Optimal Private Payoff Manipulation against Commitment in Extensive-form games

Yurong Chen[1], Xiaotie Deng[1], Yuhao Li[2]

[1] Center on Frontiers of Computing Studies, Peking University , Beijing, China
[2]Computer Science Department, Columbia University, USA
chenyurong@pku.edu.cn, xiaotie@pku.edu.cn, yuhaoli@cs.columbia.edu

1. Introduction

Two players play a game: Leader and Follower.

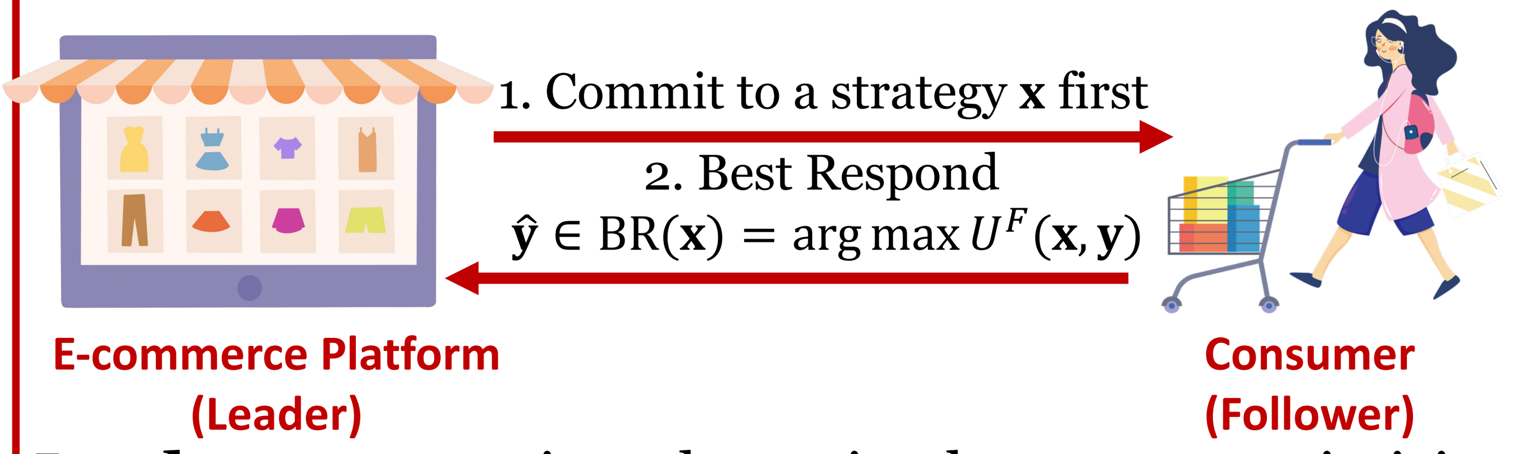

Leader can commit to the optimal strategy maximizing her own utility

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}\in A^L, \hat{\mathbf{y}}\in \mathrm{BR}(\mathbf{x})} U^L(\mathbf{x}, \hat{\mathbf{y}})$$

Strong Stackelberg Equilibrium (SSE)

- **Leader** has the first-mover advantage.
- **Follower's** payoff information U^F is needed.

Follower can provide **fake** payoff information $\tilde{U}^F$ to

- Get higher utility in the equilibrium of the new game

Practical example: customers of Uber may switch credit cards to the price charged by the platform[1].

Question here:

How to find the optimal fake $\tilde{U}^F$ that maximizes Follower's actual utility?

Previous Work:

Focus on **normal-form games**[2] and **security games**[3].

Our Work:

Extensive-form games with perfect information.

Model: Two-Phase Game

2. Settings

- **SSE** with
  - **Pure** Commitment: A^L the set of all **pure** strategies.
  - **Behavioral** Commitment: A^L the set of all **behavioral** strategies.
- **Probability Distribution** p over leaf nodes of the game tree:
  - The **game outcome** the follower aims to finally induce.
- When there exists $\tilde{U}^F$, such that
  - **Inducibility**: One SSE of game $(U^L, \tilde{U}^F)$ leads to p
  - **Strong Inducibility**: All SSEs of game $(U^L, \tilde{U}^F)$ lead to p

3. Our Contribution

For both pure and behavioral commitment settings:

1. **Characterizations** of all the (strongly) inducible distributions;
2. **Polynomial-time Algorithms** for the follower to find an (near-) optimal distribution among all the (strongly) inducible ones, and construct a corresponding follower's payoff function that induces this distribution.

WE compare the **optimal** utilities in **one** game that a follower can get in the four different settings:

3. **Inducibility v.s. Strong Inducibility**: **characterization** of the games where the two values are (nearly-)equal.
   - *Utility Supremum Equivalence (USE)* property
4. **Pure v.s. Behavioral Commitment**: The optimal value under behavioral commitment is **always no less** than that under pure commitment.

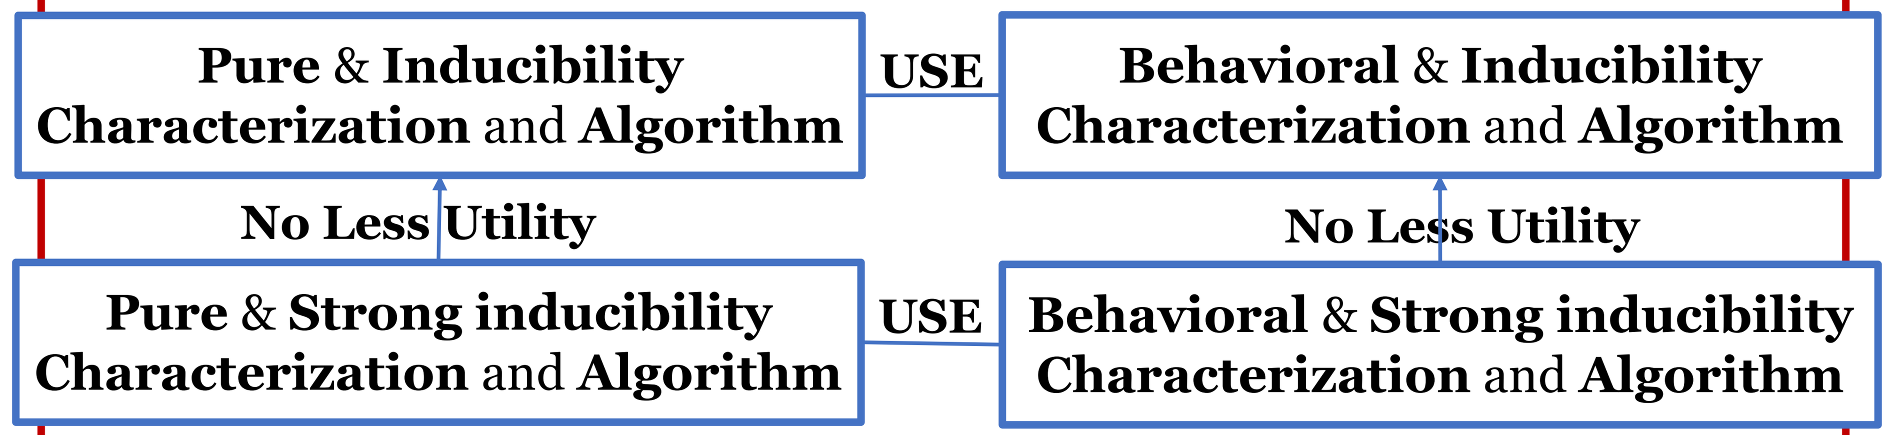

4. Pure & Inducibility

Characterization. Leaf node z is inducible iff

$$U^L(z) \geq M^L(root)$$

Where $M^L(root)$ is the leader's maximin value at root node.

Example: Follower deceives to gain better actual utility

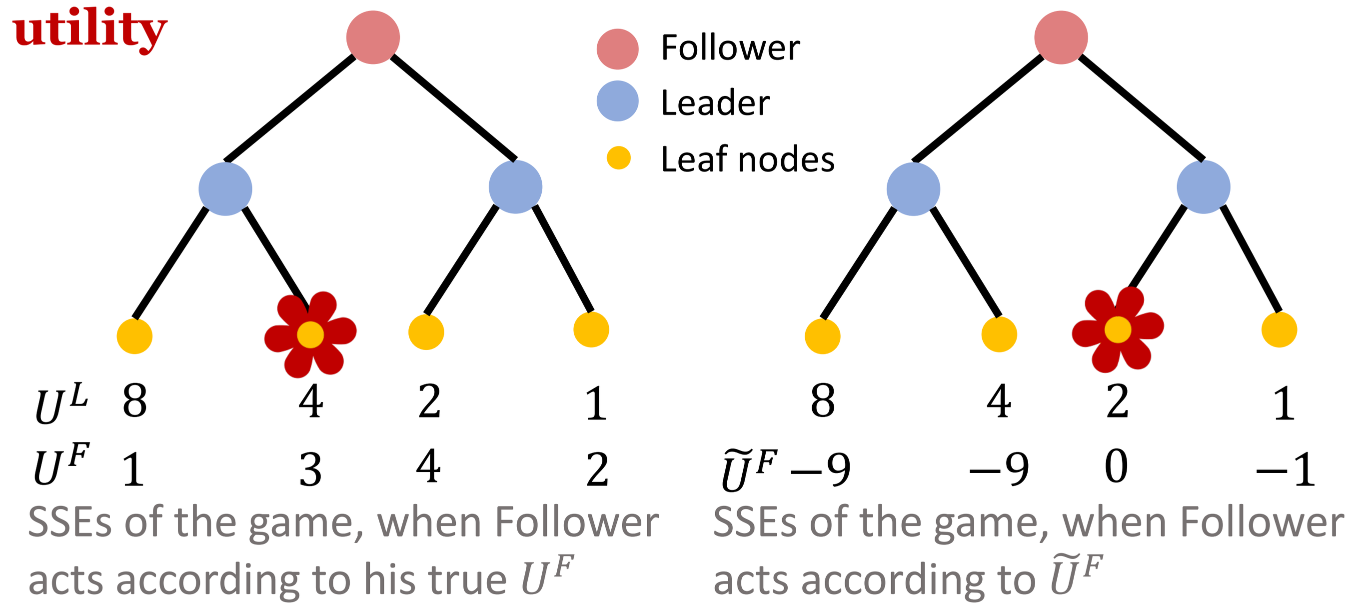

Tips for deception

- show stronger conflicts of interests where the leader can gain more utilities;
- use constant-sum subgames and games with constant worst utilities to restrict the feasible strategy profiles the leader can consider to induce via commitment;

5. For More General Settings

For Pure & Strong Inducibility, Behavioral & Inducibility and Behavioral & Inducibility settings, maximin value is **NOT ENOUGH** for characterization!

- Besides the conditions related to maximin values, the same properties need to be held on subtrees.
- And other conditions are needed.

This attributes to the Tree structure, specially owned by extensive-form games.

6. A Key Technique – "Y"-shape distribution

Definition. Distribution p is "Y"-shape if $\#\{\text{leaf node } z: p(z) \neq 0\} = 2$.

The corresponding strategy profiles of p yields distribution, s.t. edges with non-zero probabilities form a "Y".

Good Property. For any (strongly-)inducible distribution p, there exists a "Y"-shape distribution p', such that

1. p' is (strongly-)inducible;
2. $U^F(p') \geq U^F(p)$.

"Y"-shape distributions enable us to design algorithms for general settings and find the characterization for property USE

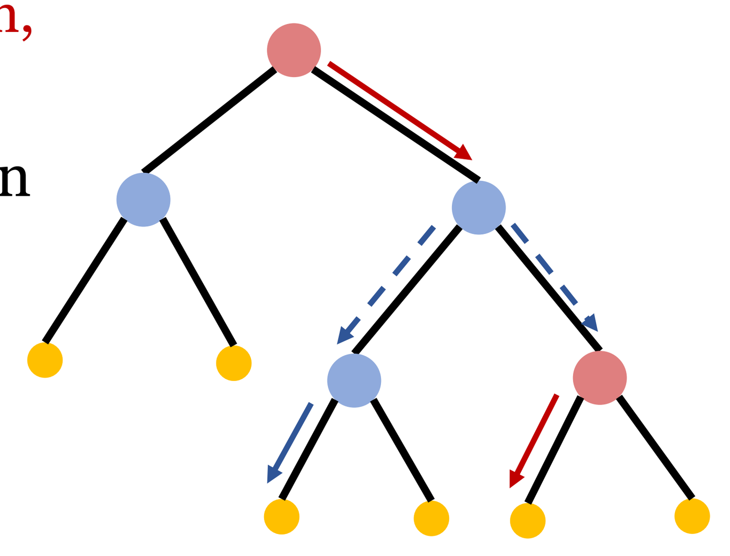

An example of Y-shape strategy pair

References:

[1] Arwa Mahdawi. 2018. Is your friend getting a cheaper Uber fare than you are? The Guardian (2018).
[2] Birmpas, Georgios, et al. "Optimally Deceiving a Learning Leader in Stackelberg Games." *Journal of Artificial Intelligence Research* 72 (2021): 507-531.
[3] Nguyen, Thanh, and Haifeng Xu. "Imitative Attacker Deception in Stackelberg Security Games." *IJCAI*. 2019.