

# Deviate or Not: Learning Coalition Structures with Multiple-bit Observations in Games

Yixuan Even Xu<sup>1</sup>, Zhe Feng<sup>2</sup>, Fei Fang<sup>1</sup>

<sup>1</sup>Carnegie Mellon University

<sup>2</sup>Google Research

{yixuanx,feif}@cs.cmu.edu, zhfef@google.com

## Abstract

We consider the Coalition Structure Learning (CSL) problem in multi-agent systems, motivated by the existence of coalitions in many real-world systems, e.g., trading platforms and auction systems. In this problem, there is a hidden coalition structure within a set of  $n$  agents, which affects the behavior of the agents in games. Our goal is to actively design a sequence of games for the agents to play, such that observations in these games can be used to learn the hidden coalition structure. In particular, we consider the setting where in each round, we design and present a game together with a strategy profile to the agents, and receive a multiple-bit observation – for each agent, we observe whether or not they would like to deviate from the specified strategy. We show that we can learn the coalition structure in  $O(\log n)$  rounds if we are allowed to design any normal-form game, matching the information-theoretical lower bound. For practicality, we extend the result to settings where we can only choose games of a specific format, and design algorithms to learn the coalition structure in these settings. For most settings, our complexity matches the theoretical lower bound up to a constant factor.

## 1 Introduction

Coalitions and collusive behavior are common in the real world. In auction systems, bidders may form coalitions to coordinate bids and exploit mechanisms for better odds of winning (Milgrom 2004). On ridesharing platforms like Uber and Lyft, some drivers might disconnect simultaneously to create an artificial shortage, triggering a price surge they can later benefit from (Hamilton 2019; Sweeney 2019; Dowling 2023). While illegal collusion, like price fixing, is regulated by agencies such as the SEC in the United States, it remains challenging to eliminate, as seen in scandals like the LIBOR Scandal (Wikipedia 2024). As AI systems become more advanced and widely used in domains like stock trading and autonomous driving, regulating these AI agents presents a growing challenge as seen by regulators and researchers (EPRS 2020; Qu, Tang, and Ma 2023).

In response to such concerns, the problem of coalition structure learning (CSL) (Xu, Ling, and Fang 2024) has recently gained attention. In CSL, a platform aims to figure out which agents are in a coalition by *designing a small number*

*of games* for the agents to play and *make observations on the outcomes* of these games. Notably, more than one coalition may exist simultaneously. In practice, the platform could represent a regulator, like an auction or ridesharing platform, aiming to detect collusion or design better mechanisms.

However, the existing study on CSL relies on a strong assumption that the platform can only observe whether a specified strategy profile is a Nash equilibrium (NE) in each round. With this single-bit observation oracle and other assumptions, optimal algorithms require  $\Theta(n \log n)$  rounds of games to learn the coalition structure, where  $n$  is the number of agents. Such an observation oracle is too restrictive for two reasons. First, for real-world systems with a large number of agents,  $\Theta(n \log n)$  rounds is still prohibitively large. Second, it may be possible for the platform to know which agents would deviate from the specified strategy profile. For example, a ridesharing platform could observe which drivers disconnect from the platform given a pricing scheme.

In this paper, we aim to address this limitation and consider a *multiple-bit observation oracle*: the platform can observe an  $n$ -dimensional binary vector in each round, where the  $i$ -th bit indicates whether or not agent  $i$  would like to deviate from the specified strategy profile. This oracle allows the platform to gain  $n$  bits of information in each round, which is significantly more informative than the single-bit observation oracle, allowing the platform to learn the coalition structure with much fewer rounds. However, from a technical perspective, switching from the single-bit observation oracle to the multiple-bit observation oracle fundamentally changes the problem. With the single-bit observation, the platform can choose the game to use to get the next bit of observation, which means it is easy to ensure that each bit of observation brings in information. However, with multi-bit observation, the platform will get  $n$ -bit information in a batch, and it is often inevitable that these  $n$ -bits contain repetitive or highly correlated information. Therefore, the key challenge in algorithm design is how to design the games so that the  $n$  bits of observation are more de-correlated and contain more information.

### 1.1 Our Results

Depending on the specific scenario of application, sometimes external constraints may limit the type of games that can be designed. For example, in an auction platform, the al-

| Type of Games | Lower Bound                            | Complexity                  | Section    |
|---------------|--|-----------------------------|------------|
| Normal-Form   | $\log_2 n - O(\log \log n)$            | $\log_2 n + 2$              | Section 3  |
| Congestion    | $\log_2 n - O(\log \log n)$            | $\log_2 n + 2$              | Appendix A |
| Graphical     | $\max(\log_2 n, n/d) - O(\log \log n)$ | $2n/d + 2 \log_2 d + 1$     | Appendix B |
| Auctions      | $\log_2 n - O(\log \log n)$            | $(1 + \log_2 n)(1 + c) + 1$ | Section 4  |

Table 1: Summary of results.

gorithm may only use auctions. Therefore, we study different settings of the type of games the algorithm may design. We summarize our results for different settings in Table 1.

**Information-theoretical lower bound.** We first show that any algorithm that learns the coalition structure requires at least  $\log_2 n - O(\log \log n)$  rounds regardless of the type of games used, where  $n$  is the number of agents.

**Normal-form games.** Normal-form games are a general class of games. We consider the case when the platform can design any normal-form game and there are no external constraints on the type of games that can be designed. We show an algorithm, *simultaneous binary search*, that can learn the coalition structure in  $\log_2 n + 2$  rounds using *directed prisoner’s dilemmas*, a type of designed games for agents to play to elicit coalition structures. The number of rounds matches the lower bound almost exactly.

**Congestion games.** Congestion games (Rosenthal 1973) are a well-studied type of games that model the allocation of resources, and has demonstrated its importance in real-world transportation system. For CSL in congestion games, we show that with a different type of games – *directed Braess’s paradoxes* – the simultaneous binary search algorithm can be adapted to learn the coalition structure with congestion games. The number of rounds required is also  $\log_2 n + 2$ .

**Graphical games.** Graphical games (Kearns, Littman, and Singh 2001) are another thoroughly researched class of games. Investigating CSL with graphical games can be useful in contexts like social networks. We consider degree- $d$  graphical games, i.e., each agent has at most  $d$  neighbors in the underlying graph. In the negative direction, we show that any algorithm for the problem requires at least  $\lceil \frac{n-1}{d} \rceil$  rounds. In the positive direction, we design an algorithm that couples simultaneous binary search with the idea of *block decomposition*, i.e., to divide the agents into blocks of size  $\lfloor \frac{d}{2} \rfloor$ . See the construction of Algorithm 5 and Fig. 5 for details of this technique. Our algorithm learns the coalition structure in  $2n/d + 2 \log_2 d + 1$  rounds. Given the two lower bounds, this algorithm is optimal up to a constant factor.

**Auctions.** Finally, we study CSL with auctions, in particular, *second-price auctions with personalized reserves*, which are widely adopted in online advertising systems. The format of auctions is much more restrictive and challenging than the previous types of games since there is only one winner in an auction, and it is difficult to elicit information from the other players. Our algorithm in this case learns the coalition structure in  $(1 + \log_2 n)(1 + c) + 1$  rounds, where  $c$  is the size of the largest coalition. When the coalitions are small, the number of rounds required is close to the lower bound.

## 1.2 Related Work

Our work lies in the general research direction of learning in games (Fudenberg and Levine 1998). In particular, our work is closely related to the line of research on Inverse Game Theory, e.g., (Waugh, Ziebart, and Bagnell 2011; Kuleshov and Schrijvers 2015; Ling, Fang, and Kolter 2018; Letchford, Conitzer, and Munagala 2009), where these previous papers aim to infer the parameters and underlying payoffs of a game by observing the behavior of players. Our work adopts this perspective but seeks to understand the *coalition structure* in games, which is not considered in prior works. By strategically and adaptively selecting defender strategies and observing the attacker’s best responses, Balcan et al. (2015); Haghtalab et al. (2016); Wu et al. (2022) show that one can uncover the underlying utility functions driving attacker behavior in Stackelberg security games. Our work learns in a similar adaptive manner but focuses on the coalition structure among agents. In addition, learning coalition structure has a fundamental difference from the above existing literature since the space of all possible coalition structures is exponentially large.

The closest related work is (Xu, Ling, and Fang 2024), where they focus on learning coalition structure through a single bit feedback, i.e., the game designed in the learning process only returns one-bit information – whether a specified strategy profile is a NE or not. In this work, we strictly generalize the model proposed by Xu, Ling, and Fang (2024) that allows us to observe multiple-bit feedback rather than just binary feedback. This flexibility strictly increases the set of possible observations and thus the difficulty of designing the games. Moreover, Bonjour, Aggarwal, and Bhargava (2022); Mazrooei, Archibald, and Bowling (2013) provide general approaches to detect *one* (collusion) coalition in multi-agent games, whereas, our work aims to find the entire coalition structure in games.

There is also a line of research on *learning with revealed preference feedback*, e.g., (Beigman and Vohra 2006; Zadimoghaddam and Roth 2012; Balcan et al. 2014; Blum, Haghtalab, and Procaccia 2014; Amin et al. 2015; Roth, Ullman, and Wu 2016), in which they target to predict the strategic agents’ behavior or optimize the profit of the decision maker, through the revealed preference feedback from strategic agents. In this paper, we assume that we can observe the best response feedback of the agents for the games, where each game is carefully designed to elicit information on the coalition structure within a set of agents. Loosely related work includes no-regret learning in games (Cesa-Bianchi and Lugosi 2006; Daskalakis, Fishelson, and Golowich 2021). Indeed, if there is a cost to perform one game to learn coalition structure, our learning pro-

cess achieves *constant* regret given the number of rounds needed to identify all coalitions only depends on the number of agents and the structure of the underlying games.

## 2 Notations and Preliminaries

Consider a set  $N = \{1, 2, \dots, n\}$  of  $n$  strategic agents. A coalition  $S \subseteq N$  is a nonempty subset of the agents, in which the agents can coordinate with each other in games. We assume there is a coalition structure  $\mathcal{S}^* = \{S_1^*, S_2^*, \dots, S_m^*\}$  among these agents, which is a set partition of  $N$ , and each agent in  $N$  belongs to exactly one of these mutually disjoint coalitions. We use  $[i]_{\mathcal{S}^*}$  to denote the coalition containing agent  $i$  under  $\mathcal{S}^*$ .  $\mathcal{S}^*$  is not public information, but agents in each coalition know the other agents in the same coalition. In the Coalition Structure Learning (CSL) problem, our goal is to actively design a sequence of games for the agents to play, such that observations in these games can be used to learn the underlying coalition structure  $\mathcal{S}^*$ . Specifically, we interact with the agents in a sequence of rounds. Each round, we design a game  $\mathcal{G}$  along with a strategy profile  $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$  in  $\mathcal{G}$  for the agents, and we make an observation  $\mathcal{O}(\mathcal{G}, \Sigma)$  about the strategy profile  $\Sigma$  in  $\mathcal{G}$ . The task is to learn the underlying coalition structure  $\mathcal{S}^*$ .

**Behavior model.** We assume that the agents are rational and strategic. When playing a game  $\mathcal{G}$ , agents in the same coalition under  $\mathcal{S}^*$  can coordinate with each other to choose actions and will share rewards after the game. Thus, they effectively play as a joint player whose action space and utility are, respectively, the Cartesian product of the action spaces and the sum of utilities of the agents in that coalition. Since the agents are not aware of the underlying coalition structure beyond their own coalitions, the information in the game is *incomplete*. We do not assume the agents’ full behavior model, but focus on the agents’ deviation decisions from the specified strategy profile.

**Multiple-bit observation oracle.** In this paper, we consider the multiple-bit observation oracle as opposed to the single-bit observation oracle considered in (Xu, Ling, and Fang 2024). The observation  $\mathcal{O}(\mathcal{G}, \Sigma)$  is an  $n$ -dimensional binary vector, where the  $i$ -th bit  $\mathcal{O}_i(\mathcal{G}, \Sigma) \in \{\text{True}, \text{False}\}$  indicates whether or not agent  $i$  would like to *deviate* from the specified strategy  $\sigma_i$  when the coalition is best responding as a whole. In particular, given a game  $\mathcal{G}$  and a strategy profile  $\Sigma$ , for each coalition  $S \in \mathcal{S}^*$ , agents in  $S$  would first determine whether  $\Sigma_S$  is a joint best response to the other agents’ specified strategy  $\Sigma_{-S}$  in  $\mathcal{G}$ . If it is, then  $\mathcal{O}_i = \text{False}$  for each agent  $i \in S$ , i.e., they would decide not to deviate. Otherwise, they would compute a possible joint best response  $\text{BR}(S, \Sigma)$  to the other agents’ specified strategy  $\Sigma_{-S}$  in  $\mathcal{G}$ . Then, they would compute the deviation decision  $\mathcal{O}_i = \mathbb{I}[\text{BR}_i(S, \Sigma) \neq \sigma_i]$  for each agent  $i \in S$ . The observation  $\mathcal{O}(\mathcal{G}, \Sigma)$  is the concatenation of the decisions of all agents in  $N$ . If there are multiple best responses for a coalition, we assume that the agents in that coalition would choose any of them in the worst case. We will use “query the observation oracle” or “query  $(\mathcal{G}, \Sigma)$ ” to refer to the process of designing a game  $\mathcal{G}$  and a strategy profile  $\Sigma$  and receiving the observation  $\mathcal{O}(\mathcal{G}, \Sigma)$ .

## 2.1 Lower Bound of the Number of Rounds

We first show a lower bound on the number of rounds required to solve the Multiple-bit CSL problem. The lower bound is based on an information-theoretical argument similar to (Xu, Ling, and Fang 2024). Note that this lower bound holds for any type of games that the algorithm may design.

**Theorem 2.1.** *Any algorithm that solves the Multiple-bit CSL problem requires at least  $\log_2 n - O(\log \log n)$  rounds of interactions with the agents in the worst case.*

*Proof.* In each round, the algorithm receives at most  $n$  bits of information. As the number of possible set partitions of  $N$  is the Bell number  $B_n$ , to distinguish between them, we need at least  $\lceil \log_2 B_n \rceil = n \log_2 n - O(n \log_2 \log_2 n)$  bits of information. The equation follows from the asymptotic expression of Bell number established in (De Bruijn 1981). The theorem then follows.  $\square$

## 3 Multiple-bit CSL with Normal-form Games

We show in this section an algorithm that solves Multiple-bit CSL with normal-form games in  $\log_2 n + 2$  rounds. Note that the number of rounds matches the lower bound in Theorem 2.1 up to  $O(\log \log n)$ . The main idea of the algorithm is to use binary search to find another agent in the same coalition for each agent simultaneously. We start by defining game-strategy pairs and a special pair that we call the *directed prisoner’s dilemma*, which we use in the algorithm.

**Definition 3.1.** *A game-strategy pair  $(\mathcal{G}, \Sigma)$  is a  $n$ -player normal-form game along with a pure strategy profile  $\Sigma$  in  $\mathcal{G}$ . Here,  $\Sigma$  is called the **specified strategy profile** of  $(\mathcal{G}, \Sigma)$ .*

**Definition 3.2.** *For  $i \in N, j \in N$ , a **directed prisoner’s dilemma**  $\mathcal{P}(i, j)$  is a game-strategy pair  $(\mathcal{G}, \Sigma)$ . In game  $\mathcal{G}$ , agent  $j$  has two actions  $\{C, D\}$ , and all other agents only have one action  $\{D\}$ . If agent  $j$  plays C, then  $i$  and  $j$  receive utility 2 and  $-1$  respectively. Otherwise they both receive 0 utility. All other agents always receive 0 utility.  $\Sigma$  is the pure strategy profile where all agents play D.*

|   |         |        |
|---|---------|--------|
|   | C       | D      |
| D | (2, -1) | (0, 0) |

Figure 1: The payoff table of agents  $i$  and  $j$  in the directed prisoner’s dilemma  $\mathcal{P}(i, j)$ . Agent  $i$  is the row player and agent  $j$  is the column player. The others only have one action and are not shown in the table.

Given the construction of **directed prisoner’s dilemma**, we have the following observation.

**Lemma 3.1.** *Let  $i, j \in N$  and  $\mathbf{O} = \mathcal{O}(\mathcal{P}(i, j))$ . Then,  $\mathbf{O}_j = \text{True}$  if and only if agent  $i$  and agent  $j$  are in the same coalition under  $\mathcal{S}^*$ . Moreover,  $\mathbf{O}_k = \text{False}$  for  $k \in N \setminus \{j\}$ .*

*Proof.* If agent  $i$  and agent  $j$  are in the same coalition under  $\mathcal{S}^*$ , then playing C is a dominant strategy for agent  $j$ ’s coalition. Otherwise, playing D is a dominant strategy. Therefore,

$O_j = \text{True}$  if and only if agent  $i$  and agent  $j$  are in the same coalition under  $\mathcal{S}^*$ . Moreover, since all agents other than  $j$  only have one action  $\{D\}$ ,  $O_k = \text{False}$  for  $k \in N \setminus \{j\}$ .  $\square$

Lemma 3.1 shows that we can use directed prisoner’s dilemmas to determine whether two agents are in the same coalition under  $\mathcal{S}^*$ . However, querying a single directed prisoner’s dilemma is not efficient as we only get 1 bit of information. To utilize the information from the observation oracle more efficiently, we will need to consider the *product* of multiple directed prisoner’s dilemmas.

**Definition 3.3** ((Xu, Ling, and Fang 2024)). *Let  $(\mathcal{G}_1, \Sigma_1), (\mathcal{G}_2, \Sigma_2)$  be two game-strategy pairs where  $A_{x,i}, u_{x,i}$  are the action set and utility function of agent  $i$  in  $\mathcal{G}_x$  respectively for  $x \in \{1, 2\}$ . Let  $\Sigma_1 = (\sigma_{1,i})_{i \in N}$  and  $\Sigma_2 = (\sigma_{2,i})_{i \in N}$ . The **product** of  $(\mathcal{G}_1, \Sigma_1)$  and  $(\mathcal{G}_2, \Sigma_2)$  is a game-strategy pair  $(\mathcal{G}_p, \Sigma_p)$ . Here,  $\mathcal{G}_p$  is a normal-form game with action set  $A_{1,i} \times A_{2,i}$  and utility function  $u_{1,i} + u_{2,i}$  for each  $i \in N$ .  $\Sigma_{p,i} = (\sigma_{1,i}, \sigma_{2,i})_{i \in N}$ . We denote the product of  $(\mathcal{G}_1, \Sigma_1)$  and  $(\mathcal{G}_2, \Sigma_2)$  as  $(\mathcal{G}_1, \Sigma_1) \times (\mathcal{G}_2, \Sigma_2)$  or  $\prod_{x=1}^2 (\mathcal{G}_x, \Sigma_x)$ . The game-strategy pairs  $(\mathcal{G}_1, \Sigma_1), (\mathcal{G}_2, \Sigma_2)$  are called the **factors** of  $(\mathcal{G}_p, \Sigma_p)$ .*

**Lemma 3.2.** *Let  $\{(i_x, j_x) \mid x \in \{1, 2, \dots, k\}\}$  be  $k$  pairs of agents and  $\mathbf{O} = \mathcal{O}(\prod_{x=1}^k \mathcal{P}(i_x, j_x))$ . Then,  $O_j = \text{True}$  if and only if there exists  $x \in \{1, 2, \dots, k\}$  such that  $j_x = j$  and  $i_x \in [j]_{\mathcal{S}^*}$ .*

*Proof.* By Definition 3.3, playing the product game is equivalent to separately playing each factor game, and summing up the resulting utilities of each agent. Therefore, agent  $j$  decides to deviate in the product game  $\prod_{x=1}^k \mathcal{P}(i_x, j_x)$  if and only if agent  $j$  decides to deviate in at least one factor game. The result then follows from Lemma 3.1.  $\square$

A graphical interpretation of Lemma 3.2 is as follows. When we query the game-strategy pair  $\prod_{x=1}^k \mathcal{P}(i_x, j_x)$ , consider a directed graph where each vertex represents an agent and each edge  $(j_x, i_x), x \in \{1, 2, \dots, k\}$  represents  $\mathcal{P}(i_x, j_x)$ . The observation tells us for each vertex if any of its outgoing neighbors in this constructed graph is in the same coalition. With Lemma 3.2, we are ready for our Algorithm 1 for Multiple-bit CSL with normal-form games.

Intuitively, given an agent  $j$ , if we want to find another agent  $i$  ( $i < j$ ) in the same coalition with  $j$ , we can query a game-strategy pair  $\prod_{k \in T_j} \mathcal{P}(k, j)$  where  $T_j = \{1, 2, \dots, j-1\}$  initially. According to Lemma 3.2, we will know whether there is such an agent in  $T_j$  from the  $j$ -th bit of the observation. If there is such an agent, we can then bisect the set  $T_j$  and apply binary search to find this agent. Moreover, this binary search can be done simultaneously for different agents. For example, given two agents  $j$  and  $j'$ , we can construct the game-strategy pairs  $\prod_{k \in T_j} \mathcal{P}(k, j) \times \prod_{k \in T_{j'}} \mathcal{P}(k, j')$  and apply binary search at the same time on  $T_j$  and  $T_{j'}$  using the  $j$ -th and  $j'$ -th bits of the observation. This is because the  $j$ -th bit in the observation vector only depends on whether there exists another agent in  $T_j$  in the same coalition with agent  $j$ , and the same argument applies to agent  $j'$ .

---

### Algorithm 1: Simultaneous Binary Search

---

**Input:** The number of agents  $n$  and the multiple-bit observation oracle  $\mathcal{O}$ .

**Output:** A coalition structure  $\mathcal{S}$  of the agents.

```

1 Query  $\prod_{i \in N, j \in N, i < j} \mathcal{P}(i, j)$  and observe  $\mathbf{O}$ ;
2 Let  $T_j \leftarrow \{1, 2, \dots, j-1\}$  if  $O_j = \text{True}$  else  $\emptyset$ 
   for each  $j \in N$ ;
3 while  $\exists j \in N$  such that  $|T_j| \geq 2$  do
4   Let  $L_j \leftarrow \{\text{the smallest } \lfloor \frac{|T_j|}{2} \rfloor \text{ elements in } T_j\}$ 
   for each  $j \in N$ ;
5   Let  $R_j \leftarrow T_j \setminus L_j$  for each  $j \in N$ ;
6   Query  $\prod_{j \in N, i \in L_j} \mathcal{P}(i, j)$  and observe  $\mathbf{O}$ ;
7   Let  $T_j \leftarrow L_j$  if  $O_j = \text{True}$  else  $R_j$ 
   for each  $j \in N$ ;
8 Let  $\mathcal{S} \leftarrow \{\{1\}, \{2\}, \dots, \{n\}\}$ ;
9 for  $j \in N$  such that  $|T_j| \neq \emptyset$  do
10  Let  $i \leftarrow \text{the only element in } T_j$ ;
11  Merge  $[i]_{\mathcal{S}}$  and  $[j]_{\mathcal{S}}$  in  $\mathcal{S}$ ;
12 return  $\mathcal{S}$ ;
```

---

Algorithm 1 leverages this intuition and works in two stages. In the first stage (Lines 1 to 7), we try to find for each agent  $j$  the smallest index of the agents in  $[j]_{\mathcal{S}^*}$ . To do this, we first query  $\prod_{i \in N, j \in N, i < j} \mathcal{P}(i, j)$  (Line 1). In this game, agent  $j$  has the incentive to deviate only when another agent  $i$  with a smaller index than  $j$  is in the same coalition with  $j$ . Therefore, using the  $j$ -th bit from the observation, we can determine whether agent  $j$  is the agent with the smallest index in  $[j]_{\mathcal{S}^*}$  (Line 2). For agents who are not, we then use binary search to locate the smallest indexed agents in their coalitions simultaneously (Lines 3 to 6). In the second stage (Lines 8 to 11), it merges each coalition in  $\mathcal{S}^*$  together according to the result of the first stage.

We provide an example of Algorithm 1 in Fig. 2.

**Theorem 3.1.** *Algorithm 1 solves Multiple-bit CSL with normal-form games in  $\log_2 n + 2$  rounds.*

*Proof.* We first show the correctness of the algorithm. To do this, we will show two claims: (i) after Lines 1 to 7, for each  $j \in N$ , if  $j$  is the agent with the smallest index in  $[j]_{\mathcal{S}^*}$ ,  $T_j = \emptyset$ , otherwise,  $T_j$  contains only the smallest index of the agents in  $[j]_{\mathcal{S}^*}$ ; (ii) after Lines 8 to 11,  $\mathcal{S} = \mathcal{S}^*$ .

For (i), if  $j$  is the agent with the smallest index in  $[j]_{\mathcal{S}^*}$ , then by Lemma 3.2,  $O_j = \text{False}$ , and  $T_j = \emptyset$  on Line 2. Throughout the algorithm,  $T_j$  remains  $\emptyset$ . Otherwise,  $T_j$  contains the smallest index of the agents in  $[j]_{\mathcal{S}^*}$  after Line 2. In the while loop (Line 3 to 7),  $T_j$  is updated to  $L_j$  if  $O_j = \text{True}$ , and to  $R_j$  otherwise. Since  $L_j$  contains the smallest  $\lfloor |T_j|/2 \rfloor$  elements in  $T_j$ , by Lemma 3.2,  $O_j = \text{True}$  if and only if the smallest index of the agents in  $[j]_{\mathcal{S}^*}$  is in  $L_j$ . Therefore, after one iteration of the loop,  $T_j$  still contains only the smallest index of the agents in  $[j]_{\mathcal{S}^*}$ , while the size of  $T_j$  is halved. The loop terminates after  $|T_j| = 1$ , and  $T_j$  contains only the smallest index of the agents in  $[j]_{\mathcal{S}^*}$ .

For (ii), since every agent  $j$  is either the agent with the

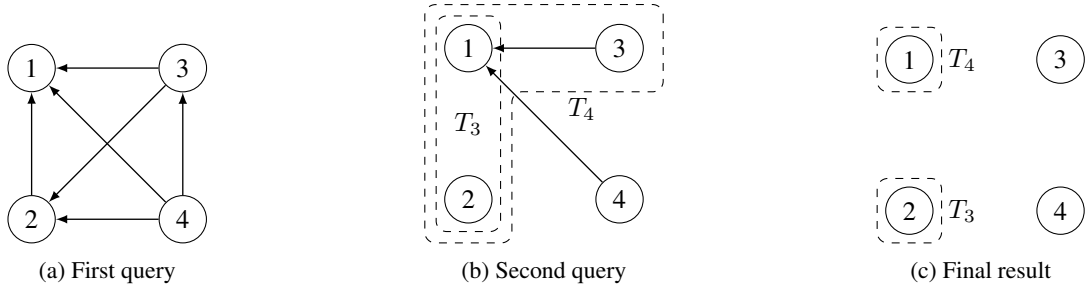


Figure 2: Example execution of Algorithm 1 when  $\mathcal{S}^* = \{\{1, 4\}, \{2, 3\}\}$ . The vertices represent the agents, the edges represent the directed prisoner’s dilemma that the algorithm queries each time, and the dashed rectangles represent the sets  $T_j$ . In the first query, the algorithm queries  $\prod_{i \in N, j \in N, i < j} \mathcal{P}(i, j)$  (Line 1) as shown in (a). Using the observations, the algorithm sets  $T_1 = T_2 = \emptyset$ ,  $T_3 = \{1, 2\}$ , and  $T_4 = \{1, 2, 3\}$  (Line 2). In the second query,  $T_3$  is bisected into  $L_3 = \{1\}$ ,  $R_3 = \{2\}$ ,  $T_4$  is bisected into  $L_4 = \{1\}$ ,  $R_4 = \{2, 3\}$ , and the algorithm queries  $\prod_{j \in N, i \in L_j} \mathcal{P}(i, j)$  (Lines 3 to 6) as shown in (b). Using the observations, the algorithm sets  $T_3 = L_3 = \{1\}$  and  $T_4 = R_4 = \{2, 3\}$  (Line 7) as shown in (c). Finally, the algorithm merges  $\{1\}$  and  $\{4\}$ , and  $\{2\}$  and  $\{3\}$  together to recover the coalition structure (Lines 8 to 11).

smallest index in  $[j]_{\mathcal{S}^*}$  or merged with that agent, each coalition is merged in  $\mathcal{S}$  after Lines 8 to 11. Thus  $\mathcal{S}$  becomes the same as  $\mathcal{S}^*$ .

Next, we show the complexity of the algorithm. The while loop in Lines 3 to 7 runs at most  $\lceil \log_2 n \rceil$  times, and each iteration requires 1 query. Together with the query on Line 1, the total number of queries is at most  $\lceil \log_2 n \rceil + 1 \leq \log_2 n + 2$ .  $\square$

**Multiple-bit CSL with Other Types of Games.** A natural next step after this section is to consider the case where the algorithm can only design games of a specific type. We will show that the idea of simultaneous binary search from Algorithm 1 can be adapted to solve the Multiple-bit CSL problem with congestion games (Appendix A), graphical games (Appendix B), or auctions (Section 4). The adaptation to congestion games is natural once we construct a new game-strategy pair that is similar in spirit to the directed prisoner’s dilemma. However, the adaptation to graphical games and auctions requires more innovations in algorithm design.

#### 4 Multiple-bit CSL with Auctions

We consider in this section the case of Multiple-bit CSL with auctions. The auction format we consider is theoretically well-studied *second-price auctions with personalized reserves* (Paes Leme, Pal, and Vassilvitskii 2016), which is widely used in practice. We first formally define this format.

**Definition 4.1.** A *second-price auction with personalized reserves*  $\mathcal{G}$  is a tuple  $(N, \mathbf{v}, \mathbf{r})$  where  $N$  is the set of agents,  $\mathbf{v} = (v_1, v_2, \dots, v_n)$  is the valuation vector, and  $\mathbf{r} = (r_1, r_2, \dots, r_n)$  is the reserve price vector. The auction proceeds as follows: Each agent  $i$  submits a bid  $\sigma_i \in \mathbb{R}_{\geq 0}$  and the agent with the highest bid wins the auction with ties broken at uniform random. If the winner  $i$  bids greater than  $r_i$ , then the winner gets the item at a price equal to the maximum between the second highest bid and  $r_i$ . The item can be reallocated among  $[i]_{\mathcal{S}^*}$ . Otherwise, the item is not allocated.

Our algorithm for the Multiple-bit CSL problem with auctions works in  $(1 + \log_2 n)(1 + c) + 1$  rounds, where  $c$  is size of the largest coalition in  $\mathcal{S}^*$ . To introduce the algorithm, we first define the game-strategy pairs that we will use in the algorithm, the *auction game gadgets*, where we will specify the valuation and reserve prices of every agents in each partition bucket.

**Definition 4.2.** Let  $\{X, Y, Z\}$  be a partition of  $N$ , i.e.,  $X \cap Y = X \cap Z = Y \cap Z = \emptyset$  and  $X \cup Y \cup Z = N$ . An *auction gadget*  $\mathcal{A}(X, Y, Z)$  is a game-strategy pair  $(\mathcal{G}, \Sigma)$  where  $\mathcal{G} = (N, \mathbf{v}, \mathbf{r})$  is a second-price auction with personalized reserves, and  $\Sigma$  is a specified bid vector for the agents, where

$$v_i = \begin{cases} 1 & (i \in X) \\ 0 & (i \in Y) \\ 0 & (i \in Z), \end{cases} \quad r_i = \begin{cases} 1 & (i \in X) \\ 0 & (i \in Y) \\ 1 & (i \in Z), \end{cases} \quad \text{and} \quad \sigma_i = 0$$

Note that in the above auction game gadgets, we do not require each agent to know all the other agents’ valuations so that the valuation of each agent is still *private* across strategic agents. Only agents in the same coalition know each other’s true valuation.

**Lemma 4.1.** For  $\{X, Y, Z\}$  that is a partition of  $N$ , let  $\mathbf{O} = \mathcal{O}(\mathcal{A}(X, Y, Z))$ . For each coalition  $S \in \mathcal{S}^*$ , if  $S \cap X \neq \emptyset$  and  $S \cap Y \neq \emptyset$ , then, there is exactly one agent  $i \in S \cap Y$  with  $O_i = \text{True}$  and  $O_j = \text{False}$  for all  $j \in S \setminus \{i\}$ . Otherwise,  $O_i = \text{False}$  for all  $i \in S$ .

*Proof.* For each coalition  $S \in \mathcal{S}^*$ . We first consider the case where  $S \cap X \neq \emptyset$  and  $S \cap Y \neq \emptyset$ . Let agent  $i \in S \cap X$  and agent  $j \in S \cap Y$ . If the coalition follows the specified strategy profile  $\Sigma_S$ , then the item will not be allocated, so the utility of the agents in  $S$  is 0. However, If agent  $j$  raises its bid to  $\sigma'_j > 0$ , then it becomes the winner of the auction, and the item will be allocated to  $j$  with price  $\max\{\text{second highest bid}, r_j\} = 0$ . The item can then be reallocated to agent  $i$ , giving utility 1 to the agents in  $S$ . Thus,  $\Sigma_S$  is not a joint best response to  $\Sigma_{-S}$ . Moreover, for agents in  $S$ , the only joint best response is to let exactly one agent

in  $S \cap Y$  raise its bid to positive, and the other agents in  $S$  keep their bids at 0. Thus, the observation  $\mathbf{O}$  is as described in the lemma.

Otherwise, either  $S \cap X = \emptyset$  or  $S \cap Y = \emptyset$ . If  $S \cap X = \emptyset$ , then no one in  $S$  has a positive valuation for the item, so the utility of the agents in  $S$  cannot be greater than 0. If  $S \cap Y = \emptyset$ , then no one in  $S$  has a reserve price that is smaller than 1, which is the largest possible valuation. Thus, the utility of the agents in  $S$  cannot be greater than 0, either. In both cases, keeping the bids at 0 is a joint best response to  $\Sigma_{-S}$ , so the observation  $\mathbf{O}$  is as described in the lemma.  $\square$

The auction gadgets are important building blocks for our algorithm that works in  $(1 + \log_2 n)(1 + c) + 1$  rounds. To help demonstrate how they can be used to solve Multiple-bit CSL, we present a simpler algorithm that works in  $n - 1$  rounds using the auction gadgets below.

---

**Algorithm 2:** Iterative Location with Auctions

---

**Input:** The number of agents  $n$  and the multiple-bit observation oracle  $\mathcal{O}$ .

**Output:** A coalition structure  $\mathcal{S}$  of the agents.

```

1 Let  $T \leftarrow N$ ;
2 Let  $\mathcal{S} \leftarrow \{\{1\}, \{2\}, \dots, \{n\}\}$ ;
3 while  $|T| \geq 2$  do
4   Let  $i \leftarrow$  the first element in  $T$ ;
5   while  $|T| \geq 2$  and  $i \in T$  do
6     Let  $X \leftarrow \{i\}, Y \leftarrow T \setminus \{i\}$  and
7        $Z \leftarrow N \setminus S$ ;
8     Query  $\mathcal{A}(X, Y, Z)$  and observe  $\mathbf{O}$ ;
9     if  $\exists j$  such that  $O_j = \text{True}$  then
10      Merge  $[i]_S$  and  $[j]_S$  in  $\mathcal{S}$ ;
11       $T \leftarrow T \setminus \{j\}$ ;
12    else
13       $T \leftarrow T \setminus \{i\}$ ;
13 return  $\mathcal{S}$ ;
```

---

To provide some intuitions, consider an agent  $i \in N$ . If we want to find another agent in the same coalition as  $i$ , we can query the auction gadget  $\mathcal{A}(\{i\}, N \setminus \{i\}, \emptyset)$ . According to Lemma 4.1, if the observation  $\mathbf{O}$  has  $O_j = \text{True}$  for some  $j \in N \setminus \{i\}$ , then  $i$  and  $j$  are in the same coalition under  $\mathcal{S}^*$ . Otherwise,  $i$  is in a coalition by itself. Then, suppose we have found an agent  $j$  in  $i$ 's coalition, and we would like to find a different agent in that coalition. We can query the auction gadget  $\mathcal{A}(\{i\}, N \setminus \{i, j\}, \{j\})$ . By moving agent  $j$  to the set  $Z$ , we ensure that  $j$  does not have the incentive to deviate, and either another agent in the same coalition as  $i$  will have  $O_k = \text{True}$  for some  $k \in N \setminus \{i, j\}$ , or we can conclude that  $i$  and  $j$  are in a coalition by themselves. We can then repeat this process until we have found all the agents in the same coalition as  $i$ . After that, we can move  $i$  to the set  $Z$  and proceed with the next agent in  $N$ . In this way, we can recover the coalition structure  $\mathcal{S}^*$ . Moreover, since each time we query an auction gadget, we move one agent to set  $Z$ , and the algorithm terminates when  $|Z| = n - 1$ , it uses  $n - 1$  queries in total.

**Theorem 4.1.** *Algorithm 2 solves the Multiple-bit CSL problem with auctions in  $n - 1$  rounds.*

We defer the proof to Appendix C. Other than the auction gadgets, we will also use another game-strategy pair in the algorithm. We prove the following lemma about the observations obtained by querying this game-strategy pair.

**Lemma 4.2.** *Let  $\mathbf{O} = \mathcal{O}((N, \mathbf{v} = \mathbf{1}, \mathbf{r} = \mathbf{0}), \Sigma = \mathbf{0})$ , where  $\mathbf{0}, \mathbf{1}$  are vectors consisting purely of 0s and 1s respectively. For each coalition  $S \in \mathcal{S}^*$ , exactly one agent  $i \in S$  has  $O_i = \text{True}$ .*

*Proof.* For each coalition  $S \in \mathcal{S}^*$ , if any agent  $i$  in  $S$  raises its bid to positive, then the item will be allocated to  $i$  with price 0, giving utility 1 to the agents in  $S$ . Thus,  $\Sigma_S$  is not a joint best response to  $\Sigma_{-S}$ . Moreover, the only joint best response for agents in  $S$  is to let exactly one agent in  $S$  raise its bid to positive, and the other agents in  $S$  keep their bids at 0. The lemma then follows.  $\square$

Then, we are ready to present our algorithm that works in  $(1 + \log_2 n)(1 + c) + 1$  rounds for Multiple-bit CSL with auctions. The algorithm is shown in Algorithm 3.

---

**Algorithm 3:** Bitwise Search with Auctions

---

**Input:** The number of agents  $n$  and the multiple-bit observation oracle  $\mathcal{O}$ .

**Output:** A coalition structure  $\mathcal{S}$  of the agents.

```

1 Query  $((N, \mathbf{v} = \mathbf{1}, \mathbf{r} = \mathbf{0}), \Sigma = \mathbf{0})$  and observe  $\mathbf{O}$ ;
2 Let  $T_x \leftarrow \{i \mid O_i = \text{True}\}$  and  $T_y \leftarrow N \setminus T_x$ ;
3 for  $b \in \{0, 1, \dots, \lfloor \log_2 n \rfloor\}$  do
4   Let  $X \leftarrow \{i \mid i \in T_x \text{ and the } b\text{-th lowest binary bit of } i \text{ is } 1\}$ ;
5   Let  $T_{\text{False}} \leftarrow T_y$ ;
6   repeat
7     Let  $Y \leftarrow T_{\text{False}}$  and  $Z \leftarrow N \setminus (X \cup Y)$ ;
8     Query  $\mathcal{A}(X, Y, Z)$  and observe  $\mathbf{O}$ ;
9     Let
10       $T_{\text{True}} \leftarrow \{i \mid i \in T_{\text{False}} \text{ and } O_i = \text{True}\}$ ;
11       $T_{\text{False}} \leftarrow T_{\text{False}} \setminus T_{\text{True}}$ ;
12   until  $T_{\text{True}} = \emptyset$ ;
13 Let  $O_i^{(b)} \leftarrow \mathbb{I}[i \notin T_{\text{False}}]$  for each  $i \in T_y$ ;
14 Let  $\mathcal{S} \leftarrow \{\{1\}, \{2\}, \dots, \{n\}\}$ ;
15 for  $i \in T_y$  do
16   Let  $j \leftarrow \sum_{b=0}^{\lfloor \log_2 n \rfloor} 2^b \cdot O_i^{(b)}$ ;
17   Merge  $[i]_S$  and  $[j]_S$  in  $\mathcal{S}$ ;
17 return  $\mathcal{S}$ ;
```

---

The algorithm starts by querying the game-strategy pair  $((N, \mathbf{v} = \mathbf{1}, \mathbf{r} = \mathbf{0}), \Sigma = \mathbf{0})$  (Line 1). Using the obtained observation, by Lemma 4.2, the algorithm identifies exactly one agent in each coalition, groups them into set  $T_x$  and the rest of the agents into set  $T_y$  (Lines 2). Essentially, set  $T_x$  contains a set of agents who are definitely not in the same coalition with each other. Furthermore, for each agent  $i$  in  $T_y$ , there is exactly one ‘‘teammate’’ of  $i$  in  $T_x$ . We denote the index of this agent as  $\alpha(i)$ .




Figure 3: Example execution of the bitwise search (Lines 1 to 12) in Algorithm 3 when  $\mathcal{S}^* = \{\{1, 4\}, \{2, 3\}, \{5, 6\}\}$ . The vertices represent the agents, and the dashed lines represent the sets used in the algorithm. Using the first query, the algorithm identifies one agent in each coalition and groups them as  $T_x = \{1, 3, 5\}$ . The rest are  $T_y = \{2, 4, 6\}$  (Lines 1 to 2). Then, as shown above, for each  $b \in \{0, 1, 2\}$ , the algorithm picks the set of agents in  $T_x$  with the  $b$ -th lowest binary bit as 1 as  $X$ , and partitions  $T_y$  into  $T_{\text{True}}$ , each of which is cooperating with some agent in  $X$ , and  $T_{\text{False}}$ , each of which is not cooperating with any agent in  $X$  (Lines 3 to 12).

An important idea behind Algorithm 3 is *bitwise search*, i.e., to determine each binary bit of  $\alpha(i)$  separately. To do this, for each  $b \in \{0, 1, \dots, \lfloor \log_2 n \rfloor\}$ , the algorithm picks the set of agents in  $T_x$  such that the  $b$ -th lowest binary bits of their index are 1 as  $X$  (Line 4), and tries to tell which agents in  $T_y$  are in the same coalition as some agent in  $X$  by querying auction gadgets. Intuitively, if we let  $Y = T_y$ ,  $Z = N \setminus (X \cup Y)$  and query  $\mathcal{A}(X, Y, Z)$ , by Lemma 4.1, for each coalition that contains an agent in  $X$ , one of its members in  $Y$  will decide to deviate. Thus, we can move these agents to  $Z$  like how we did in Algorithm 2, until no agents in  $Y$  decide to deviate. The remaining agents in  $Y$  are then the agents that are not in the same coalition as any agent in  $X$ . The algorithm leverages this idea to determine each binary bit of  $\alpha(i)$  separately, and stores the results in  $O_i^{(b)}$  (Lines 5 to 12). Note that for each  $b$ , the algorithm uses at most  $1 + c$  queries, where  $c$  is the size of the largest coalition in  $\mathcal{S}^*$ . We illustrate the execution of the bitwise search with an example in Fig. 3.

Finally, the algorithm recovers  $\alpha(i)$  for each agent  $i$  in  $T_y$  by summing up the values of the binary bits in  $O_i^{(b)}$  (Line 15), and merges  $i$  with  $\alpha(i)$  accordingly (Line 16).

**Theorem 4.2.** *Algorithm 3 solves the Multiple-bit CSL problem with auctions in  $(1 + \log_2 n)(1 + c) + 1$  rounds, where  $c$  is size of the largest coalition in  $\mathcal{S}^*$ .*

*Proof.* We first show the correctness of the algorithm. By Lemma 4.2, after Line 2,  $T_x$  contains exactly one agent in each coalition. This means that for each agent  $i \in T_y$ , there is exactly one agent  $j \in T_x$  such that  $i$  and  $j$  are in the same coalition. We denote this agent  $j$  as  $\alpha(i)$ . Then, we will show two claims: (i) after Lines 3 to 12, for each  $b \in \{0, 1, \dots, \lfloor \log_2 n \rfloor\}$  and each  $i \in T_y$ ,  $O_i^{(b)} = \text{True}$  if and only if the  $b$ -th lowest binary bit of  $\alpha(i)$  is 1, and (ii) after Lines 13 to 16,  $\mathcal{S} = \mathcal{S}^*$ .

For (i), we only need to show that after the repeat loop in Lines 6 to 11,  $T_{\text{False}}$  contains exactly the agents  $i \in T_y$  such that the  $b$ -th lowest binary bit of  $\alpha(i)$  is 0. If an agent  $i \in T_y$  has the  $b$ -th lowest binary bit of  $\alpha(i)$  as 0, then  $\alpha(i) \notin X$  after Line 4. Thus, by Lemma 4.1,  $O_i$  will never be True

on Line 8, and  $i \in T_{\text{False}}$  after the repeat loop. If an agent  $i \in T_y$  has the  $b$ -th lowest binary bit of  $\alpha(i)$  as 1, then  $\alpha(i) \in X$  after Line 4. By Lemma 4.1, in each iteration of the repeat loop, one agent  $j \in [i]_{\mathcal{S}^*} \cap T_{\text{False}}$  will have  $O_j = \text{True}$ , and will be removed from  $T_{\text{False}}$ . Since the repeat loop terminates only when no agent is removed from  $T_{\text{False}}$  in one iteration,  $i \notin T_{\text{False}}$  after the loop.

For (ii), given each binary bit of  $\alpha(i)$ , the algorithm computes the index of the agent in  $T_y$  by summing up the values of the binary bits on Line 15. Then, the algorithm merges  $i$  with  $\alpha(i)$  accordingly on Line 16. The correctness of the algorithm then follows.

For the complexity, the algorithm uses 1 query on Line 1, and for each  $b \in \{0, 1, \dots, \lfloor \log_2 n \rfloor\}$ , the algorithm uses at most  $1 + c$  queries in the repeat loop in Lines 6 to 11. This shows that the total number of queries used by the algorithm is at most  $(1 + \log_2 n)(1 + c) + 1$ .  $\square$

## 5 Conclusion and Future Work

In this paper, we study the Coalition Structure Learning (CSL) problem under the multiple-bit observation oracle. We consider various settings of the type of games that the algorithm may design, including normal-form games, congestion games, graphical games, and auctions. In each of the settings, we present an algorithm that learns the coalition structure with a sublinear number of rounds. We also complement our algorithmic results with lower bounds on the number of rounds required to learn the coalition structure, demonstrating their optimality in most of the settings. Compared to the existing work, our results significantly reduce the number of rounds required to learn the coalition structure, and greatly improve the potential applicability of the CSL problem in real-world systems. Regarding CSL, there are various different settings for future work that are not covered by this paper. It would be interesting to consider the CSL problem (i) when the underlying coalition structure is dynamic, (ii) when the agents are aware of the algorithm and respond strategically, (iii) when the algorithm can only observe a subset of the agents in each round, and (iv) when the observations are noisy or the agents are boundedly rational.

## Acknowledgements

This work was supported in part by NSF grant IIS-2046640 (CAREER) and NSF IIS-2200410.

## References

- Amin, K.; Cummings, R.; Dworkin, L.; Kearns, M.; and Roth, A. 2015. Online learning and profit maximization from revealed preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, 770–776. AAAI Press. ISBN 0262511290.
- Balcan, M.-F.; Blum, A.; Haghtalab, N.; and Procaccia, A. D. 2015. Commitment Without Regrets: Online Learning in Stackelberg Security Games. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, 61–78. New York, NY, USA: Association for Computing Machinery. ISBN 9781450334105.
- Balcan, M.-F.; Daniely, A.; Mehta, R.; Urner, R.; and Vazirani, V. V. 2014. Learning Economic Parameters from Revealed Preferences. In Liu, T.-Y.; Qi, Q.; and Ye, Y., eds., *Web and Internet Economics*, 338–353. Cham: Springer International Publishing.
- Beigman, E.; and Vohra, R. 2006. Learning from revealed preference. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, EC '06, 36–42. New York, NY, USA: Association for Computing Machinery. ISBN 1595932364.
- Blum, A.; Haghtalab, N.; and Procaccia, A. D. 2014. Learning Optimal Commitment to Overcome Insecurity. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Bonjour, T.; Aggarwal, V.; and Bhargava, B. 2022. Information theoretic approach to detect collusion in multi-agent games. In Cussens, J.; and Zhang, K., eds., *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180 of *Proceedings of Machine Learning Research*, 223–232. PMLR.
- Cesa-Bianchi, N.; and Lugosi, G. 2006. *Prediction, Learning, and Games*. Cambridge University Press.
- Daskalakis, C.; Fishelson, M.; and Golowich, N. 2021. Near-optimal no-regret learning in general games. *Advances in Neural Information Processing Systems*, 34: 27604–27616.
- De Bruijn, N. G. 1981. *Asymptotic methods in analysis*, volume 4. Courier Corporation.
- Dowling, J. 2023. How Uber drivers trigger fake surge price periods when no delays exist. *Drive*.
- EPRS. 2020. The ethics of artificial intelligence: Issues and initiatives. *European Parliament Research Service, Scientific Foresight Unit (STOA), Panel for the Future of Science and Technology*.
- Fudenberg, D.; and Levine, D. K. 1998. *The Theory of Learning in Games*, volume 1 of *MIT Press Books*. The MIT Press.
- Haghtalab, N.; Fang, F.; Nguyen, T. H.; Sinha, A.; Procaccia, A. D.; and Tambe, M. 2016. Three strategies to success: learning adversary models in security games. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, 308–314. AAAI Press. ISBN 9781577357704.
- Hamilton, I. A. 2019. Uber Drivers Are Reportedly Colluding to Trigger ‘Surge’ Prices Because They Say the Company Is Not Paying Them Enough. *Business Insider*.
- Kearns, M.; Littman, M. L.; and Singh, S. 2001. Graphical Models for Game Theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial, Intelligence*, 2001, 253–260.
- Kuleshov, V.; and Schrijvers, O. 2015. Inverse Game Theory: Learning Utilities in Succinct Games. In Markakis, E.; and Schäfer, G., eds., *Web and Internet Economics*, 413–427. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Letchford, J.; Conitzer, V.; and Munagala, K. 2009. Learning and Approximating the Optimal Strategy to Commit To. In Mavronicolas, M.; and Papadopoulou, V. G., eds., *Algorithmic Game Theory*, 250–262. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Ling, C. K.; Fang, F.; and Kolter, J. Z. 2018. What Game Are We Playing? End-to-end Learning in Normal and Extensive Form Games. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, IJCAI-18, 396–402. International Joint Conferences on Artificial Intelligence Organization.
- Mazrooei, P.; Archibald, C.; and Bowling, M. 2013. Automating Collusion Detection in Sequential Games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 27(1): 675–682.
- Milgrom, P. R. 2004. *Putting auction theory to work*. Cambridge University Press.
- Paes Leme, R.; Pal, M.; and Vassilvitskii, S. 2016. A field guide to personalized reserve prices. In *Proceedings of the 25th international conference on world wide web*, 1093–1102.
- Qu, A.; Tang, Y.; and Ma, W. 2023. Adversarial Attacks on Deep Reinforcement Learning-based Traffic Signal Control Systems with Colluding Vehicles. *ACM Transactions on Intelligent Systems and Technology*, 14(6): 1–22.
- Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2: 65–67.
- Roth, A.; Ullman, J.; and Wu, Z. S. 2016. Watch and learn: optimizing from revealed preferences feedback. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, 949–962. New York, NY, USA: Association for Computing Machinery. ISBN 9781450341325.
- Sweeney, S. 2019. Uber, Lyft drivers manipulate fares at Reagan National causing artificial price surges. *WJLA*.
- Waugh, K.; Ziebart, B. D.; and Bagnell, J. A. 2011. Computational rationalization: the inverse equilibrium problem. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML'11, 1169–1176. Madison, WI, USA: Omnipress. ISBN 9781450306195.

Wikipedia. 2024. Libor scandal. [https://en.wikipedia.org/wiki/Libor\\_scandal](https://en.wikipedia.org/wiki/Libor_scandal). Accessed: 2024-08-11.

Wu, J.; Shen, W.; Fang, F.; and Xu, H. 2022. Inverse Game Theory for Stackelberg Games: the Blessing of Bounded Rationality. *Advances in Neural Information Processing Systems*, 35: 32186–32198.

Xu, Y. E.; Ling, C. K.; and Fang, F. 2024. Learning Coalition Structures with Games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9944–9951.

Zadimoghaddam, M.; and Roth, A. 2012. Efficiently Learning from Revealed Preference. In Goldberg, P. W., ed., *Internet and Network Economics*, 114–127. Berlin, Heidelberg: Springer Berlin Heidelberg.