'18] considers analogous monopoly pricing problem, but buyer's type is about his belief of the state, not his utility uncertainty
- ✓ [CV, ITCS'21] develops algorithms for the above model
- ✓ [BCVZ, EC'22] characterizes how much revenue loss the seller incurs without using price discrimination

(Some) Other Models of Information Pricing

- ✓ [BBS, AER'18] considers analogous monopoly pricing problem, but buyer's type is about his belief of the state, not his utility uncertainty
- ✓ [CV, ITCS'21] develops algorithms for the above model
- ✓ [BCVZ, EC'22] characterizes how much revenue loss the seller incurs without using price discrimination
- ✓ [Li, EC'22] studies what happens if buyer can obtain additional signals outside

(Some) Other Models of Information Pricing

- ✓ [BBS, AER'18] considers analogous monopoly pricing problem, but buyer's type is about his belief of the state, not his utility uncertainty
- ✓ [CV, ITCS'21] develops algorithms for the above model
- ✓ [BCVZ, EC'22] characterizes how much revenue loss the seller incurs without using price discrimination
- ✓ [Li, EC'22] studies what happens if buyer can obtain additional signals outside
- ✓ [Yang, AER'22] studies how to sell information about market segments to a producer for further the producer's revenue via price discrimination

Selling consumer data for profit: Optimal market-segmentation design and its consequences Kai Hao Yang - American Economic Review, 2022

(Some) Other Models of Information Pricing

- ✓ [BBS, AER'18] considers analogous monopoly pricing problem, but buyer's type is about his belief of the state, not his utility uncertainty
- ✓ [CV, ITCS'21] develops algorithms for the above model
- ✓ [BCVZ, EC'22] characterizes how much revenue loss the seller incurs without using price discrimination
- ✓ [Li, EC'22] studies what happens if buyer can obtain additional signals outside
- ✓ [Yang, AER'22] studies how to sell information about market segments to a producer for further the producer's revenue via price discrimination
- ✓ [BDHN, working paper'22] studies selling information to multiple and competitive players.

Selling Information in Competitive Environments. Alessandro Bonatti, Munther Dahleh, Thibaut Horel, and Amir Nouripour, working paper.

(Some) Other Models of Information Pricing

- ✓ [BBS, AER'18] considers analogous monopoly pricing problem, but buyer's type is about his belief of the state, not his utility uncertainty
- ✓ [CV, ITCS'21] develops algorithms for the above model
- ✓ [BCVZ, EC'22] characterizes how much revenue loss the seller incurs without using price discrimination
- ✓ [Li, EC'22] studies what happens if buyer can obtain additional signals outside
- ✓ [Yang, AER'22] studies how to sell information about market segments to a producer for further the producer's revenue via price discrimination
- ✓ [BDHN, working paper'22] studies selling information to multiple and competitive players.
- ✓ [BB, ARE'19] gives an excellent survey from the economic perspective about markets for information
    - Seems to lack survey from algorithmic/operational perspective…

Markets for information: An introduction. Dirk Bergemann and Alessandro Bonatti, Annual Review of Economics.

Outline

- Info Design with Monetary Transfers – Pricing of Information
- Info Design in Optimal Stopping
- Info Design in Principal-Agent Problems
- Info Design without Commitment – Cheap Talk

An Example: Pandora's Box [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

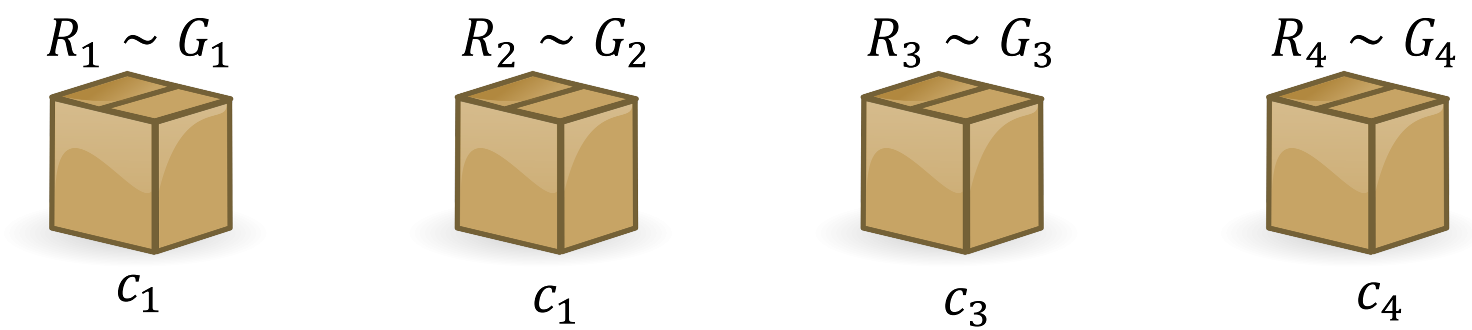

❖ Numerous applications: look for startups to fund, open house, find channels to subscribe

An Example: Pandora's Box [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?

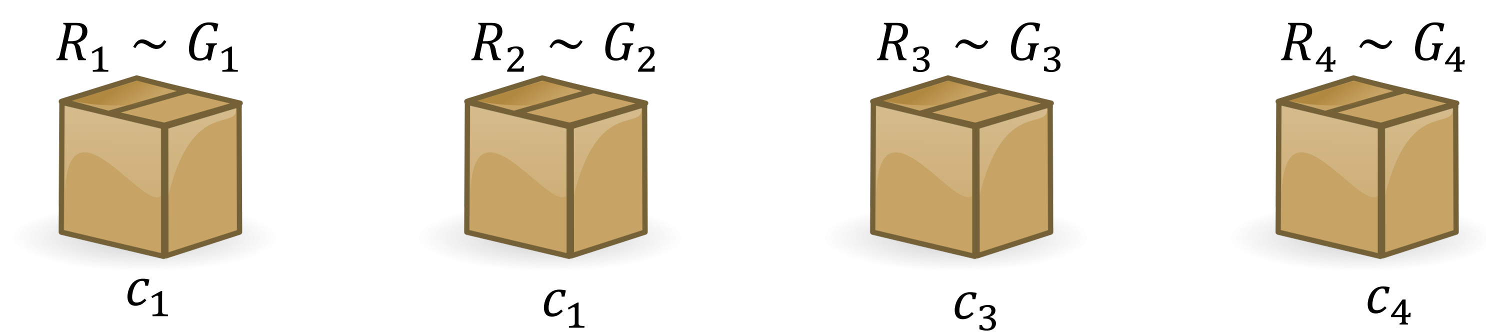

❖ Numerou  ons: look for startups to fund, open house, find chan scribe

An Example: Pandora's Box [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?

$R_1 \sim G_1$ $\quad$ $R_2 \sim G_2$ $\quad$ $R_3 \sim G_3$ $\quad$ $R_4 \sim G_4$

✓ There is an elegant greedy policy that is optimal for this problem

An Example: Pandora's Box [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?

$R_1 \sim G_1$ $\quad$ $R_2 \sim G_2$ $\quad$ $R_3 \sim G_3$ $\quad$ $R_4 \sim G_4$

- ✓ There is an elegant greedy policy that is optimal for this problem
- ✓ Many other optimal stopping problem has similar structure, but have different reward selection criteria and box order constraints

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

OpenAI SPACEX ByteDance FTX

Venture capital searches for a good startup to invest

Pandora’s Box with Strategic Boxes

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

c_1 $\quad$ c_1 $\quad$ c_3 $\quad$ c_4

Search for a PhD admission through their open houses

Pandora's Box with Strategic Boxes

$r_1 \sim G_1$

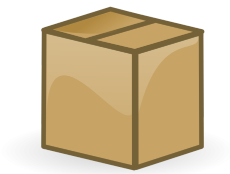

c_1

$r_2 \sim G_2$

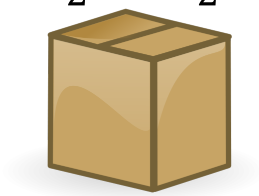

c_1

$r_3 \sim G_3$

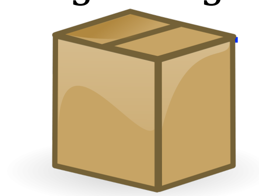

c_3

$r_4 \sim G_4$

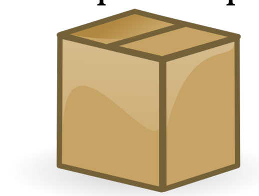

c_4

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

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

c_1 c_1 c_3 c_4

Competitive Information Design [DFHTX, SODA'23]

- ➢ 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

Competitive Information Design for Pandora's Box. Bolin Ding, Yiding Feng, Chien-Ju Ho, Wei Tang and Haifeng Xu, SODA 2023

Main Results

Result 1 [Information Order in Pandora's Box]. Let $U(G_i, G_{-i})$ denote agent's optimal utility in Pandora's Box. Then G_i is more informative than 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.

Main Results

Result 1 [Information Order in Pandora's Box]. Let $U(G_i, G_{-i})$ denote agent's optimal utility in Pandora's Box. Then G_i is more informative than H, if and only if

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

Result 2. Fully characterizes symmetric equilibria in symmetric environments

- Equilibrium characterization reveals conceptual messages about transparency in Pandora's Box
  - More competition → more transparency
  - Larger inspection cost → more transparency
- Strictly generalize [Au/Kawai, GEB'20; Hwang et al. 2019], which study special case with 0 search cost.

Open Directions

- Many decision-making/searching problems involve costly information acquisition
  - Very often, information providers are strategic
  - Examples: secretary problem, option trading, house selling, parking, etc.
- Many open problems:
  - Immediate: equilibrium in asymmetric environment (cost and rewards)?
  - Generally: Informational design in many other optimal stopping problems

Outline

- Info Design with Monetary Transfers – Pricing of Information
- Info Design in Optimal Stopping
- Info Design in Principal-Agent Problems
- Info Design without Commitment – Cheap Talk

Standard Principal-Agent Setups

Standard Principal-Agent Setups

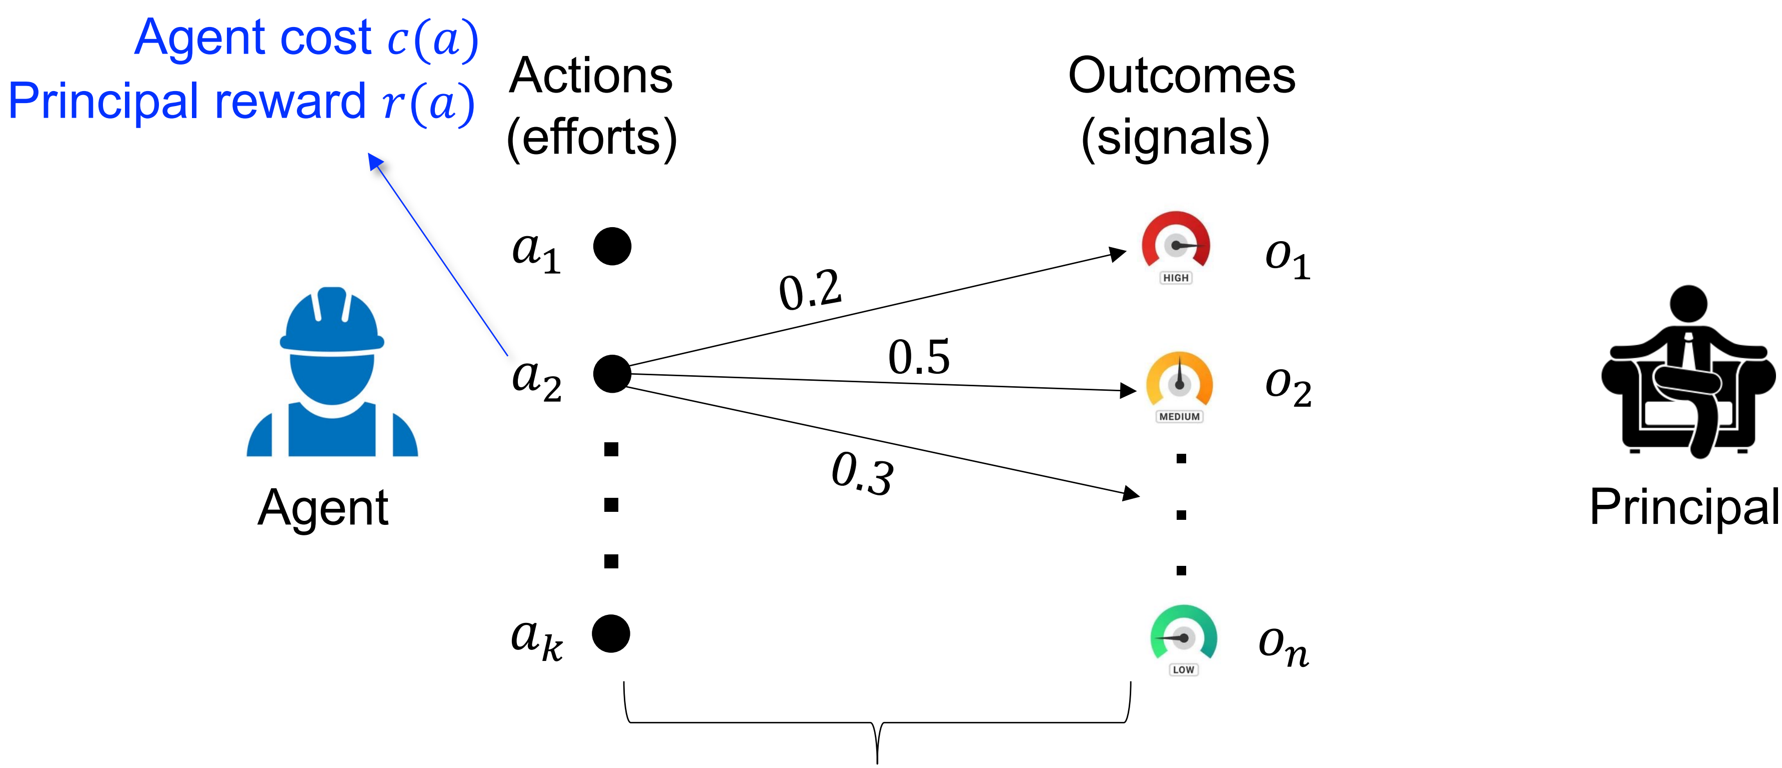

A stochastic mapping that determines how much information the observed outcome carries about agent’s underlying action

This is a signaling scheme (a.k.a., “monitoring technology”)

Standard Principal-Agent Setups

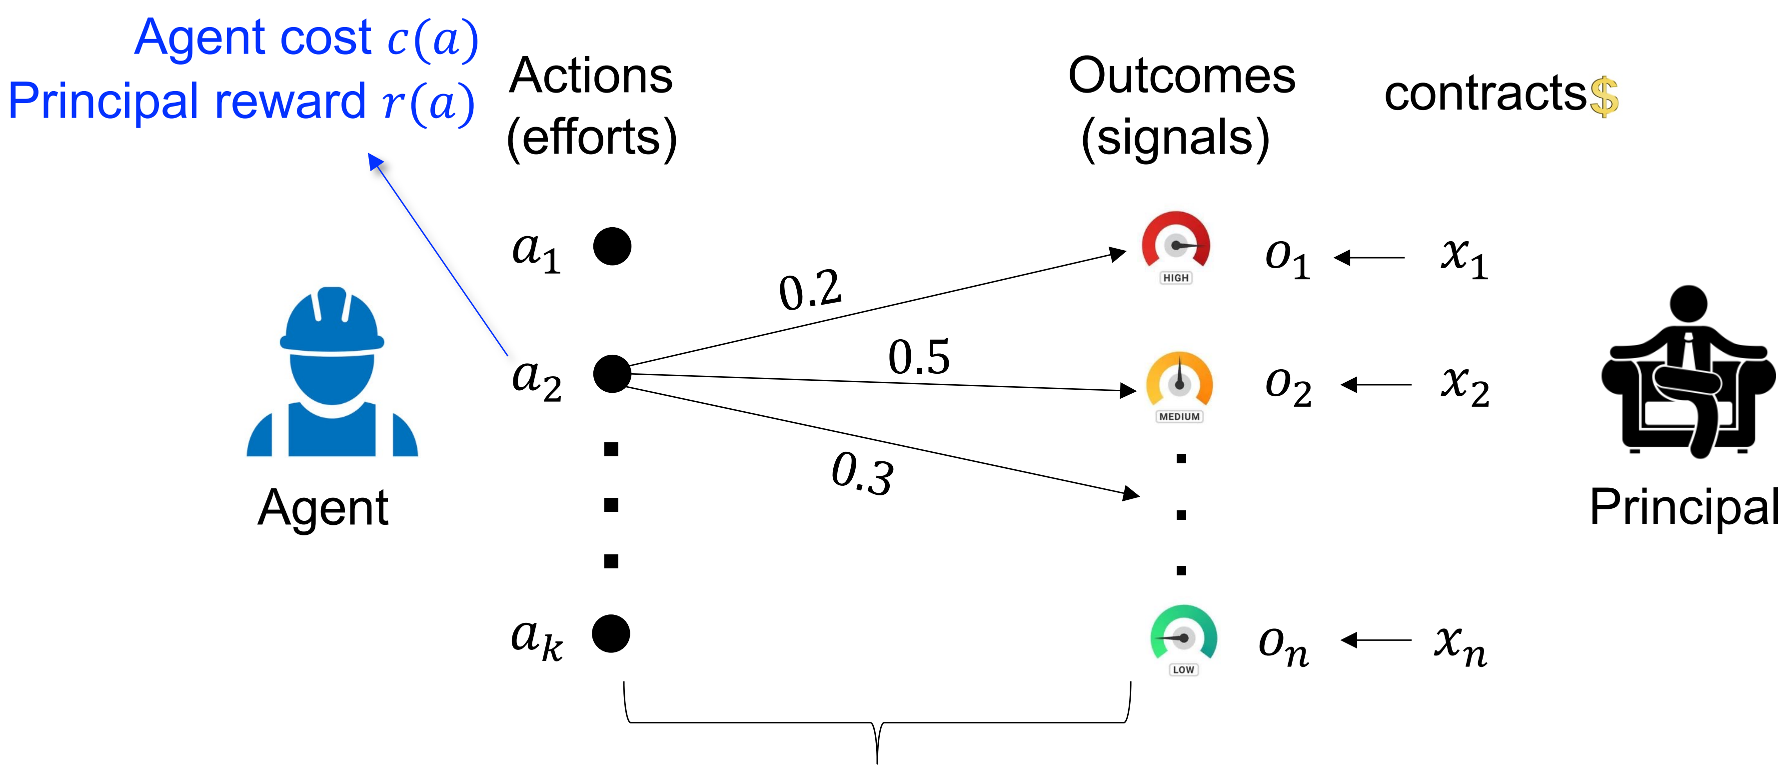

A stochastic mapping that determines how much information the observed outcome carries about agent's underlying action

This is a signaling scheme (a.k.a., "monitoring technology")

Standard Principal-Agent Setups

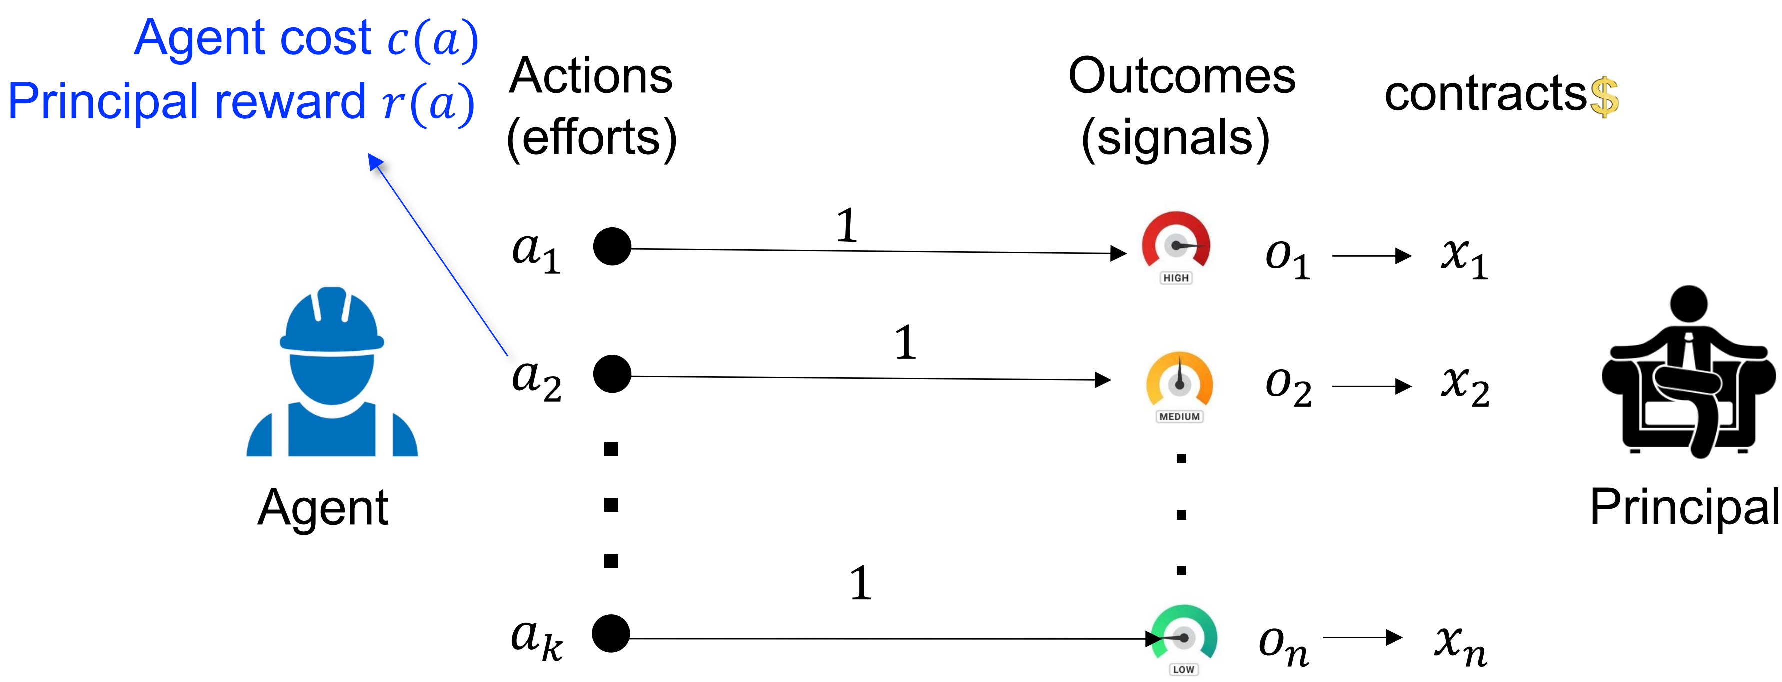

Q: What happens if fully informative?

- Optimal contract sets $x_i = \text{cost}(a_i)$
- Agent gets 0 surplus

Standard Principal-Agent Setups

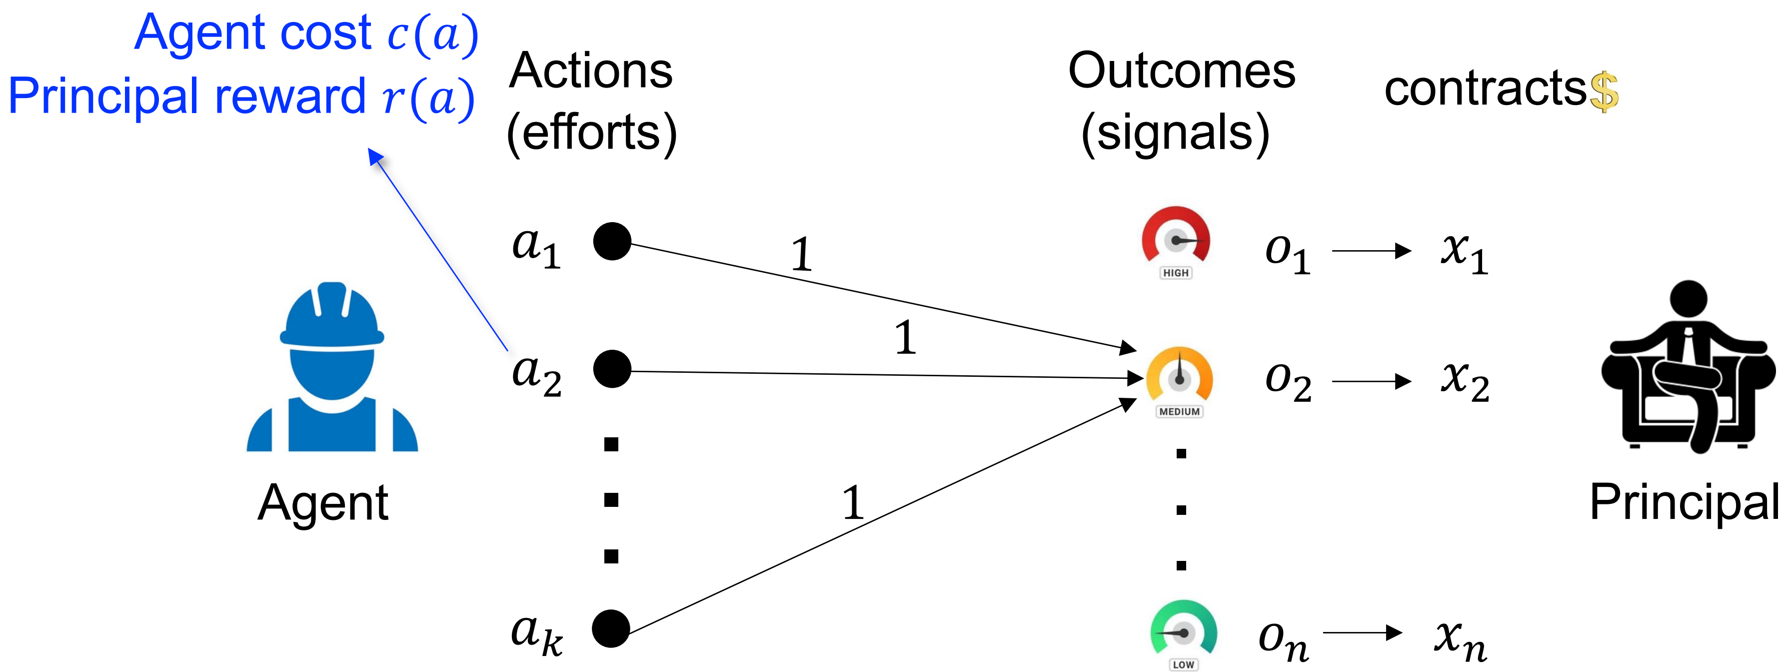

Q: What about completely non-informative scheme?

- Agent will take least-cost action, due to indistinguishable outcomes
- Contract = least action cost
- Potentially inefficient social outcome

Standard Principal-Agent Setups

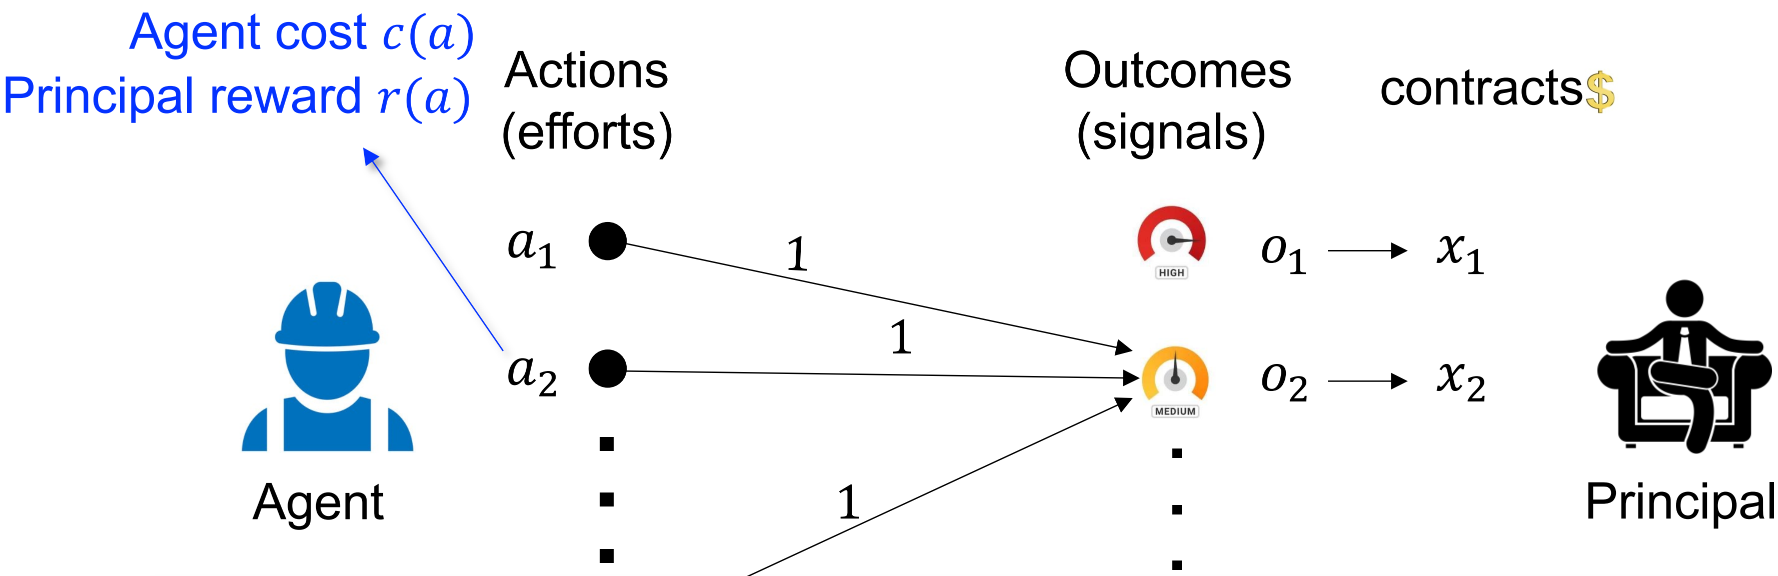

Lessons learned: This information structure affects total welfare, as well as what fraction of the welfare each player can get

Standard Principal-Agent Setups

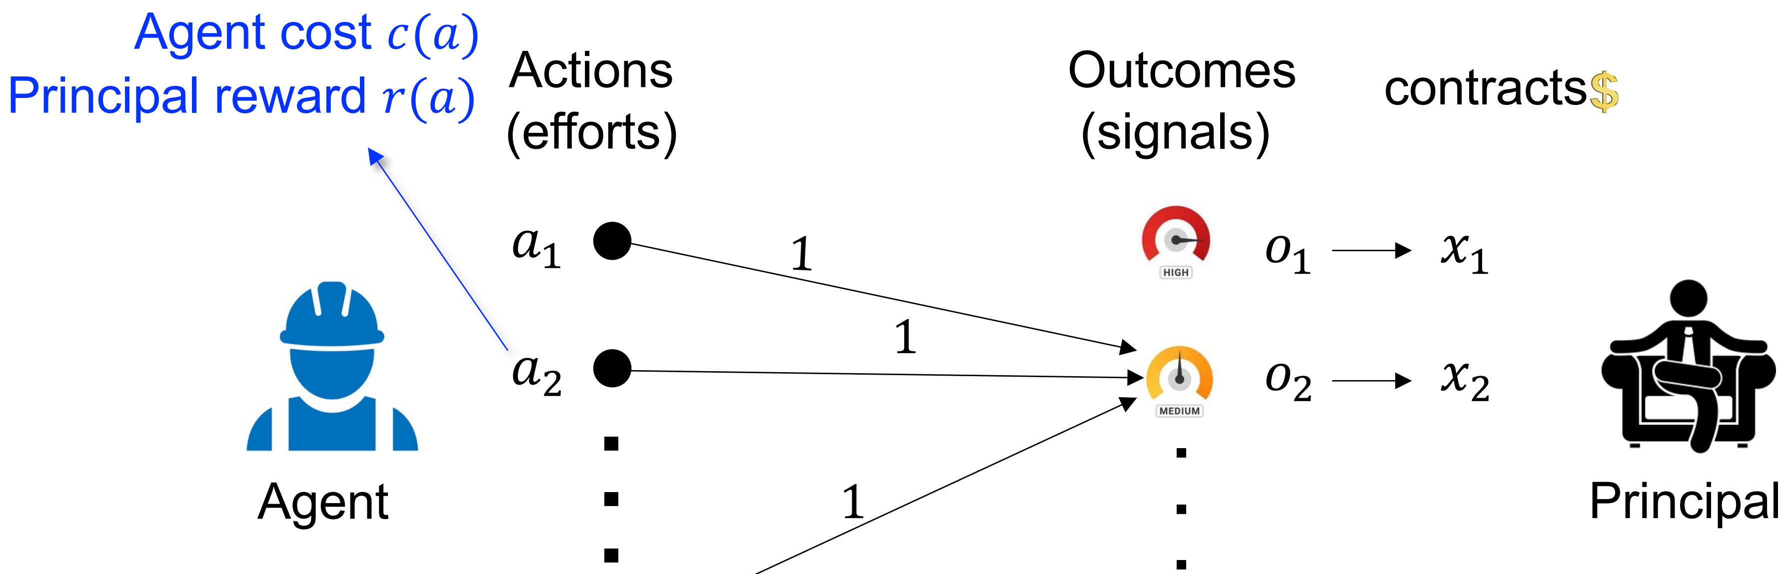

Lessons learned: This information structure affects total welfare, as well as what fraction of the welfare each player can get

- Most previous works assume information structure is fixed exogenously
- In many applications, a planner/regulator can design it!
  - Company monitoring policy
  - Freelancing worker-task matching platforms (e.g., Upwork)

Standard Principal-Agent Setups

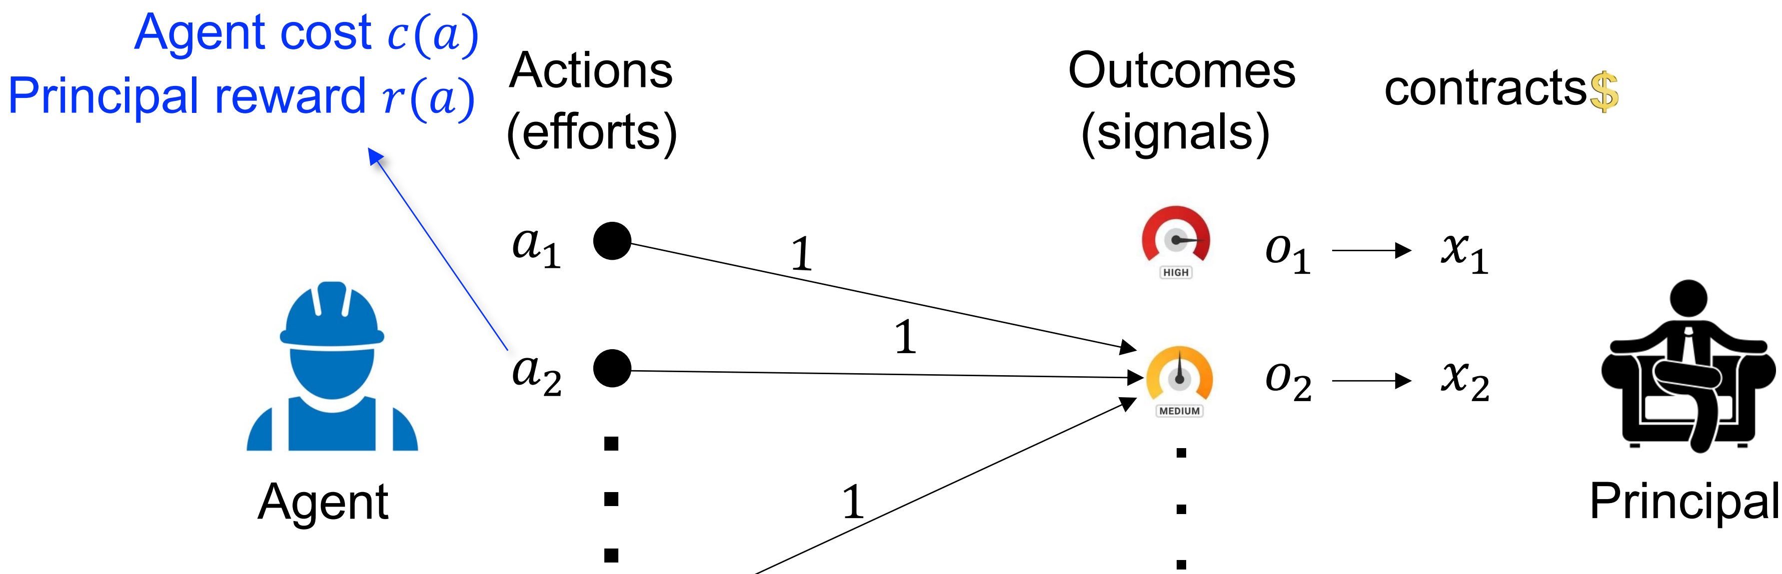

Lessons learned: This information structure affects total welfare, as well as what fraction of the welfare each player can get

Q1: understand how info structure affects social outcome

Q2: optimize info structure for designer objective under constraints

Standard Principal-Agent Setups

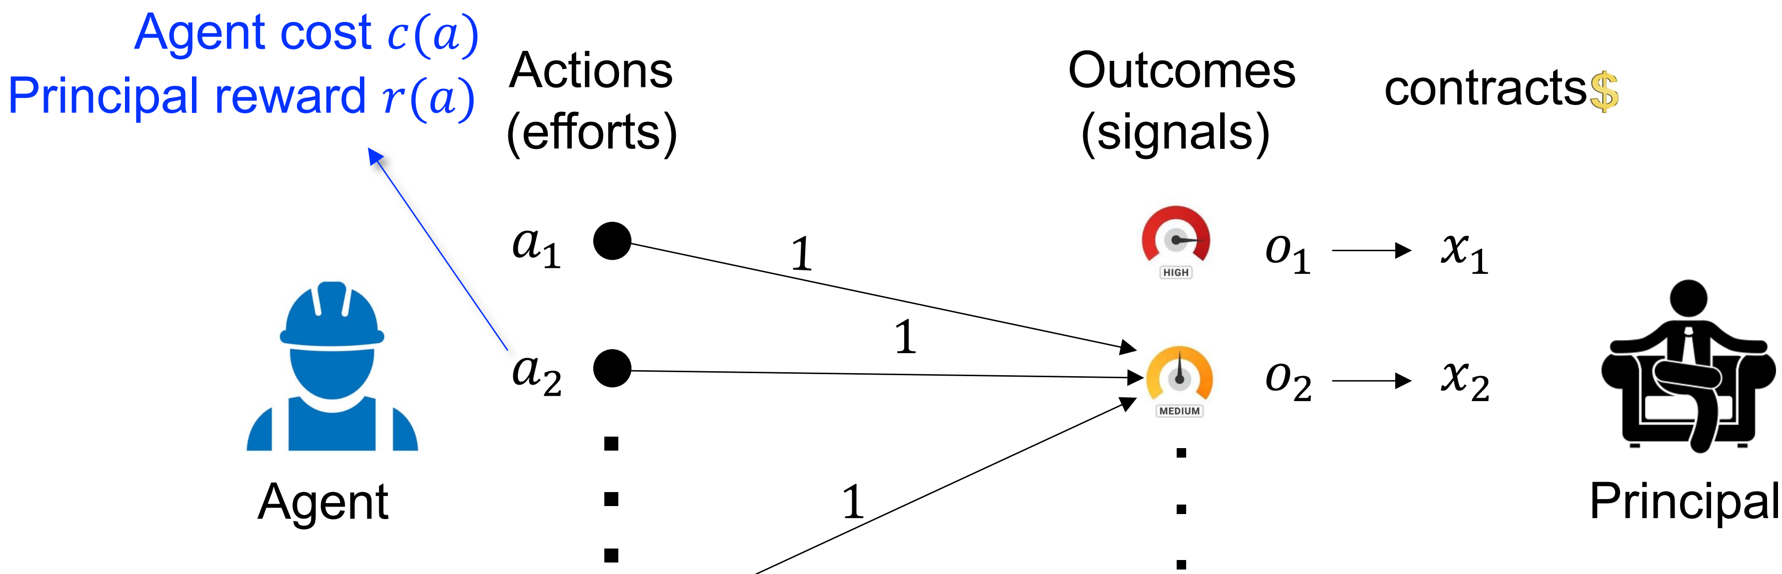

Lessons learned: This information structure affects total welfare, as well as what fraction of the welfare each player can get

Q1: understand how info structure affects social outcome [BTXZ'23]

Q2: optimize info structure for designer objective under constraints

Information Design in the Principal-Agent Problem. Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, Konstantin Zabarnyi, working paper 2023

Standard Principal-Agent Setups

Q1: characterize how info structure affects social outcome [BTXZ'23]

- Similar in spirit to a seminal work by Bergemann/Brooks/Morris [AER'15] on "*The Limits of Price Discrimination*" (adverse selection vs moral hazard)

Information Design in the Principal-Agent Problem. Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, Konstantin Zabarnyi, working paper 2023

Standard Principal-Agent Setups

Q1: characterize how info structure affects social outcome [BTXZ'23]

- Similar in spirit to a seminal work by Bergemann/Brooks/Morris [AER'15] on "*The Limits of Price Discrimination*" (adverse selection vs moral hazard)
- Here, we can fully characterize what agent action and (principal, agent) utility pairs are inducible via information design
  - Can account for risk-neural or risk-averse agents
  - Can account for some natural constraints on information structures
- Similar questions are studied in previous Econ literature, though in different models with different focuses
  - See more discussions in [BTXZ'23] and many references therein

Information Design in the Principal-Agent Problem. Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, Konstantin Zabarnyi, working paper 2023

Standard Principal-Agent Setups

Q1: characterize how info structure affects social outcome [BTXZ’23]

- Similar in spirit to a seminal work by Bergemann/Brooks/Morris [AER’15] on “*The Limits of Price Discrimination*” (adverse selection vs moral hazard)

- Here, we can fully characterize what agent action and (principal, agent) utility pairs are inducible via information design
  - Can account for risk-neural or risk-averse agents
  - Can account for some natural constraints on information structures

- Similar questions are studied in previous Econ literature, though in different models with different focuses
  - See more discussions in [BTXZ’23] and many references therein

- **Q2** about optimizing information structure is mostly open
  - For example, guarantee ”fair” welfare share, or maximize weighted combination of principal and agent utilities
  - Account for design constraints on info structures

Information Design in the Principal-Agent Problem. Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, Konstantin Zabarnyi, working paper 2023

Outline

- Info Design with Monetary Transfers – Pricing of Information
- Info Design in Optimal Stopping
- Info Design in Principal-Agent Problems
- Info Design without Commitment – Cheap Talk

Though commitment is natural in many applications, it can be unrealistic in others (e.g., security/warship, bargaining)

- Cheap talk – information design without commitment [Crawford/Sobel'82]

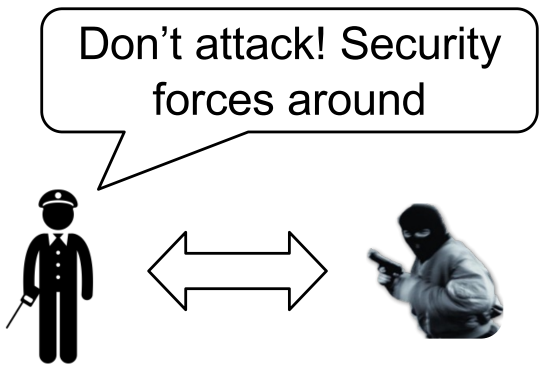

Strategic Information Transmission, Vincent P. Crawford, Joe Sobel, Econometrica 1982.

Though commitment is natural in many applications, it can be unrealistic in others (e.g., security/warship, bargaining)

- Cheap talk – information design without commitment [Crawford/Sobel'82]

Word state

θ_1

Sender

θ_2

$\vdots$

θ_k

actions

a_1

a_2

$\vdots$

a_n

Receiver

Strategic Information Transmission, Vincent P. Crawford, Joe Sobel, Econometrica 1982.

Though commitment is natural in many applications, it can be unrealistic in others (e.g., security/warship, bargaining)

- Cheap talk – information design without commitment [Crawford/Sobel'82]

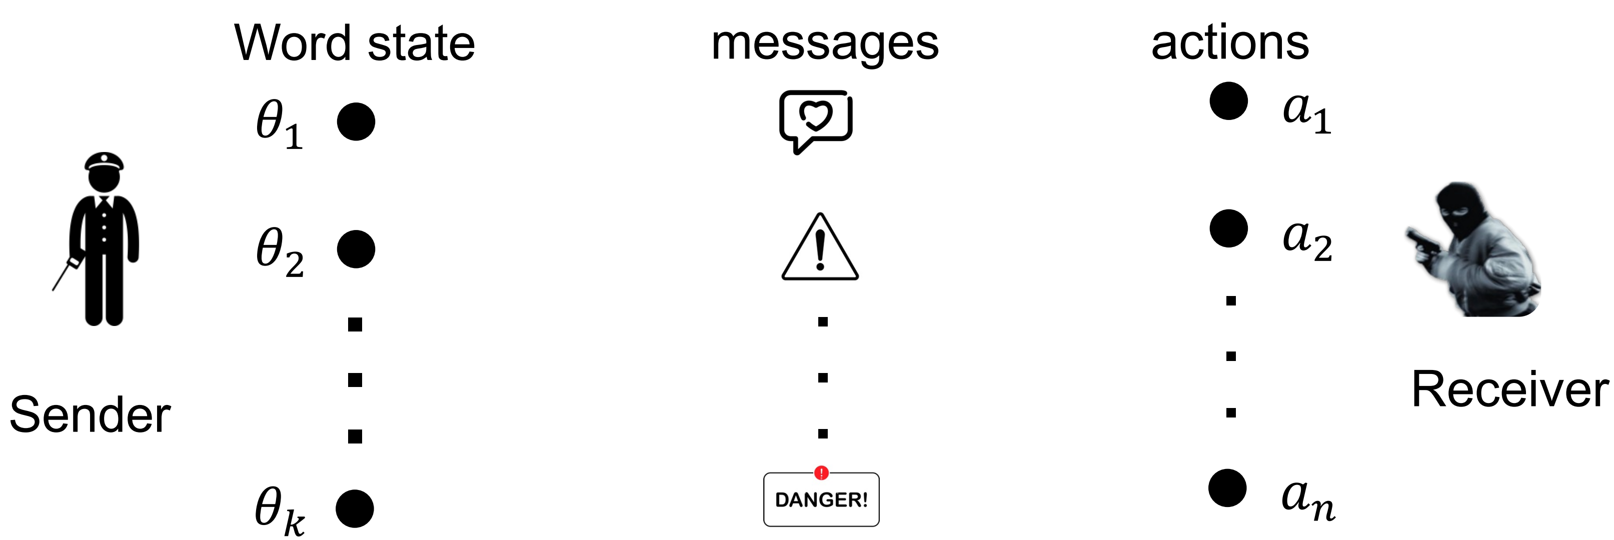

Strategic Information Transmission, Vincent P. Crawford, Joe Sobel, Econometrica 1982.

Though commitment is natural in many applications, it can be unrealistic in others (e.g., security/warship, bargaining)

- Cheap talk – information design without commitment [Crawford/Sobel'82]

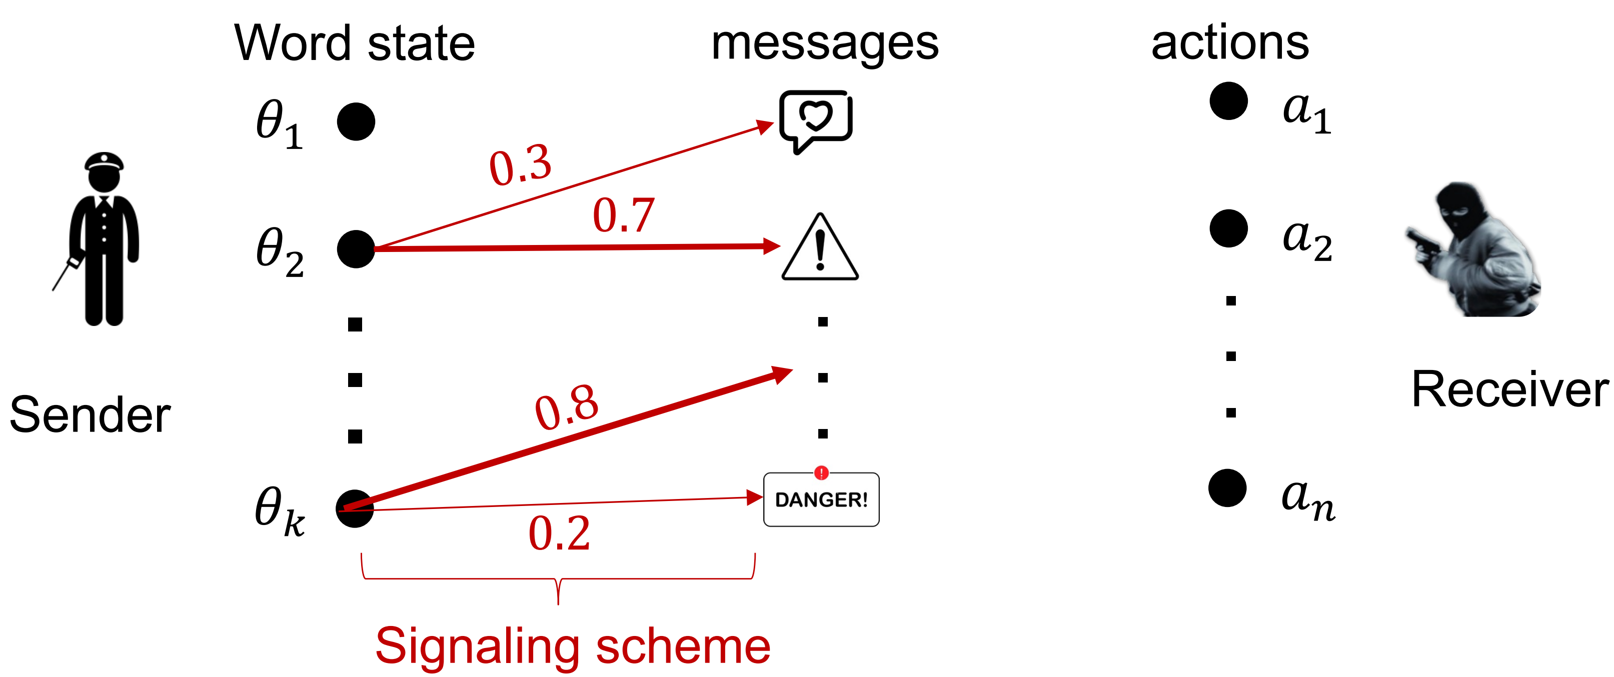

Strategic Information Transmission, Vincent P. Crawford, Joe Sobel, Econometrica 1982.

Though commitment is natural in many applications, it can be unrealistic in others (e.g., security/warship, bargaining)

- Cheap talk – information design without commitment [Crawford/Sobel'82]

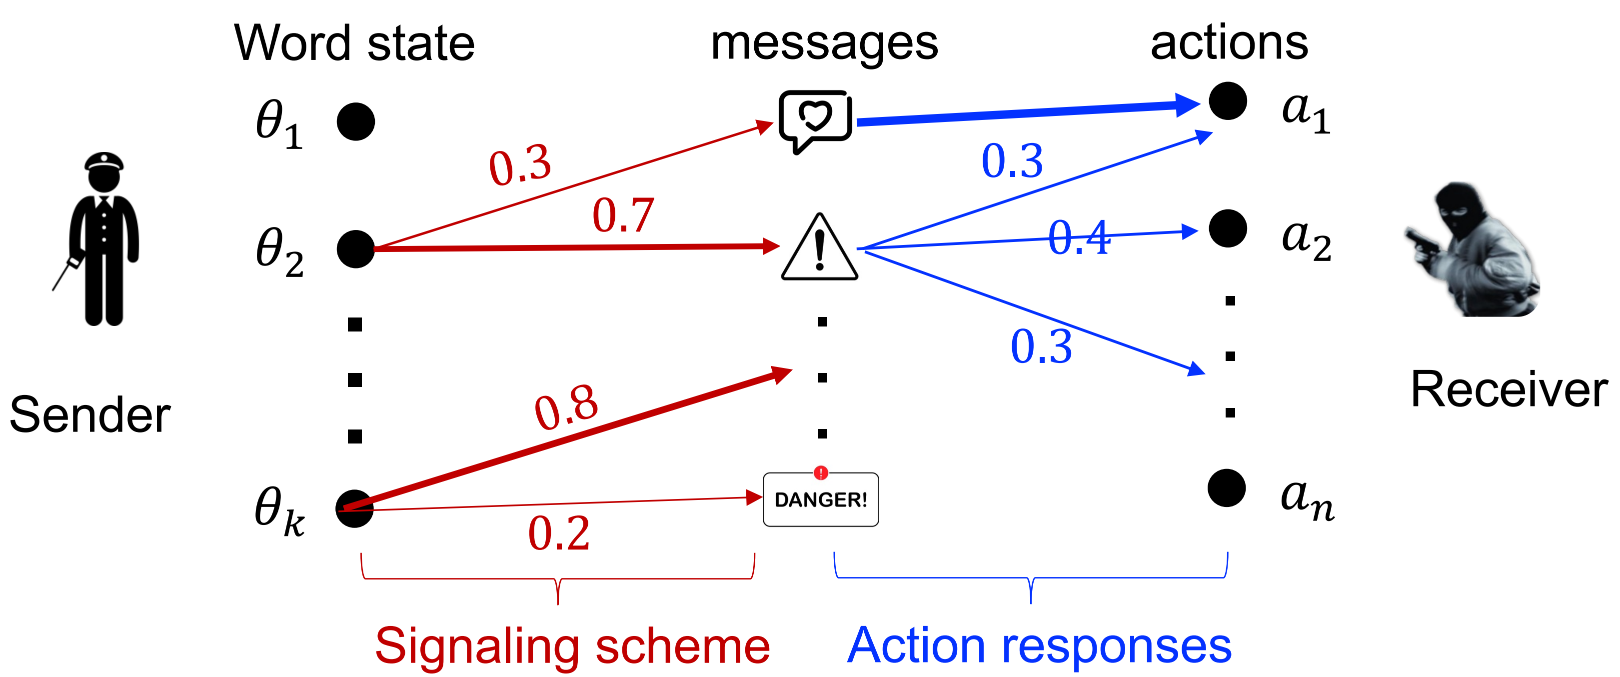

Strategic Information Transmission, Vincent P. Crawford, Joe Sobel, Econometrica 1982.

Though commitment is natural in many applications, it can be unrealistic in others (e.g., security/warship, bargaining)

- Cheap talk – information design without commitment [Crawford/Sobel'82]
- (signaling scheme, action response) forms a cheap talk equilibrium if no player wants to unilaterally deviate

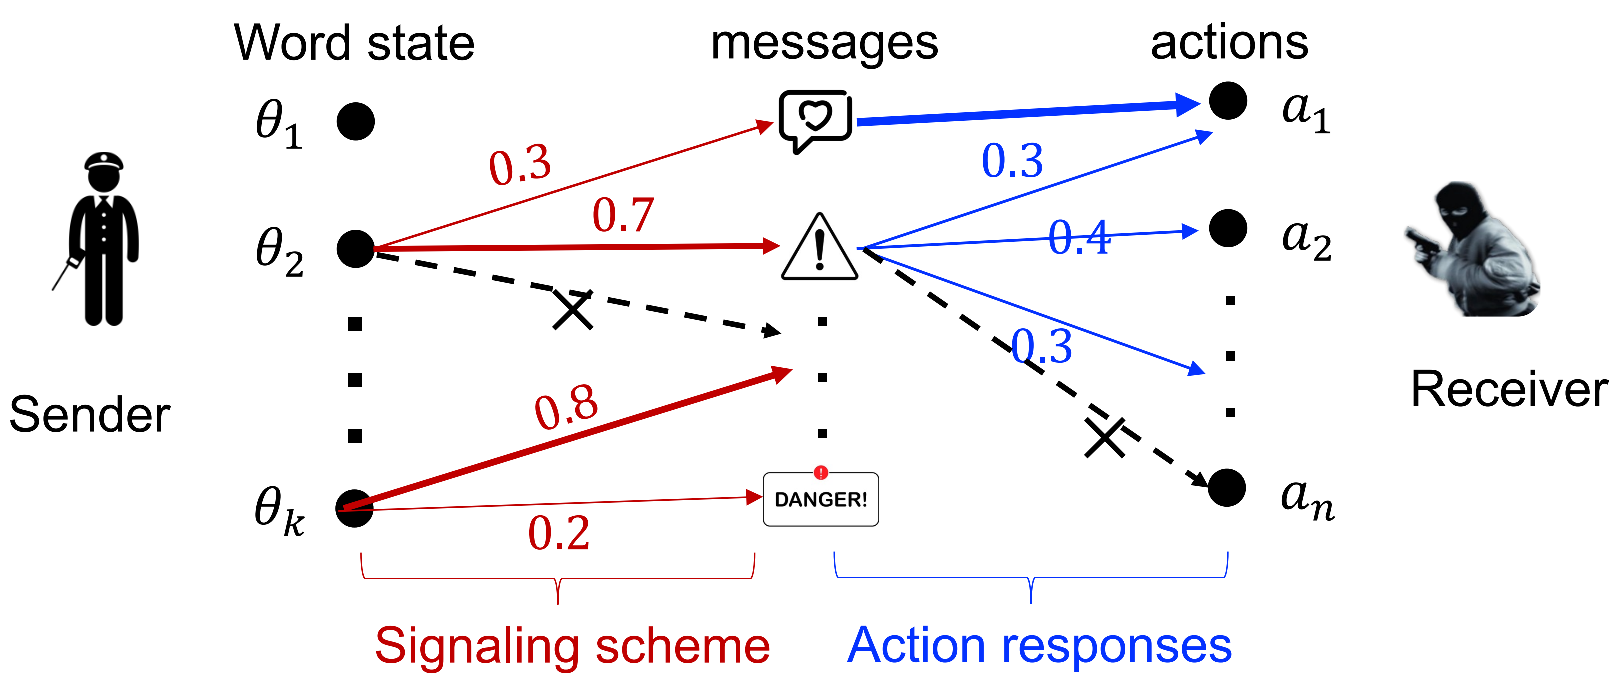

Strategic Information Transmission, Vincent P. Crawford, Joe Sobel, Econometrica 1982.

Though commitment is natural in many applications, it can be unrealistic in others (e.g., security/warship, bargaining)

- Cheap talk – information design without commitment [Crawford/Sobel'82]
- (signaling scheme, action response) forms a cheap talk equilibrium if no player wants to unilaterally deviate
  - Can mathematically formulate these as constraints on strategies

Open Algorithmic Questions

1. Complexity of computing one (or the optimal) cheap talk equilibrium in discrete-state-discrete-action game?
2. Under what conditions can the cheap talk equilibrium be efficiently computed?
   - help to explain when cheap talk is easy, and when it is not

Thank You

Questions?

haifengxu@uchicago.edu

SIGMA LAB
Strategic IntelliGence in Machine Agents