Learning Efficient Truthful Mechanisms for Trading Networks

Takayuki Osogami[1] , **Segev Wasserkrug**[1] , **Elisheva S. Shamash**[2*]
[1]IBM Research
[2]Technion - Israel Institute of Technology
osogami@jp.ibm.com, segevw@il.ibm.com, elisheva.shamash@gmail.com

Abstract

Trading networks are an indispensable part of today's economy, but to compete successfully with others, they must be efficient in maximizing the value they provide to the external market. While the prior work relies on truthful disclosure of private information to achieve efficiency, we study the problem of designing mechanisms that result in efficient trading networks by incentivizing firms to truthfully reveal their private information to a third party. Additional desirable properties of such mechanisms are weak budget balance (WBB; the third party needs not invest) and individual rationality (IR; firms get non-negative utility). Unlike combinatorial auctions, there may not exist mechanisms that simultaneously satisfy these properties *ex post* for trading networks. We propose an approach for computing or learning truthful and efficient mechanisms for given networks in a Bayesian setting, where WBB and IR, respectively, are relaxed to *ex ante* and *interim* for a given distribution over the private information. We incorporate techniques to reduce computational and sample complexity. We empirically demonstrate that the proposed approach successfully finds the mechanisms with the relaxed properties for trading networks where achieving *ex post* properties is impossible.

1 Introduction

Trading networks [Hatfield *et al.*, 2013] are becoming increasingly ubiquitous and essential in vital areas such as transportation and supply chain. A trading network consists of the firms who trade through bi-lateral contracts and aims at benefiting those firms and the markets external to the network. A major goal of a trading network is thus *efficiency* in the sense of maximizing the value (social welfare) that the trading network creates.

Prior work on trading networks primarily investigates solution concepts such as stability and competitive equilibrium [Hatfield *et al.*, 2013; Candogan *et al.*, 2021], which result in efficiency. To compute these solutions for a given trading network, however, existing algorithm require the information about how the firms value trades (which we define as the types of the firms), private information which therefore needs to be truthfully revealed by the firms. As we will show, however, the firms typically have incentive to be untruthful.

We thus seek to provide a mechanism for a given trading network such that the firms have the incentive to be truthful, which in turn leads to efficiency. To this end, we allow the firms to make (possibly negative) payment to an Independent Party (IP). In this setting, we desire mechanisms that satisfies weak budget balance (WBB; non-negative utility of IP) and individual rationality (IR; non-negative utility of firms) in addition to efficiency and dominant-strategy incentive compatibility (DSIC; promotion of truthfulness). These four properties are standard in mechanism design [Parkes, 2001]. In particular, for combinatorial auctions, the VCG mechanism with the Clarke pivot rule is known to satisfy all of the four properties *ex post*, i.e. surely for any types, [Nisan, 2007].

However, trading networks are fundamentally more complex than combinatorial auctions. In particular, a trade transfers a good from a seller to a buyer, and the seller who has negative valuation on the trade is compensated by the payment from the buyer. Therefore, as we describe in the sequel, it is impossible to simultaneously satisfy the four properties *ex post* for a broad class of trading networks, making the design of mechanisms for such networks non-trivial, in contrast to combinatorial auctions.

We thus relax the WBB and IR properties, and restrict ourselves to a Bayesian setting, in which there is a (prior) discrete distribution over the types. In this setting, we seek to compute a mechanism that satisfies WBB and IR *ex ante* and *interim* respectively (i.e., in expectation with respect to the distribution of types). To ensure DSIC and efficiency, we utilize the class of Groves mechanisms and show how to compute the pivot rule of the Groves mechanism via a linear program (LP) that encodes WBB and IR as constraints.

While this LP successfully computes a mechanism with the four desirable properties, it involves two shortcomings. First, the LP needs the exact knowledge of the distribution of types. Second, the LP becomes intractable as the number of trades, players, or types increases. To mitigate these shortcomings, we provide a *mechanism learning* approach, which only requires a sample of types instead of their exact distribution. We further propose several techniques to reduce the computa-

*This work was performed at IBM Research.

tional and sample complexity of the mechanism learning. In particular, we manually design special Groves mechanisms, which achieve some but not all of the four properties *ex post*, and demonstrate how the dimension of the feature vector of the types can be reduced based on the knowledge of those special mechanisms.

Our contribution is an automated mechanism design (AMD; [Conitzer and Sandholm, 2002]) approach for computing or learning desirable mechanisms for trading networks (Section 4-5). This consists of a) formulating an LP whose solution gives a Groves mechanism that satisfies all of the four desirable properties for a trading network with a given distribution of types, b) integrating machine learning techniques into the LP to extend the cases where such mechanisms can be designed, and c) reducing the number of variables in LP or the dimension of the feature vectors in the corresponding learning approach based on special Groves mechanisms. Furthermore, we provide empirical analysis to support the effectiveness of our AMD approaches (Section 6). For reproducibility, the source code is open-sourced at https://github.com/IBM/mechanism-learning.

1.1 Related Work

Trading networks and their efficiency have been extensively studied [Hatfield *et al.*, 2013; Candogan *et al.*, 2021; Ostrovsky, 2008; Hatfield and Kominers, 2012; Hatfield *et al.*, 2015]), but no approaches are known to simultaneously achieve the four properties in trading networks. In fact, there are few studies on incentive compatibility in trading networks with an exception of [Schlegel, 2022], who establishes incentive compatibility but only for those firms who are buyers (or sellers) in all trades.

From methodological perspectives, the prior work most related to ours is some of the AMD approaches for combinatorial auctions, where constrained optimization problems are formulated and approximated with the sample from the distribution of types. However, in [Duetting *et al.*, 2019; Rahme *et al.*, 2021], DSIC is encoded as a constraint of the optimization problem and is not necessarily guaranteed due to the sample approximation unless the dataset covers the full support of the distribution, while we always ensure DSIC via the Groves mechanism (hence, we do not require the knowledge of the support, which is a fundamental advantage over [Duetting *et al.*, 2019; Rahme *et al.*, 2021]). On the other hand, [Manisha *et al.*, 2018; Tacchetti *et al.*, 2022] learn the mechanism (specifically, the rule of redistributing payment from IP to players) that minimizes the expected revenue of IP, while ensuring the four properties via the VCG mechanism. In trading networks, however, it is possible to guarantee only two properties via the Groves mechanism, and we encode the other two properties as constraints of our LP.

[Alon *et al.*, 2021] also study an approach of computing (no learning) a mechanism for a principal-agent model as a solution to LP, while ensuring some of the properties via the VCG mechanism. Earlier work along this line of computing a mechanism within a restricted class of mechanisms includes [Likhodedov and Sandholm, 2004], who maximize expected revenue in combinatorial auctions. Sample approximation of this method is studied by [Likhodedov and Sandholm, 2005].

2 Trading Networks

Following [Hatfield *et al.*, 2013], we model a trading network by a tuple $(\mathcal{N}, \Omega, v)$, where $\mathcal{N}$ is a set of players (or firms), Ω is a set of bi-lateral trades, and $v \equiv (v_i)_{i\in\mathcal{N}}$, where each $v_i : 2^\Omega \to \mathbb{R}$ is the type (valuation) of player i. That is, $v_i(\Phi)$ represents the (possibly negative) value of $\Phi \subseteq \Omega$ for the player with type v_i (i.e., a player's valuation is allowed to depend on the trades between other players). For notational convenience, we denote a trading network $(\mathcal{N}, \Omega, v)$ as $\mathcal{T}(v)$ when specific $\mathcal{N}$ and Ω are unimportant or clear from the context. Without loss of generality, we assume that the value of no trade is zero for any player (i.e., $v_i(\emptyset) = 0, \forall i \in \mathcal{N}$). Each trade is associated with a seller and a buyer in $\mathcal{N}$, and the buyer makes non-negative payment to the seller. Let $p(\omega)$ be the payment associated with $\omega \in \Omega$. We refer to a pair $(\omega, p(\omega))$ as a contract. A pair of Φ and p determines a set of contracts, $\{(\omega, p(\omega)) : \omega \in \Phi\}$.

One's goal with a trading network is to determine a set of trades $\Phi \subseteq \Omega$ to be conducted together with the payment associated with each trade in a way that certain properties are satisfied. A particularly important metric is the total valuations associated with the trades to be conducted (i.e., $\sum_{i\in\mathcal{N}} v_i(\Phi)$), since it is the value that the trading network produces to the external markets. The total valuations may be considered as net payment from the external markets to the trading network, who buys raw materials and sells final products. Therefore, we say that the set of trades $\Phi^\star$ is efficient for a trading network $\mathcal{T}(v)$ if

$$\sum_{i\in\mathcal{N}} v_i(\Phi^\star) \geq \sum_{i\in\mathcal{N}} v_i(\Phi), \forall \Phi \in 2^\Omega. \tag{1}$$

The efficiency may also be represented in terms of the utilities of players. Specifically, by conducting the set of contracts $\{(\omega, p(\omega)) : \omega \in \Phi\}$, the player $i \in \mathcal{N}$ gets the following (quasi-linear) utility:

$$\tilde{u}_i((\Phi, p); \mathcal{T}(v)) = v_i(\Phi) + \sum_{\omega\in\Phi_{i\to}} p(\omega) - \sum_{\omega\in\Phi_{\to i}} p(\omega),$$

where $\Phi_{i\to}$ is the subset of Φ where i is the seller, and $\Phi_{\to i}$ is the one where i is the buyer. Since the total payment from sellers equals the total payment to buyers, we can show (see (26)-(29) in [Osogami *et al.*, 2022]) that, for any p, (1) is equivalent to

$$\sum_{i\in\mathcal{N}} \tilde{u}_i((\Phi^\star, p); \mathcal{T}(v)) \geq \sum_{i\in\mathcal{N}} \tilde{u}_i((\Phi, p); \mathcal{T}(v)), \forall \Phi \in 2^\Omega. \tag{2}$$

The prior work has developed algorithms for finding the set of contracts that have the properties stronger than efficiency (e.g., competitive equilibrium and stability) under certain conditions (e.g., full substitutability [Hatfield *et al.*, 2019]). However, they all rely on the knowledge of the types.

In practice, however, the types are private information and need to be revealed by the players. Then, the players can have the incentive to be untruthful, which in turn leads to inefficiency, as the following example shows:

Example 1. *Consider a trading network with a single potential trade, $\Omega = \{\omega\}$, between two players, $\mathcal{N} = \{\mathrm{S}, \mathrm{B}\}$. With ω, the seller (S) incurs a production cost C_S, and the buyer (B) gets a profit $1-C_\mathrm{B}$ (retail price minus handling cost): i.e., $v_\mathrm{S}(\Phi) = -C_\mathrm{S}\, I\{\Phi = \Omega\}$ and $v_\mathrm{B}(\Phi) = (1-C_\mathrm{B})\, I\{\Phi = \Omega\}$, where $I\{\cdot\}$ is the indicator function. Consider the payment, $p(\omega) = \frac{1+C_\mathrm{S}-C_\mathrm{B}}{2}$, that equally shares the net profit. Given the types (i.e., C_S and C_B), we can achieve efficiency by letting $\Phi = \Omega$ if $C_\mathrm{S} + C_\mathrm{B} < 1$ and $\Phi = \emptyset$ otherwise. When $C_\mathrm{S} + C_\mathrm{B} < 1$, the utility of each player is $\frac{1-C_\mathrm{S}-C_\mathrm{B}}{2}$. However, then each player i has the incentive to untruthfully declare higher cost $C_i + \varepsilon$ to get higher utility $\frac{1-C_\mathrm{S}-C_\mathrm{B}+\varepsilon}{2}$. If the declared cost is too high, the trade is not conducted. Since $1 - C_\mathrm{S} - C_\mathrm{B} = v_\mathrm{S}(\Omega) + v_\mathrm{B}(\Omega) > v_\mathrm{S}(\emptyset) + v_\mathrm{B}(\emptyset) = 0$, untruthfulness indeed leads to inefficiency.*

3 Mechanisms for Trading Networks

We consider the trading networks where the types are private information of the respective players and study the mechanisms that promote truthfulness, leading to efficiency. To facilitate truthfulness, we allow the mechanism to require each player i to make (possibly negative) payment $\tau_i \in \mathbb{R}$ to an independent party (IP). A trading network with IP is denoted by $\mathcal{T}^+(v) = (\mathcal{N}^+, \Omega, v)$, where $\mathcal{N}^+ \equiv \mathcal{N} \cup \{\mathrm{IP}\}$ denotes the set of all players and IP. Also, let $\mathcal{V}_i$ be the discrete space of types for each $i \in \mathcal{N}$ and let $\mathcal{V} \equiv \times_{i\in\mathcal{N}} \mathcal{V}_i$ be the product space. Moreover, let $\mathcal{T}^+(\mathcal{V}) \equiv (\mathcal{N}, \Omega, \mathcal{V}) \equiv \{\mathcal{T}^+(v) : v \in \mathcal{V}\}$ be the set of trading networks with IP under $\mathcal{V}$. We will simply refer to $\mathcal{T}^+(v)$ or $\mathcal{T}^+(\mathcal{V})$ as a trading network.

We study the Bayesian setting where there exists a (prior) distribution q over $\mathcal{V}$ such that the players have types v with probability $q(v)$. Throughout, we assume that the true type v_i of each player $i \in \mathcal{N}$ is known by that player. However, we we make varying assumptions on other knowledge about the types, which are made explicit in each result in the following.

3.1 Direct Mechanisms and Desirable Properties

We study direct mechanisms, where players and IP act according to the following protocol:

1. Each player i declares a possibly untruthful type $\hat{v}_i$
2. IP determines the set of contracts, $\{(\omega, \pi(\omega; \hat{v})) : \omega \in \phi(\hat{v})\}$, and (possibly negative) payment, $\tau_i(\hat{v})$, from each player i to IP

A direct mechanism of a trading network $\mathcal{T}(\mathcal{V})$ is specified by an outcome rule (ϕ, τ, π), where $\phi : \mathcal{V} \to 2^\Omega$ is the allocation rule that maps declared types, $\hat{v} \equiv (\hat{v}_i)_{i\in\mathcal{N}}$, to a set of trades $\phi(\hat{v})$ to be conducted; τ determines the rule of payment to IP, $\tau_i \in \mathcal{V} \to \mathbb{R}$, for each $i \in \mathcal{N}$; $\pi : 2^\Omega \times \mathcal{V} \to \mathbb{R}$ is the payment rule that determines the payment associated with each trade, depending on the declared types. In the following, we refer to a direct mechanism simply as a mechanism.

When $\hat{v}$ is declared by the players in the trading network $\mathcal{T}^+(v)$, each player $i \in \mathcal{N}$ gets the following utility under the mechanism (ϕ, τ, π):

$$u_i(\hat{v}; (\phi, \tau, \pi), \mathcal{T}^+(v)) \tag{3}$$
$$= v_i(\phi(\hat{v})) + \sum_{\omega\in\phi(\hat{v})_{i\to}} \pi(\omega; \hat{v}) - \sum_{\omega\in\phi(\hat{v})_{\to i}} \pi(\omega; \hat{v}) - \tau_i(\hat{v}).$$

We denote the net-payment to IP (or utility of IP) by

$$u_\mathrm{IP}(\hat{v}; (\phi, \tau, \pi), \mathcal{T}^+(v)) = \sum_{i\in\mathcal{N}} \tau_i(\hat{v}). \tag{4}$$

We consider the four desirable properties of a mechanism that is standard in mechanism design [Parkes, 2001]: Dominant Strategy Incentive Compatibility (DSIC), Efficiency, Weak Budget Balance (WBB), and Individual Rationality (IR). These properties are known to be simultaneously achievable in combinatorial auctions, but we will see that such positive results do not carry over to trading networks. We formally define each property in the following, but note that, under the Bayesian setting, we can discuss both *ex post* properties (which hold surely for any $v \in \mathcal{V}$) and *ex ante* or *interim* properties (which hold in expectation with respect to the distribution q over $\mathcal{V}$).

DSIC is an *ex post* property and ensures that the best strategy of each player is truthfully revealing its type regardless of the strategies of the other players. This gives a clear course of action for each player. Formally, we say that a mechanism (ϕ, τ, π) for a trading network $\mathcal{T}^+(\mathcal{V})$ is DSIC if the profile of truthful strategies form a dominant-strategy equilibrium:

$$u_i((v_i, v_{-i}); (\phi, \tau, \pi), \mathcal{T}^+(v))$$
$$\geq u_i((v_i', v_{-i}); (\phi, \tau, \pi), \mathcal{T}^+(v)), \forall (v, v_i') \in \mathcal{V} \times \mathcal{V}_i \tag{5}$$

where v_{-i} is the strategy profile of players except i. We will discuss the corresponding *ex ante* property of Bayesian Nash Incentive-Compatibility (BNIC) only in relation to the prior work, but BNIC only ensures that a truthful player can maximize its *expected* utility under the additional condition that other players are truthful.

Efficiency is also an *ex post* property and ensures that the trades are allocated in a way that they maximize total value. Following the definition without IP in (1), we say that a mechanism (ϕ, τ, π) for $\mathcal{T}^+(\mathcal{V})$ is Efficient if

$$\sum_{i\in\mathcal{N}} v_i(\phi(v)) \geq \sum_{i\in\mathcal{N}} v_i(\Phi), \forall (v, \Phi) \in \mathcal{V} \times 2^\Omega. \tag{6}$$

We will not discuss the corresponding *ex ante* property.

WBB ensures non-negative utility of IP. We say that a mechanism (ϕ, τ, π) for $\mathcal{T}^+(\mathcal{V})$ is *ex post* WBB if

$$u_\mathrm{IP}(v; (\phi, \tau, \pi), \mathcal{T}^+(v)) \geq 0, \forall v \in \mathcal{V} \tag{7}$$

and *ex ante* WBB if

$$\mathbb{E}\left[u_\mathrm{IP}(v; (\phi, \tau, \pi), \mathcal{T}^+(v))\right] \geq 0, \tag{8}$$

where the expectation is with respect to the distribution of v. That is, if the players are truthful (which is guaranteed with DSIC), the utility of IP is non-negative surely under *ex post* WBB and in expectation under *ex ante* WBB.

IR ensures that every player gets non-negative utility. IR can also be *ex post* or *ex ante*. However, when players know their own types, each player should require non-negative expected utility *given its own type*, and such a property is called *interim* IR. Formally, we say that a mechanism (ϕ, τ, π) for $\mathcal{T}^+(\mathcal{V})$ is *ex post* IR if

$$u_i(v; (\phi, \tau, \pi), \mathcal{T}^+(v)) \geq 0, \forall v \in \mathcal{V} \tag{9}$$

for each $i \in \mathcal{N}$ and *interim* IR if

$$\mathbb{E}\left[u_i(v; (\phi, \tau, \pi), \mathcal{T}^+(v)) \mid v_i\right] \geq 0, \forall v_i \in \mathcal{V}_i \qquad (10)$$

for each $i \in \mathcal{N}$. That is, if the players are truthful (guaranteed with DSIC), a truthful player gets non-negative utility surely under *ex post* IR and in expectation under *interim* IR.

Note that *ex post* properties are more desirable than the corresponding *ex ante* properties. For example, maximizing expected utility may not be the objective of risk-sensitive players. Also, the optimality of the truthful strategy under BNIC relies on the truthfulness of the other players, which is not required under DSIC.

3.2 No Payment Between Players

As far as the utilities of the forms in (3)-(4) are concerned, we do not lose generality by ignoring the payment among players (i.e., $\pi(\omega; \hat{v}) = 0, \forall(\omega, \hat{v})$). Intuitively, if there is a payment $p(\omega)$ from a buyer to a seller, we can let the buyer pay $p(\omega)$ to IP and let IP pay $p(\omega)$ to the seller without changing the net-payment of the two players and IP. Formally, we prove the following theorem in [Osogami *et al.*, 2022].

Theorem 1. *For any mechanism (ϕ, τ, π) and a payment rule π' for $\mathcal{T}^+(\mathcal{V}) = (\mathcal{N}^+, \Omega, \mathcal{V})$, there exists a rule of payment to IP τ' such that each of DSIC, Efficiency, WBB (either ex ante or ex post), and IR (either ex ante or interim) is satisfied with (ϕ, τ, π) iff it is satisfied with (ϕ, τ', π').*

Theorem 1 simplifies our analysis, since it implies that we only need to consider mechanisms with no payment between players (i.e., let $\pi' = \pi_0$, where π_0 allocates zero payment to any trade) as long as the desirable properties are the ones that depend only on the utility. Namely, for any set of properties that depend on the utility, there exists a mechanism with no payment between players that achieves those properties if and only if there exists a mechanism (with an arbitrary payment rule) that achieves those properties. Hence, we will only consider designing a mechanism (ϕ, τ, π_0) without payment between players and denote such a mechanism by (ϕ, τ).

3.3 Impossibility

We show that, for a discrete set of valuations, no mechanism (ϕ, τ) can achieve all of the four properties *ex post* for trading networks with a single potential trade except for trivial cases. Theorem 1 then implies that no mechanisms (ϕ, τ, π), with any payment rule π, can simultaneously achieve those properties except for trivial cases. Formally,

Definition 1. *We say that a trading network $(\mathcal{N}^+, \{\omega\}, \mathcal{V})$ with a single trade between two players, $|\mathcal{N}| = 2$, has non-trivial $\mathcal{V}$ if it satisfies both of the following conditions: i) there exists $v \in \mathcal{V}$ such that $\sum_{j \in \mathcal{N}} v_j(\{\omega\}) < 0$, and ii) for any $v_i \in \mathcal{V}_i$, there exists $v_{-i} \in \mathcal{V}_{-i}$ such that $\sum_{j \in \mathcal{N}} v_j(\{\omega\}) > 0$, where $\{i, -i\} \in \mathcal{N}$.*

In other words, $\mathcal{V}$ is said to be trivial if at least one of the two conditions in Definition 1 are violated. When condition i) is violated, the trade should always be conducted (i.e., $\phi(\hat{v}) = \{\omega\}, \forall \hat{v} \in \mathcal{V}$) for Efficiency to hold, and all of the four properties are trivially satisfied with a constant payment rule (see [Osogami *et al.*, 2022]). When condition ii) is violated, there exists a player i who has the type v_i that makes it impossible to make the social welfare, $v_i(\{\omega\}) + v_{-i}(\{\omega\})$, positive no matter what types v_{-i} that the other player has. Such player i should not participate in the trading network from the perspective of social welfare, even if all of the four properties may be achieved with some mechanisms.

For other non-trivial cases, we have impossibility:

Theorem 2. *For any trading network $(\mathcal{N}^+, \{\omega\}, \mathcal{V})$ with a single trade between two players having non-trivial $\mathcal{V}$, no mechanisms can achieve all of DSIC, Efficiency, ex post WBB, and ex post IR.*

Proof. We outline a proof here (see [Osogami *et al.*, 2022] for a complete proof). For each player, consider the two types that respectively give the lowest and highest valuation on ω. We prove that no mechanisms can achieve the four properties by showing that we cannot satisfy all of the necessary conditions associated with the players of those two types. □

[Myerson and Satterthwaite, 1983] show similar impossibility but assumes that the distribution over $\mathcal{V}$ has absolutely continuous density, which does not hold e.g. for discrete $\mathcal{V}$ [Othman and Sandholm, 2009]. Our impossibility theorem does not require such an assumption, but ours is with DSIC and hence weaker in that respect than Myerson-Satterthwaite's, which establishes impossibility with BNIC. In Section 6, we will experiment with instances of trading networks where satisfying the four properties *ex post* is impossible due to Theorem 2.

4 Computing Mechanisms

The impossibility suggested by Theorem 2, together with the incentives for players to untruthfully report their valuations, motivates us to study the mechanisms that achieve weaker properties of *ex ante* WBB and *interim* IR in addition to DSIC and Efficiency. These weaker properties are still meaningful in practice and are sufficient for risk-neutral players.

To guarantee DSIC and Efficiency, we rely on the Groves mechanism. For any trading network $(\mathcal{N}^+, \Omega, \mathcal{V})$, a mechanism (ϕ, τ') is called a Groves mechanism if it can be represented by the use of a pivot rule, $h \equiv (h_i)_{i \in \mathcal{N}}$ where $h_i : \mathcal{V}_{-i} \to \mathbb{R}$, as follows:

$$\phi(v) = \phi^\star(v) \in \operatorname*{argmax}_{\Phi \subseteq \Omega} \sum_{i \in \mathcal{N}} v_i(\Phi) \qquad (11)$$

$$\tau'_i(v) = h_i(v_{-i}) - \sum_{j \neq i} v_j(\phi^\star(v)) \qquad (12)$$

for any $v \in \mathcal{V}$. Here, the pivot rule h_i determines the payment, in addition to the second term of the right-hand side of (12), in a way that it depends only on the types of other players (and not the type of player i). The Groves mechanism is guaranteed to satisfy Efficiency and DSIC, and we may arbitrarily choose h to achieve other properties. We refer to h as a Groves mechanism when it does not cause confusion.

There may or may not exist an h that satisfies both *ex ante* WBB and *interim* IR. When it does not exist, one could seek to find the h that minimizes the violation of these conditions.

However, here we focus on studying whether such an h exists and finding one if it does. Also, when there are multiple h that satisfy both *ex ante* WBB and *interim* IR, we prefer the one with stronger budget balance (i.e., the total payment to IP should be as close to zero as possible).

Therefore, we proposed to find h as the solution to the following optimization problem:

$$\min_h \sum_{i \in \mathcal{N}} \mathbb{E}\left[h_i(v_{-i})\right] \tag{13}$$

$$\text{s.t.} \sum_{i \in \mathcal{N}} \mathbb{E}\left[h_i(v_{-i})\right] \geq (|\mathcal{N}| - 1) \sum_{i \in \mathcal{N}} \mathbb{E}\left[v_i(\phi^\star(v))\right] \tag{14}$$

$$\mathbb{E}\left[h_i(v_{-i}) \mid v_i\right] \leq \sum_{j \in \mathcal{N}} \mathbb{E}\left[v_j(\phi^\star(v)) \mid v_i\right], \quad \forall v_i \in \mathcal{V}_i, \forall i \in \mathcal{N} \tag{15}$$

where $\mathbb{E}$ is the expectation with respect to the distribution of types. The constraint (14) is on *ex ante* WBB (see (55)-(57) in [Osogami *et al.*, 2022] for the derivation). The constraint (15) is on *interim* IR (see (58)-(60) in [Osogami *et al.*, 2022]). For a given trading network $(\mathcal{N}^+, \Omega, \mathcal{V})$, the right-hand sides of (14)-(15) are constants that can be computed by the use of the allocation rule $\phi^\star$ and the distribution of types. When $\mathcal{N}$ and $\mathcal{V}$ are finite, the optimization problem (13)-(15) is a linear program (LP). The LP has $n\, m^{n-1}$ variables (i.e., $h_i(v_{-i})$ for $v_{-i} \in \mathcal{V}, i \in \mathcal{N}$) and $n\, m + 1$ constraints, when there are $n = |\mathcal{N}|$ players, and each player i has $m = |\mathcal{V}_i|$ types.

Example 2. *Consider the trading network in Example 1, but we now assume that each player has one of two types: the production cost of the seller of type X is C_{S}^X for $X \in \{\mathrm{L}, \mathrm{H}\}$, and the handling cost of the buyer of type Y is C_{B}^Y for $Y \in \{\mathrm{L}, \mathrm{H}\}$. Let $C_i^{\mathrm{L}} < C_i^{\mathrm{H}}$ for $i \in \{\mathrm{S}, \mathrm{B}\}$. By Theorem 2, there is no mechanism that simultaneously achieves the four properties ex post when $c_{\mathrm{S}}^{\mathrm{H}} + c_{\mathrm{B}}^{\mathrm{H}} > 1$ and $c_{\mathrm{S}}^X + c_{\mathrm{B}}^Y < 1$ for $(X, Y) \in \{(\mathrm{H}, \mathrm{L}), (\mathrm{L}, \mathrm{H}), (\mathrm{L}, \mathrm{L})\}$. Under the Groves mechanism, the amount of payment to IP from S of type X and B of type Y is respectively given by*

$$\tau'_S(X, Y) = h_{\mathrm{S}}(Y) - (1 - C_{\mathrm{B}}^Y)\, q^\star(X, Y) \tag{16}$$

$$\tau'_B(X, Y) = h_{\mathrm{B}}(X) + C_{\mathrm{S}}^X\, q^\star(X, Y), \tag{17}$$

where $q^\star(X, Y) \equiv I\{\phi^\star(X, Y) = \Omega\}$ is the number of goods to be traded when the types of S and B are X and Y respectively. Note that the pivot rule of each player depends on the type of the other player and can be determined by solving the LP (13)-(15)*, which involves four variables $h \equiv (h_{\mathrm{S}}(\mathrm{L}), h_{\mathrm{S}}(\mathrm{H}), h_{\mathrm{B}}(\mathrm{L}), h_{\mathrm{B}}(\mathrm{H}))$ and four constraints (see* (66)-(69) *in [Osogami* et al.*, 2022]).*

Although the solution to the LP (13)-(15) gives non-trivial Groves mechanisms that are of practical interest, it is intractable except for small trading networks. For example, i) there can be infinitely many possible types v_i, resulting in infinitely many variables in the LP. Also, ii) one may not know the exact distribution of types, which we need to compute the expectation in (13)-(15). In addition, iii) it may be hard to compute the efficient allocation $\phi^\star(v)$ for a given $v \in \mathcal{V}$.

5 Learning Mechanisms

Among these three challenges, i) and ii) are fundamental, since one cannot even represent the optimization problem with infinitely many variables or without knowing the distribution of types. Here, we propose learning techniques to overcome these challenges. For the computational complexity of $\phi^\star$, efficient algorithms are known under some conditions on $\mathcal{V}$ [Candogan *et al.*, 2021; Iwata *et al.*, 2005].

5.1 Learning Pivot Rules

We now address a less restricted setting where the mechanism designer (IP) has the sample $\mathcal{D}$ of types from the distribution q, rather than the exact knowledge on q. The expectation in (13)-(15) can then be replaced with sample average, resolving Challenge ii). IP may collect such $\mathcal{D}$ by running a Groves mechanism that is not necessarily *ex ante* WBB (hence, IP needs investment). Since players act truthfully under the Groves mechanism, the collected $\mathcal{D}$ follows q. IP may also keep collecting sample while learning mechanisms repeatedly (see [Osogami *et al.*, 2022]).

Recall that the variables in (13)-(15) correspond to the output values of the pivot rule, $h_i(v_{-i})$ for $v_{-i} \in \mathcal{V}_{-i}$ and $i \in \mathcal{N}$, which constitute the codomain of the functions $h = (h_i)_{i \in \mathcal{N}}$. To deal with infinitely many variables that stem from the functions with infinite codomain, we approximate those functions with machine learning models, $h^\theta \equiv (h_i^{\theta_i})_{i \in \mathcal{N}}$ where $h_i^{\theta_i} : \mathcal{V}_{-i} \to \mathbb{R}$ with parameter θ_i for each i.

Although such machine learning has been successfully applied in the prior work of AMD, in our case, it is not sufficient to replace the expectation with sample average and h with h^θ. In particular, the sample average in (15) is unreliable or cannot be obtained, since the set $\{\tilde{v} \in \mathcal{D} \mid \tilde{v}_i = v_i\}$ is often small and can be empty for some v_i. We solve this by learning, for each $i \in \mathcal{N}$, a regressor $\hat{g}_i$ (e.g., Gaussian process regressor) that maps $v_i \in \mathcal{V}_i$ to $\sum_{j \in \mathcal{N}} \mathbb{E}[v_j(\phi^\star(v)) \mid v_i]$, where the training data is $\hat{D}_i \equiv \{(\tilde{v}_i, \sum_{j \in \mathcal{N}} \tilde{v}_j(\phi^\star(\tilde{v})))\}_{\tilde{v} \in \mathcal{D}}$.

We can now reduce (13)-(15) to the following constrained non-linear optimization and learn a locally optimal θ with the augmented Lagrangian method:

$$\min_\theta \sum_{\tilde{v} \in \mathcal{D}} \sum_{i \in \mathcal{N}} h_i^{\theta_i}(\tilde{v}_{-i}) \tag{18}$$

$$\text{s.t.} \sum_{\tilde{v} \in \mathcal{D}} \sum_{i \in \mathcal{N}} h_i^{\theta_i}(\tilde{v}_{-i}) \geq (|\mathcal{N}| - 1) \sum_{\tilde{v} \in \mathcal{D}} \sum_{i \in \mathcal{N}} \tilde{v}_i(\phi^\star(\tilde{v})) \tag{19}$$

$$\frac{1}{|\{v \in \mathcal{D} \mid v_i = \tilde{v}_i\}|} \sum_{v \in \mathcal{D} | v_i = \tilde{v}_i} h_i^{\theta_i}(\tilde{v}_{-i}) \leq \hat{g}_i(\tilde{v}_i), \quad \forall \tilde{v}_i \in \mathcal{D}_i, \forall i \in \mathcal{N} \tag{20}$$

where $\mathcal{D}_i \equiv \{v_i \in \mathcal{V}_i \mid \exists \tilde{v} \in \mathcal{D} \text{ s.t. } \tilde{v}_i = v_i\}$ is the set of the types of player i that appear in $\mathcal{D}$. More precisely, we use a lower confidence bound for $\hat{g}(\tilde{v}_i)$ in (20) and an upper confidence bound for the right-hand side of (19). The learning problem (18)-(20) matches the optimization problem (13)-(15) in the limit of $|\mathcal{D}| \to \infty$ when the law of large numbers applies to the sample averages (e.g., $\tilde{v}_i(\phi^\star(\tilde{v}))$ has finite variance), the regressors are asymptotically consistent (i.e., $\hat{g}_i(v_i) \to \mathbb{E}[v_j(\phi^\star(v)) \mid v_i]$), the optimal h is in the class of h^θ (i.e., realizable), the support of the distribution q covers the whole $\mathcal{V}$, and $\mathcal{V}$ is finite.

While the use of h^θ alleviates Challenge i), the issue of computational complexity still remains. Namely, each $h_i^{\theta_i}$ takes types v_{-i} as its input, but each type, $v_j : 2^\Omega \to \mathbb{R}$, is a function. A question is how to represent those functions. One may for example assume that the types are in a certain parametric family of functions (e.g., neural networks) and give their parameters as the input to $h_i^{\theta_{-i}}$ [Faccio *et al.*, 2021].

We take an alternative approach of representing a function with its output values [Harb *et al.*, 2020]. Observe that $v_i : 2^\Omega \to \mathbb{R}$ is fully characterized by a $2^{|\Omega|}$-dimensional valuation-vector, $(v_i(\Phi))_{\Phi\subseteq\Omega}$, that represents the output values for all of the possible inputs. Each $h_i : \mathcal{V}_{-i} \to \mathbb{R}$ is then a function that maps $|\mathcal{N}|-1$ vectors, each having $2^{|\Omega|}$ dimensions, to a real number; namely, $h_i : \mathbb{R}^{2^{|\Omega|}\,(|\mathcal{N}|-1)} \to \mathbb{R}$ (see Figure 3(a) in [Osogami *et al.*, 2022]).

5.2 Variable Reduction

The valuation-vector has exponentially many dimensions and does not fully resolve Challenge i). We now show, based on game theoretic analysis, that some of the features of the valuation-vector are particularly important to achieve some of the desirable properties. We will then form a low dimensional valuation-vector with these important features to better address Challenge i).

To this end, we have designed special Groves mechanisms:

$$h_i^{\mathrm{WBB}}(v) \equiv \max_{\Phi\subseteq\Omega} \sum_{j\neq i} v_j(\Phi), \forall i \in \mathcal{N} \tag{21}$$

$$h_i^{\mathrm{IR}}(v) \equiv \max_{\Phi\subseteq\Omega_{-i}} \sum_{j\neq i} v_j(\Phi), \forall i \in \mathcal{N}, \tag{22}$$

where $\Omega_{-i} \equiv \Omega \setminus (\Omega_{i\to} \cup \Omega_{\to i})$ is the set of trades where player i is neither the seller nor the buyer. We will utilize these mechanisms, because it can be shown that h^{WBB} is always *ex post* WBB, and h^{IR} is *ex post* IR when no players reduce the social welfare, where a player reducing the social welfare is defined as follows:

Definition 2. *In a trading network $(\mathcal{N}^+, \Omega, \mathcal{V})$, a player $i \in \mathcal{N}$ with type $v_i \in \mathcal{V}_i$ is called a negative player if there exists $v_{-i} \in \mathcal{V}_{-i}$ such that*

$$\max_{\Phi\subseteq\Omega_{-i}} \sum_{j\neq i} v_j(\Phi) > \max_{\Phi\subseteq\Omega} \sum_{j\in\mathcal{N}} v_j(\Phi). \tag{23}$$

Formally,

Theorem 3. *For any trading network $(\mathcal{N}^+, \Omega, \mathcal{V})$, the Groves mechanism with h^{WBB} in (21) satisfies DSIC, Efficiency, and ex post WBB, and the Groves mechanism with h^{IR} in (22) satisfies DSIC, Efficiency, and ex post IR if $(\mathcal{N}^+, \Omega, \mathcal{V})$ has no negative players.*

Proof. We outline a proof here (see [Osogami *et al.*, 2022] for a complete proof). Since DSIC and Efficiency are guaranteed with the Groves mechanism, we show that the payment to IP is non-negative for each player, which implies *ex post* WBB, and that the utility of each player is non-negative, which implies *ex post* IR, when there are no negative players. □

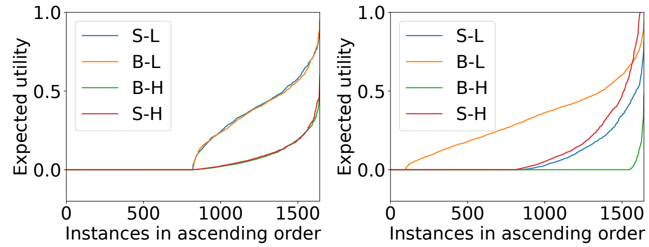

(a) Full variables (b) Reduced variables

Figure 1: The expected utility of each player (S-L is the seller of low cost; B-H is the buyer of high costs; S-H and B-L are defined analogously), plotted in ascending order, under the mechanisms computed by our approach with (a) full or (b) reduced variables.

Since h_i^{WBB} and h_i^{IR}, respectively, achieve *ex post* WBB and *interim* IR, we combine them with the expectation to simultaneously achieve the weaker properties of *ex ante* WBB and *interim* IR. Observe that an arbitrary combination of two functions can be represented by a function that takes the input of the two functions. Since

$$h_i^{\mathrm{WBB}}(v) = \sum_{j\neq i} v_j(\Phi^\star), \text{ where } \Phi^\star \in \operatorname*{argmax}_{\Phi\subseteq\Omega} \sum_{j\neq i} v_j(\Phi)$$

$$h_i^{\mathrm{IR}}(v) = \sum_{j\neq i} v_j(\Phi^{\star,i}), \text{ where } \Phi^{\star,i} \in \operatorname*{argmax}_{\Phi\subseteq\Omega_{-i}} \sum_{j\neq i} v_j(\Phi),$$

we see that $(v_j(\Phi^\star))_{j\neq i}$ and $(v_j(\Phi^{\star,i}))_{j\neq i}$ are the input of h_i^{WBB} and h_i^{IR}, respectively, and hence are the important features of the valuation-vector. We thus seek to find the optimal function within the class of functions that take those important features as input (see Figure 3(b) in [Osogami *et al.*, 2022]). We may also consider a larger class of functions by allowing additional (e.g., random) elements of the original valuation-vector as input, trading off the quality of approximation against computational complexity.

6 Experiments

To show the relevance of the proposed approaches, we conduct experiments on a large number of trading networks where it is impossible to achieve all of the four desirable properties *ex post*. Specifically, we generate 1,642 non-trivial instances of trading networks uniformly at random in the setting of Example 2, which involves a single potential trade between a seller and a buyer (see [Osogami *et al.*, 2022] for details of how those instances are generated). By Theorem 2, there exist no mechanisms that satisfy all of the four desirable properties *ex post* for those non-trivial instances. We apply the proposed approaches to these trading networks and study whether they find the mechanisms that satisfy the four properties if WBB is *ex ante* and IR is *interim*. We run all experiments on a workstation having 64 GB memory and 4.0 GHz CPU.

We first validate our computational approach by applying it to each of the randomly generated instances. Here, we solve the optimization problem in (13)-(15) via CPLEX 22.1.

Figure 1(a) shows the expected utility of each player for the 1,642 instances plotted in ascending order. Observe that, for

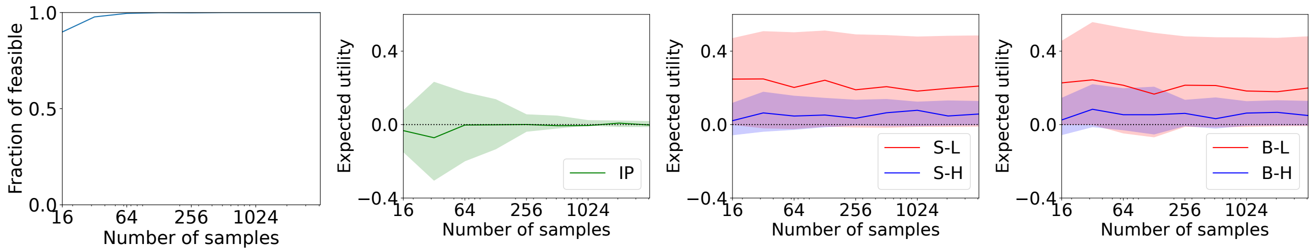

(a) Feasibility (b) Budget balance (c) Supplier utility (d) Buyer utility

Figure 2: The performance of the mechanisms learned by our approach against the number of the samples from the distribution. Panel (a) shows the fraction of the instances for which the learning problem has feasible solutions. For those feasible instances, panel (b) shows the expected budget balance of IP, and panel (c)-(d) shows expected utility of each player. The solid line is the average expected utility, and the shaded area is the [15.9%, 84.1%]-quantile.

all instances, every player (supplier or buyer of any type) expects non-negative utility. The corresponding expected budget balance (utility) of IP is essentially zero for all instances and only shown in Figure 4(a) in [Osogami *et al.*, 2022]. Therefore, the proposed approach finds the mechanism that achieves both *ex ante* WBB and *interim* IR (in addition to DSIC and Efficiency) for every randomly generated instance where achieving the four properties *ex post* is impossible.

This however does not mean that one can *always* achieve *ex ante* WBB and *interim* IR. There indeed exist instances for which *ex ante* WBB or *interim* IR cannot be achieved. Empirically, however, such infeasible instances appear to have zero Lebesgue measure and are not generated in our random process. A caveat is that, when an instance is close to an infeasible instance, our approach tends to find the mechanism that requires a large amount of payment to or from IP (see Figure 5 in [Osogami *et al.*, 2022]).

Next, we validate the effectiveness of our learning approach with a focus on the impact of sample approximation. Here, for each of the 1,642 instances, we take varying number of samples from the distribution of types to solve the learning problem (18)-(20). We use the lower confidence bound of a Gaussian process regressor to evaluate the right-hand side of (20) and add the sample standard deviation to the right-hand side of (19) to form an upper confidence bound. We seek to find an exact h without functional approximation h^θ.

Figure 2 shows (a) the fraction of the instances for which the learning problem has feasible solutions as well as the expected utility of (b) IP, (c) the supplier, and (d) the buyer, under the mechanisms learned with our approach. In (b)-(d), the solid lines show the average over the feasible instances, and the shaded area represents the [15.9%, 84.1%]-quantile, which coincide with the confidence intervals set in (19)-(20).

Figure 2(a) shows that the learning problems do not necessarily have feasible solutions (while the corresponding optimization problems are all feasible). With sufficient amount of data (≥ 512 in this case), however, all of the learning problems become feasible. For those feasible learning problems, our approach learns the mechanisms that satisfy *ex ante* WBB and *interim* IR with high probability (Figure 2(b)-(d)). Although these properties are sometimes violated, they become more frequently satisfied as the sample size increases.

Finally, we study the effectiveness of variable reduction in the settings of the previous experiments. See [Osogami *et al.*, 2022] for how the variables are reduced in these settings.

Figure 1(b) shows the expected utility of each player under the mechanisms *computed* with variable reduction (see Figure 7 in [Osogami *et al.*, 2022] for details). By comparing against the corresponding results in Figure 1(a) without variable reduction, we find different mechanisms depending on whether the variables are reduced or not, as is suggested by the different expected utilities of the players. However, the mechanisms computed with reduced variables still achieve *ex ante* WBB and *interim* IR for all of the random instances.

Figure 8 in [Osogami *et al.*, 2022] shows the performance of the mechanisms *learned* with variance reduction. Similar to Figure 2 (no variable reduction), *ex ante* WBB and *interim* IR are achieved with high probability as long as the learning problems are feasible. Variable reduction, however, reduces the solution space, and the instances that are feasible without variable reduction can become infeasible. Nevertheless, the fraction of feasible instances increases with the sample size (and essentially all of the instances become feasible with the exact distribution, as with the computational approach).

7 Conclusion

We have shown that, by computing or learning appropriate mechanisms, it can be made possible to achieve DSIC, Efficiency, *ex ante* WBB, and *interim* IR in trading networks where achieving these four properties *ex post* is impossible. Since this paper proposes the first AMD approach for trading networks, we have discussed and formalized the problem of mechanism design for trading networks, including fundamental characteristics such as Theorem 1 and Theorem 2.

There are some directions of future work that can be built upon these foundations. For example, while we learn mechanisms from previously collected (offline) data, the prior work has investigated the approaches of learning mechanisms while collecting data in an online manner (see [Osogami *et al.*, 2022]). Such online methods involves the additional challenge of the tradeoff between exploration and exploitation, and it is interesting to study such online methods for trading networks. It is also interesting to study sample complexity of learning mechanisms for trading networks (see [Osogami *et al.*, 2022] for such prior work on combinatorial auctions).

Ethical Statement

There are no ethical issues.

References

[Alon *et al.*, 2021] Tal Alon, Ron Lavi, Elisheva S Shamash, and Inbal Talgam-Cohen. Incomplete information VCG contracts for common agency. *arXiv preprint arXiv:2105.14998*, 2021.

[Candogan *et al.*, 2021] Ozan Candogan, Markos Epitropou, and Rakesh V. Vohra. Competitive equilibrium and trading networks: A network flow approach. *Operations Research*, 69(1):114–147, 2021.

[Conitzer and Sandholm, 2002] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, page 103–110, August 2002.

[Duetting *et al.*, 2019] Paul Duetting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1706–1715. PMLR, 09–15 Jun 2019.

[Faccio *et al.*, 2021] Francesco Faccio, Louis Kirsch, and Jürgen Schmidhuber. Parameter-based value functions. In *International Conference on Learning Representations*, 2021.

[Harb *et al.*, 2020] Jean Harb, Tom Schaul, Doina Precup, and Pierre-Luc Bacon. Policy evaluation networks. *CoRR*, abs/2002.11833, 2020.

[Hatfield and Kominers, 2012] John William Hatfield and Scott Duke Kominers. Matching in networks with bilateral contracts. *American Economic Journal: Microeconomics*, 4(1):176–208, 2012.

[Hatfield *et al.*, 2013] John William Hatfield, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp. Stability and competitive equilibrium in trading networks. *Journal of Political Economy*, 121(5):966–1005, 2013.

[Hatfield *et al.*, 2015] John William Hatfield, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp. Full substitutability in trading networks. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 39–40, 2015.

[Hatfield *et al.*, 2019] John William Hatfield, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp. Full substitutability. *Theoretical Economics*, 14(4):1535–1590, 2019.

[Iwata *et al.*, 2005] Satoru Iwata, Satoko Moriguchi, and Kazuo Murota. A capacity scaling algorithm for m-convex submodular flow. *Mathematical Programming*, 103(1):181–202, May 2005.

[Likhodedov and Sandholm, 2004] Anton Likhodedov and Tuomas Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proceedings of the 19th National Conference on Artifical Intelligence*, AAAI'04, page 232–237. AAAI Press, 2004.

[Likhodedov and Sandholm, 2005] Anton Likhodedov and Tuomas Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proceedings of the 20th National Conference on Artifical Intelligence*, pages 267–274, 01 2005.

[Manisha *et al.*, 2018] Padala Manisha, C V Jawahar, and Sujit Gujar. Learning optimal redistribution mechanisms through neural networks. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems*, page 345–353, July 2018.

[Myerson and Satterthwaite, 1983] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.

[Nisan, 2007] Noam Nisan. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, chapter 9, page 209–241. Cambridge University Press, 2007.

[Osogami *et al.*, 2022] Takayuki Osogami, Segev Wasserkrug, and Elisheva S. Shamash. Mechanism learning for trading networks. *arXiv:2208.09222*, 2022.

[Ostrovsky, 2008] Michael Ostrovsky. Stability in supply chain networks. *American Economic Review*, 98(3):897–923, 2008.

[Othman and Sandholm, 2009] Abraham Othman and Tuomas Sandholm. How pervasive is the myerson-satterthwaite impossibility? In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, IJCAI'09, page 233–238, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc.

[Parkes, 2001] David Christopher Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, University of Pennsylvania, 2001.

[Rahme *et al.*, 2021] Jad Rahme, Samy Jelassi, and S. Matthew Weinberg. Auction learning as a two-player game. In *International Conference on Learning Representations*, 2021.

[Schlegel, 2022] Jan Christoph Schlegel. The structure of equilibria in trading networks with frictions. *Theoretical Economics*, 17(2):801–839, May 2022.

[Tacchetti *et al.*, 2022] Andrea Tacchetti, DJ Strouse, Marta Garnelo, Thore Graepel, and Yoram Bachrach. Learning truthful, efficient, and welfare maximizing auction rules. In *ICLR 2022 Workshop on Gamification and Multiagent Solutions*, 2022.