
Benefits of Permutation-Equivariance in Auction Mechanisms

Tian Qin^{1,2*} Fengxiang He^{2*} Dingfeng Shi² Wenbing Huang^{3,4} Dacheng Tao²

¹University of Science and Technology of China ²JD Explore Academy, JD.com Inc.

³Gaoling School of Artificial Intelligence, Renmin University of China

⁴Beijing Key Laboratory of Big Data Management and Analysis Methods

tqin@mail.ustc.edu.cn, fengxiang.f.he@gmail.com,
shidingfeng@buaa.edu.cn, hwenbing@126.com, dacheng.tao@gmail.com

Abstract

Designing an incentive-compatible auction mechanism that maximizes the auctioneer’s revenue while minimizes the bidders’ ex-post regret is an important yet intricate problem in economics. Remarkable progress has been achieved through learning the optimal auction mechanism by neural networks. In this paper, we consider the popular *additive valuation* and *symmetric valuation* setting; *i.e.*, the valuation for a set of items is defined as the sum of all items’ valuations in the set, and the valuation distribution is invariant when the bidders and/or the items are permuted. We prove that permutation-equivariant neural networks have significant advantages: the permutation-equivariance decreases the expected ex-post regret, improves the model generalizability, while maintains the expected revenue invariant. This implies that the permutation-equivariance helps approach the theoretically optimal *dominant strategy incentive compatible* condition, and reduces the required sample complexity for desired generalization. Extensive experiments fully support our theory. To our best knowledge, this is the first work towards understanding the benefits of permutation-equivariance in auction mechanisms.

1 Introduction

Optimal auction design [30] has wide applications in economics, including computational advertising [18], resource allocation [16], and supply chain [4]. In an auction, every bidder has a private valuation profile over all items, and accordingly, submits her bid profile. An auctioneer collects the bids from all bidders, and determines a feasible item allocation to the bidders as well as the prices the bidders need to pay. Consequently, every bidder receives her utility. From the auctioneer’s perspective, the optimal auction mechanism is required to maximize her *revenue*, defined as the sum of all bidders’ payments. From the aspect of the bidders, the optimal auction mechanism needs to incentivize every bidder to bid their truthful valuation profiles (truthful bidding). This is summarized as the *dominant strategy incentive compatible* (DSIC) condition; *i.e.*, truthful bidding is always the dominant strategy for every bidder [25].

The optimal auction mechanism can be approximated via neural networks [12, 29, 10]. The “approximation error”, or the “distance” to the DSIC condition, is usually measured by the *ex-post regret*, defined as the gap between the bidder’s utility of truthful bidding and the utility when her bid profile is only the best to herself (selfish bidding), while the bid profiles of all other bidders are fixed in both cases [12]. When a bidder’s ex-post regret is 0, truthful bidding is her dominant strategy. Therefore, the optimal auction design can be modeled as a linear programming problem, where the object is to maximize the expected revenue subject to the expected ex-post regret being 0 for all bidders [12].

*Both authors contributed equally.

Another major consideration in learning the optimal mechanism is the generalizability to unseen data, usually measured by the *generalization bound*, *i.e.*, the upper bound of the gap between the expected revenue/ex-post regret and their empirical counterparts on the training data [12].

In this paper, we consider the popular setting of *additive valuation* and *symmetric valuation* [12, 29, 10]. The additive valuation condition defines the valuation for a set of items as the sum of the valuations for all items in this set. The symmetric valuation condition assumes the joint distribution of all bidders' valuation profiles to be invariant when bidders and/or items are permuted. This setting covers many applications in practice. For example, when items are independent, the additive valuation condition holds. Moreover, if the auction is anonymous or the order of the items is not prior-known, the symmetric valuation condition holds.

We demonstrate that permutation-equivariant models have significant advantages in learning the optimal auction mechanism as follows. **(1)** We prove that the permutation-equivariance in auction mechanisms decreases the expected ex-post regret while maintaining the expected revenue invariant. Conversely, and equivalently, the permutation-equivariance promises a larger expected revenue, when the expected ex-post regret is fixed. **(2)** We show that the permutation-equivariance of auction mechanisms reduces the required sample complexity for desirable generalizability. We prove that the $l_{\infty,1}$ -distance between any two mechanisms in the mechanism space decreases when they are projected to the permutation-equivariant mechanism (sub-)space. This smaller distance implies a smaller covering number of the permutation-equivariant mechanism space, which further leads to a small generalization bound [12].

We further provide an explanation for the learning process of non-permutation-equivariant neural networks (NPE-NNs). In learning the optimal auction mechanism by an NPE-NN, we show that an extra positive term exists in the quadratic penalty of the ex-post regret based on the result **(1)**. This term serves as a regularizer to penalize the “non-permutation-equivariance”. Moreover, this regularizer also interferes the revenue maximization, and thus affects the learning performance of NPE-NNs. This further explains the advantages of permutation-equivariance in auction design.

Experiments in extensive auction settings are conducted to verify our theory. We design permutation-equivariant versions of RegretNet (RegretNet-PE and RegretNet-test) by projecting the RegretNet [12] to the permutation-equivariant mechanism space in the training and test stage respectively. The empirical results show that permutation-equivariance helps: (1) significantly improve the revenue while maintain the same ex-post regret; (2) record the same revenue with a significantly lower ex-post regret; and (3) narrow the generalization gaps between the training ex-post regret and its test counterpart. These results fully support our theory.

Related works. Myerson completely solves the optimal auction design problem in one-item auctions [23]. However, solutions are not clear when the number of bidders/items exceeds one [12]. Initial attempts have been presented on the characterization of optimal auction mechanisms [20, 26, 9] and algorithmic solutions [1, 2, 27]. Remarkable advances have been made in the required sample complexity for learning the optimal auction mechanism in various settings, including single-item auctions [5, 21, 16], multi-item single-bidder auctions [11], combinatorial auctions [3, 32], and allocation mechanisms [24]. Machine learning-based auction design (automated auction design) have obtained considerable progress [6, 7, 31]. The optimal auction design is modeled as a linear programming problem [6, 7]. However, early works suffer from the scalability issues that the number of the constraints grows exponentially when the bidder number and the item numbers increase. To address this issue, recent works propose to learn the optimal auction mechanism by deep learning. RegretNet is designed for multi-bidder and multi-item settings [12]. Then, RegretNet is developed to meet more restrictive constraints, such as the budget condition [14] and the certifying strategyproof condition [8]. Rahme *et al.* [29] propose the first equivariant neural network-based auction mechanism design method with significant empirical advantages. Ivanov *et al.* [17] propose a RegretFormer which (1) introduces attention layers to RegretNet to learn permutation-equivariant auction mechanisms, and (2) adopts a new interpretable loss function to control the revenue-regret trade-off. Duan *et al.* [10] extend the applicable domain to contextual settings. All these works make remarkable contributions in designing new algorithms from the empirical aspect only. However, the theoretical foundations are still elusive. To our best knowledge, our paper is the first work on theoretically studying the benefits of permutation equivariance in auction design via deep learning.

2 Notations and Preliminaries

Auction. Suppose n bidders are bidding m items in an auction. Every bidder i has her bidder-context (feature) $x_i \in \mathcal{X}$, while every item j is associated with its item-context (feature) $y_j \in \mathcal{Y}$. The bidder i has a private valuation $v_{ij} \in \mathcal{V} \subset \mathbb{R}_{\geq 0}$ for the item j , which is sampled from a conditioned distribution $\mathbb{P}(\cdot | x_i, y_j)$. The value profile $v_i = (v_{i1}, \dots, v_{im})$ is unknown to the auctioneer. For the simplicity, we define $x = (x_1^T, \dots, x_n^T)^T$, $y = (y_1, \dots, y_m)$, $v = (v_1^T, \dots, v_n^T)^T$, $v_{-i} = (v_1^T, \dots, v_{i-1}^T, v_{i+1}^T, \dots, v_n^T)^T$, and $(v'_i, v_{-i}) = (v_1^T, \dots, (v'_i)^T, \dots, v_n^T)^T$.

Every bidder submits a bid profile b_i to the auctioneer according to her valuation profile. Then, the auctioneer determines a feasible item allocation $g(b, x, y)$ and corresponding payments $p(b, x, y)$ as per an auction mechanism (g, p) . Consequently, every bidder receives her utility as

$$u_i(v_i, b, x, y) = \sum_{j=1}^m g_{ij}(b, x, y) \cdot v_{ij} - p_i(b, x, y).$$

The auction mechanism (g, p) consists of an allocation rule $g : \mathbb{R}^{n \times m} \times \mathcal{X}^n \times \mathcal{Y}^m \rightarrow \mathbb{R}^{n \times m}$ and a payment rule $p : \mathbb{R}^{n \times m} \times \mathcal{X}^n \times \mathcal{Y}^m \rightarrow \mathbb{R}^{n \times m}$, where g_{ij} is the probability of allocating item j to the bidder i , and $p_i = \sum_{j=1}^m p_{ij}$ is the price that the bidder i should pay. To avoid allocating an item over once, the allocation rule is constrained such that $\sum_{i=1}^n g_{ij}(b, x, y) \leq 1$ for all $j \in [m]$. Every v_i in our notations can be replaced by b_i . Thus, we can define the similar notations: $b_{-i} = (b_1^T, \dots, b_{i-1}^T, b_{i+1}^T, \dots, b_n^T)^T$, $(b_i, v_{-i}) = (v_1^T, \dots, (b_i)^T, \dots, v_n^T)^T$ and $(v_i, b_{-i}) = (b_1^T, \dots, (v_i)^T, \dots, b_n^T)^T$.

Optimal auction mechanism. An auction mechanism (g, p) is defined to be *dominant strategy incentive compatible* (DSIC), if truthful bidding is always a dominant strategy of every bidder; i.e.,

$$u_i(v_i, (v_i, b_{-i}), x, y) \geq u_i(v_i, b, x, y),$$

for all $i \in [n]$, $v, b \in \mathcal{V}^{n \times m}$, $x \in \mathcal{X}^n$ and $y \in \mathcal{Y}^m$. In addition, an auction mechanism (g, p) is called *individually rational* (IR), if for any bidder-contexts $x \in \mathcal{X}^n$, any item-contexts $y \in \mathcal{Y}^m$, any bidder $i \in [n]$ $x \in \mathcal{X}^n$, valuation profile and bid profile $v, b \in \mathcal{V}^{n \times m}$, truthful bidding always leads to a non-negative utility, i.e.,

$$u_i(v_i, (v_i, b_{-i}), x, y) \geq 0.$$

If an auction mechanism is DSIC and IR, a rational bidder with an obvious dominant strategy will play it (bidding truthfully). Moreover, an optimal auction mechanism is required to maximize the auctioneer's expected revenue $rev = \mathbb{E}_{(v, x, y)} [\sum_{i=1}^n p_i(v, x, y)]$.

Auction design. The ex-post regret $reg_i(v, x, y)$ for the bidder i is defined as

$$\max_{b'_i \in \mathcal{V}^m} u_i(v_i, (b'_i, v_{-i}), x, y) - u_i(v_i, v, x, y).$$

An auction mechanism (g, p) is DSIC, if and only if $\sum_{i=1}^n reg_i(v, x, y) = 0$ for any value profile $v \in \mathcal{V}^{n \times m}$, bidder-context $x \in \mathcal{X}^n$, and item-context $y \in \mathcal{Y}^m$. Suppose the payment rule p satisfies $p_i(b, x, y) \leq \sum_{j=1}^m g_{ij}(b, x, y) b_{ij}$, which implies that each bidder has a non-negative utility. Then, the auction design can be modeled as a linear programming problem that maximizes the expected revenue $\mathbb{E}_{(v, x, y)} [\sum_{i=1}^n p_i(v, x, y)]$ subject to the expected ex-post regret $\mathbb{E}_{(v, x, y)} [\sum_{i=1}^n reg_i(v, x, y)] = 0$. Without loss of generality, the ex-post regret may refer to the average of all bidders' ex-post regrets.

Definition 2.1. Suppose the network's parameter is ω , and the bidder i 's empirical payment and ex-post regret are defined as $\frac{1}{L} \sum_{l=1}^L p_i^\omega(v^{(l)}, x^{(l)}, y^{(l)})$, and $reg_i(\omega) = \frac{1}{L} \sum_{l=1}^L reg_i^\omega(v^{(l)}, x^{(l)}, y^{(l)})$, where the sample set $\{(v^{(l)}, x^{(l)}, y^{(l)})\}_{l=1}^L$ is *i.i.d.* sampled from the following prior distribution, $\mathbb{P}(v, x, y) = \prod_{i,j=1}^{n,m} \mathbb{P}(v_{ij} | x_i, y_j) \mathbb{P}_{X_i}(x_i) \mathbb{P}_{Y_j}(y_j)$.

Equivariant mapping. We define a mapping f as *G-equivariant* if $\psi_g \circ f = f \circ \rho_g$ for two chosen group linear representations ρ and ψ and any g in group G .

Definition 2.2 (Permutation-Equivariant Mapping). A permutation-equivariant mapping is defined to be $f : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^{n \times m}$ that for any instance $x \in \mathbb{R}^{n \times m}$, and permutation matrices $\sigma_n \in \mathbb{R}^{n \times n}$ and $\sigma_m \in \mathbb{R}^{m \times m}$, we have $f(\sigma_n x \sigma_m) = \sigma_n f(x) \sigma_m$.

In this paper, we consider the bidder-permutation $\sigma_n \in \mathbb{R}^{n \times n}$ and item-permutation $\sigma_m \in \mathbb{R}^{m \times m}$. Specifically, we define a mapping f is bidder-symmetric or item-symmetric, if $f(\sigma_n x) = \sigma_n f(x)$ or $f(x\sigma_m) = f(x)\sigma_m$, respectively. Moreover, we define an auction mechanism (g, p) as bidder-symmetric or item-symmetric, if the allocation rule g and the payment rule p are both bidder-symmetric or item-symmetric.

Orbit averaging. For any feature mapping $f : \mathcal{F} \rightarrow \mathcal{G}$, the orbit averaging \mathcal{Q} on f is defined as $\mathcal{Q}f = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \psi_g^{-1} \circ f \circ \rho_g$, where ρ and ψ are two chosen group representations acting on the feature spaces \mathcal{F} and \mathcal{G} , respectively. Orbit averaging can project any mapping to be equivariant:

Proposition 2.3. *Orbit averaging \mathcal{Q} is a projection to the equivariant mapping space $\{f : \psi \circ f = f \circ \rho\}$, i.e., $\psi \circ \mathcal{Q}f = \mathcal{Q}f \circ \rho$ and $\mathcal{Q}^2 = \mathcal{Q}$. In particular, if f is already equivariant, then $\mathcal{Q}f = f$.*

Moreover, $\mathcal{Q}u$ and $\mathcal{Q}reg$ refer to the utility and the ex-post regret induced by $\mathcal{Q}g$ and $\mathcal{Q}p$. For the simplicity, we denote the orbit averagings that modify the auction mechanism to be bidder-symmetric, item-symmetric, and bidder/item-symmetric by bidder averaging \mathcal{Q}_1 , item averaging \mathcal{Q}_2 , and bidder-item aggregated averaging \mathcal{Q}_3 . Besides, a detailed proof of the feasibility of the projected mechanisms can be found in Appendix A.1.

Hypothesis complexity. The generalizability to unseen data is usually measured by the generalization bound, which depends on the hypothesis set’s complexity. To characterize the complexity of the hypothesis set, we introduce the following definitions of *covering number* $\mathcal{N}_{\infty,1}$ and its corresponding distance $l_{\infty,1}$. Based on the covering number, we can obtain a generalization bound in Theorem 3.6.

Definition 2.4 ($l_{\infty,1}$ -distance). Let \mathcal{X} be a feature space and \mathcal{F} a space of functions from \mathcal{X} to \mathbb{R}^n . The $l_{\infty,1}$ -distance on the space \mathcal{F} is defined as $l_{\infty,1}(f, g) = \max_{x \in \mathcal{X}} (\sum_{i=1}^n |f_i(x) - g_i(x)|)$.

Definition 2.5 (Covering number). Covering number $\mathcal{N}_{\infty,1}(\mathcal{F}, r)$ is the minimum number of balls with radius r that can cover \mathcal{F} under $l_{\infty,1}$ -distance.

3 Theoretical Results

This section presents the theoretical results. For simplicity, we view $p = (p_1, \dots, p_n)^T$ as a $n \times 1$ matrix to present the prices the bidders should pay. We first prove that the permutation-equivariance induces the same expected revenue and a smaller expected ex-post regret in Section 3.1. Next in Section 3.2, we prove that the permutation-equivariant mechanism space has a smaller covering number, which promises a smaller required sample complexity and a better generalization. Detailed proofs are omitted from the main text and given in supplementary materials due to space limitation.

3.1 Benefits for Revenue and Ex-Post Regret

In this section, we discuss the benefits for the revenue and the ex-post regret in the conditions of bidder-symmetry and item-symmetry separately, and then discuss the benefits when both of them hold. Based on these results, we also study the learning process of non-permutation-equivariant neural networks for auction design.

3.1.1 Benefits in the Bidder/Item-Symmetry Condition

When the bidders come from the same distribution, the joint valuation distribution f is invariant under bidder-permutation, i.e. $f(\sigma_n v, \sigma_n x, y) = f(v, x, y)$ for any $\sigma_n \in S_n$. Meanwhile, when the items are indistinguishable, the joint distribution f is invariant under item-permutation, i.e., $f(v\sigma_m, x, y\sigma_m) = f(v, x, y)$ for any $\sigma_m \in S_m$. Both conditions do not always hold simultaneously. In this section, we study them separately.

To measure the “non-permutation-equivariance” of the mechanism, we introduce the conception of *regret gap* between the projected mechanism and the original mechanism as below,

$$\Delta(g, p; v, x, y) = \max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n [\mathcal{Q}.u]_i(v_i, (v'_i, v_{-i}), x, y),$$

where v is the valuation profiles, v_i is the valuation profile of bidder i , x is the bidder-context, y is the item-context, and the orbit averaging \mathcal{Q} can be the bidder averaging \mathcal{Q}_1 or the item averaging \mathcal{Q}_2 .

The bidder averaging \mathcal{Q}_1 and the item averaging \mathcal{Q}_2 acting on the allocation rule g and the payment rule p , respectively, are as below,

$$\begin{aligned}\mathcal{Q}_1 g(v, x, y) &= \frac{1}{n!} \sum_{\sigma_n \in S_n} \sigma_n^{-1} g(\sigma_n v, \sigma_n x, y), \quad \mathcal{Q}_1 p(v, x, y) = \frac{1}{n!} \sum_{\sigma_n \in S_n} \sigma_n^{-1} p(\sigma_n v, \sigma_n x, y), \\ \mathcal{Q}_2 g(v, x, y) &= \frac{1}{m!} \sum_{\sigma_m \in S_m} g(v \sigma_m, x, y \sigma_m) \sigma_m^{-1}, \quad \text{and} \quad \mathcal{Q}_2 p(v, x, y) = \frac{1}{m!} \sum_{\sigma_m \in S_m} p(v \sigma_m, x, y \sigma_m).\end{aligned}$$

We thus can prove the following theorem that characterizes the benefits of permutation-equivariance for revenue and ex-post regret in the condition of bidder/item-symmetry.

Theorem 3.1 (Benefits for revenue and ex-post regret in the condition of bidder/item-symmetry). *When the valuation distribution is invariant under permutations of bidders/items, the projected mechanism has the same expected revenue and a smaller expected ex-post regret, that is,*

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q} \cdot p]_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right], \quad \text{and} \quad (1)$$

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \text{reg}_i(v, x, y) \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q} \cdot \text{reg}]_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\Delta \cdot (g, p; v, x, y) \right] \geq 0, \quad (2)$$

where p is the payment rule, reg is the ex-post regret, and \mathcal{Q} is the bidder/item averaging.

A smaller expected ex-post regret implies this mechanism is closer to the *dominant strategy incentive compatible* condition. Conversely, and equivalently, when the expected ex-post regrets are fixed, the projected auction mechanism has a larger expected revenue. For any auction mechanism, in the bidder/item-symmetry condition, we can project it through the bidder/item averaging.

Remark 3.2. The mechanism space can be decomposed into the direct sum of the permutation-equivariant mechanism space $\{\mathcal{M} : \mathcal{Q}\mathcal{M} = \mathcal{M}\}$ and the complementary space $\{\mathcal{N} : \mathcal{Q}\mathcal{N} = 0\}$ [13]. Thus, a mechanism \mathcal{M} has a unique decomposition: $\mathcal{M} = \mathcal{Q}\mathcal{M} + \mathcal{N}$. The pure permutation-equivariant part $\mathcal{Q}\mathcal{M}$ contains all and only the ‘‘permutation-equivariance’’ of the mechanism \mathcal{M} . The pure non-permutation-equivariant part \mathcal{N} is independent from the permutation-equivariance. In this way, we may study the influence of permutation-equivariance by comparing the mechanism \mathcal{M} and its permutation equivariant part $\mathcal{Q}\mathcal{M}$.

3.1.2 Interplay between Bidder-Symmetry and Item-Symmetry.

If the valuation distribution is invariant under both bidder-permutation and item-permutation, we can project the mechanism to be permutation-equivariant with respect to both bidder and item in two steps (by mapping $\mathcal{Q}_1 \circ \mathcal{Q}_2$ or mapping $\mathcal{Q}_2 \circ \mathcal{Q}_1$). Consequently, the projected mechanism has the same expected revenue and a smaller expected ex-post regret. Equivalently, we can also project an auction mechanism to be bidder-symmetric and item-symmetric immediately by the bidder-item aggregated averaging \mathcal{Q}_3 as below,

$$\begin{aligned}\mathcal{Q}_3 g(v, x, y) &= \frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} \sigma_n^{-1} g(\sigma_n v \sigma_m, \sigma_n x, y \sigma_m) \sigma_m^{-1}, \quad \text{and} \\ \mathcal{Q}_3 p(v, x, y) &= \frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} \sigma_n^{-1} p(\sigma_n v \sigma_m, \sigma_n x, y \sigma_m).\end{aligned}$$

We can prove that the bidder-item aggregated averaging \mathcal{Q}_3 is the composition of the orbit averaging operators \mathcal{Q}_1 and \mathcal{Q}_2 , as shown in the following lemma. This lemma shows that the order of \mathcal{Q}_1 and \mathcal{Q}_2 would not influence their composition.

Lemma 3.3. *The bidder-item aggregated averaging is the composition of bidder averaging and item averaging: $\mathcal{Q}_3 = \mathcal{Q}_1 \circ \mathcal{Q}_2 = \mathcal{Q}_2 \circ \mathcal{Q}_1$.*

Based on this lemma, we can prove the following theorem on the benefits of permutation-equivariance for revenue and ex-post regret in the condition of both bidder-symmetry and item-symmetry.

Theorem 3.4 (Benefits for revenue and ex-post regret in the condition of both bidder-symmetry and item-symmetry). *When the valuation distribution is invariant under both item-permutation and bidder-permutation, then the projected mechanism has a same expected revenue and a smaller expected ex-post regret, that is,*

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \mathcal{Q}_3 p_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \text{ and}$$

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \text{reg}_i(v, x, y) \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_3 \text{reg}]_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\Delta_3(g, p; v, x, y) \right] \geq 0,$$

where p is the payment rule, reg is the ex-post regret, and \mathcal{Q}_3 is the bidder-item aggregated averaging.

The difference between bidder-symmetry and item-symmetry is significant in practice. For example, for a symmetric valuation distribution, when the mechanism is already bidder-symmetric but not item-symmetric, we can project it to be item-symmetric to gain an extra benefit from item-symmetry. That means, the two regret gaps induced by \mathcal{Q}_1 and \mathcal{Q}_2 are “additive” as below,

$$\Delta_3(g, p; v, x, y) = \Delta_1(g, p; v, x, y) + \Delta_2(\mathcal{Q}_1 g, \mathcal{Q}_1 p; v, x, y).$$

In general, $\mathbb{E}[\Delta_2(g, p; v, x, y)] \neq \mathbb{E}[\Delta_2(\mathcal{Q}_1 g, \mathcal{Q}_1 p; v, x, y)]$ and thus $\mathbb{E}[\Delta_3(g, p; v, x, y)] \neq \mathbb{E}[\Delta_1(g, p; v, x, y)] + \mathbb{E}[\Delta_2(g, p; v, x, y)]$. Thus, the benefits from bidder-symmetry and item-symmetry are “additive” but not strictly “independent”.

3.1.3 Insights on Training Non-Permutation-Equivariant Mechanism

Because the expected revenue is always the same for the original mechanism and the projected permutation-equivariant mechanism, we only consider the gradient caused by the expected ex-post regret. We can decompose the original expected ex-post regret into the sum of the expected ex-post regret of the projected mechanism and the expectation of the regret gap as below,

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \text{reg}_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_3 \text{reg}]_i(v, x, y) \right] + \mathbb{E}_{(v,x,y)} \left[\Delta_3(g, p; v, x, y) \right].$$

The regret gap $\Delta_3(g, p; \cdot)$ follows from the “non-permutation-equivariance” of the mechanism \mathcal{M} . When the distance $l(\mathcal{M}, \mathcal{QM})$ tends to 0, the regret gap converges to 0. When the auction mechanism has a negligible ex-post regret, the expectation of the regret gap is also close to 0. That means, the mechanism is close to being permutation-equivariant. However, even using a symmetric dataset or adopting data augmentation in training, the learned mechanism will not be permutation-equivariant in general [19]. As a result, to achieve negligible ex-post regret, the non-permutation-equivariant models need to learn more samples to approach permutation-equivariance. That is because the non-permutation-equivariant part (expected regret gap) would mislead the gradient of the expected regret but have no benefit to the expected revenue and the expected ex-post regret.

On the other hand, the regret gap can be viewed as a regularizer in the ex-post regret to penalize the “non-permutation-equivariance” of the mechanism. When the optimizer tries to minimize the ex-post regret, the auction mechanism approaches to be permutation-equivariant. Therefore, if the mechanism achieves a negligible ex-post regret, it is almost to be permutation-equivariant. This result can explain why RegretNet struggles to find permutation-equivariant auction mechanisms [29]. However, in complex settings, it will be harder for non-permutation-equivariant models to approach the negligible ex-post regret. It can explain why the permutation-equivariant models show a significant improvement in complex settings, compared with that they have similar performances in simple settings [10, 17], which shows the great importance of adopting permutation-equivariant models in complex settings.

3.2 Benefits for Generalization

In this section, we study permutation-equivariance from the aspect of generalizability [22, 15], which characterizes the performance gap of a learned mechanism on collected training data and unseen data.

We first study the covering number of the permutation-equivariant mechanism space. Let $\mathcal{U} = \{u^\omega : \omega \in \Omega\}$ and $\mathcal{P} = \{p^\omega : \omega \in \Omega\}$ be the spaces of all possible utilities and payment rules, and

$\mathcal{Q}\mathcal{U} = \{\mathcal{Q}.u : u \in \mathcal{U}\}$ and $\mathcal{Q}\mathcal{P} = \{\mathcal{Q}.p : p \in \mathcal{P}\}$ the spaces of all projected utilities and payment rules. In addition, let $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$ and $\mathcal{N}_{\infty,1}(\mathcal{P}, r)$ be the minimum numbers of balls with radius r that can cover \mathcal{U} and \mathcal{P} under $l_{\infty,1}$ -distance, respectively. We obtain the following result, which indicates the projected permutation-equivariant mechanism space has smaller covering numbers.

Theorem 3.5 (Covering number of the permutation-equivariant mechanism space). *The space of all projected bidder-symmetric mechanisms has smaller covering numbers, that is,*

$$\mathcal{N}_{\infty,1}(\mathcal{Q}_1\mathcal{U}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, r) \quad \text{and} \quad \mathcal{N}_{\infty,1}(\mathcal{Q}_1\mathcal{P}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{P}, r).$$

The space of all projected item-symmetric mechanisms has smaller covering numbers, that is,

$$\mathcal{N}_{\infty,1}(\mathcal{Q}_2\mathcal{U}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, r) \quad \text{and} \quad \mathcal{N}_{\infty,1}(\mathcal{Q}_2\mathcal{P}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{P}, r).$$

Intuitively, the orbit averaging \mathcal{Q} narrows the distance between two mechanisms: $l(\mathcal{Q}\mathcal{M}, \mathcal{Q}\mathcal{M}') \leq l(\mathcal{M}, \mathcal{M}')$, for any two mechanisms. Then, any r -cover \mathcal{A} for space \mathcal{U} or space \mathcal{P} induces an r -cover $\mathcal{Q}\mathcal{A}$ for space $\mathcal{Q}\mathcal{U}$ or space $\mathcal{Q}\mathcal{P}$.

Combining with Lemma 3.3, we have the following results,

$$\begin{aligned} \mathcal{N}_{\infty,1}(\mathcal{Q}_3\mathcal{U}, r) &= \mathcal{N}_{\infty,1}(\mathcal{Q}_1\mathcal{Q}_2\mathcal{U}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{Q}_2\mathcal{U}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, r) \quad \text{and} \\ \mathcal{N}_{\infty,1}(\mathcal{Q}_3\mathcal{P}, r) &= \mathcal{N}_{\infty,1}(\mathcal{Q}_1\mathcal{Q}_2\mathcal{P}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{Q}_2\mathcal{P}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{P}, r). \end{aligned}$$

We then prove that two generalization bounds of permutation-equivariant mechanisms, which characterize the gap between the expected revenue/ex-post regret and their empirical counterparts. Similar generalization results are existing in previous works [10, 12].

Theorem 3.6 (Generalization bounds of permutation-equivariant mechanisms). *If for any bidder, her valuation satisfies that $v_i(S) \leq 1$ for any $S \subset [m]$, then with probability at least $1 - \delta$, we have the following inequalities with $\epsilon \geq \sqrt{\frac{9m^2}{2L} (\log \frac{4}{\delta} + \max\{\log \mathcal{N}_{\infty,1}(\mathcal{P}, \frac{\epsilon}{3}), \log \mathcal{N}_{\infty,1}(\mathcal{U}, \frac{\epsilon}{6})\})}$,*

$$\left| \mathbb{E} \left[\sum_{i=1}^n p_i^\omega(v, x, y) \right] - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) \right| \leq \epsilon, \quad \text{and} \quad (3)$$

$$\left| \mathbb{E} \left[\sum_{i=1}^n \text{reg}_i^\omega(v, x, y) \right] - \sum_{i=1}^n \widehat{\text{reg}}_i(\omega) \right| \leq \epsilon, \quad (4)$$

where L is the number of samples, \mathcal{U} and \mathcal{P} are the spaces of all possible utilities and payment rules.

Equivalently, we can rewrite this result in the form of the sample complexity,

Corollary 3.7. *For any $\epsilon > 0$, $\delta \in (0, 1)$, and mechanism parameter ω , when the sample complexity $L \geq \frac{9m^2}{2\epsilon^2} \left(\log \frac{4}{\delta} + \max \left\{ \log \mathcal{N}_{\infty,1}(\mathcal{P}, \frac{\epsilon}{3}), \log \mathcal{N}_{\infty,1}(\mathcal{U}, \frac{\epsilon}{6}) \right\} \right)$, with probability at least $1 - \delta$, the generalization bounds, eqs. (3) and (4), hold.*

Remark 3.8. Combining Theorem 3.5, we have proved that the permutation-equivariance can improve the generalizability.

4 Experiments

This section presents our experimental results. More details and results are presented in the supplementary materials.

Model architecture. We project RegretNet [12] to the permutation-equivariant mechanism space via employing bidder-item aggregated averaging for the bidder-symmetry and item-symmetry condition. The projected model is called RegretNet-PE. We also project the well-trained RegretNet, called RegretNet-Test. Specifically, RegretNet is an auction mechanism defined as (g^ω, p^ω) , in which both the allocation rule g^ω and the payment rule p^ω are neural networks that consist of three fully-connected layers, and ω is the overall model parameter of the auction mechanism. The detailed architecture is given in the supplementary materials.

Comparison with EquivariantNet. RegretNet uses two feed-forward fully-connected networks to learn the allocation rule and payment rule, respectively. We denote the weight matrix in the

Table 1: Experimental results. " $n \times m$ Uniform" refers that there are n bidders and m items, and the valuations are i.i.d. drawn from the uniform distribution $U[0, 1]$. To simplify, we multiply all results by a factor of 10^5 for the ex-post regret and generalization error (GE).

Method	2×1 Uniform			3×1 Uniform			5×1 Uniform		
	Revenue	Regret	GE	Revenue	Regret	GE	Revenue	Regret	GE
Optimal	0.417	0	-	0.531	0	-	0.672	0	-
RegretNet	0.415	17.4	6.00	0.535	18.3	11.4	0.658	15.9	6.40
RegretNet-Test	0.415	16.3	-	0.535	13.3	-	0.658	6.50	-
RegretNet-PE	0.420	14.6	3.90	0.541	16.4	10.2	0.677	13.2	5.10

Table 2: Experimental results. " $n \times m$ Uniform" refers that there are n bidders and m items, and the valuations are i.i.d. sampled from the uniform distribution $U[0, 1]$.

Method	1×2 Uniform		2×2 Uniform	
	Revenue	Regret	Revenue	Regret
RegretNet	0.562	0.00061	0.870	0.00070
EquivariantNet	0.551	0.00013	0.873	0.00100
RegretNet-Test	0.562	0.00052	0.870	0.00054
RegretNet-PE	0.563	0.00037	0.913	0.00067

layer ℓ as $W^{(\ell)}$. Both EquivariantNet and RegretNet-PE inherit the architecture of RegretNet (with some modifications), but utilize different approaches to realize the permutation-equivariance. EquivariantNet applies parameter-sharing in every layer during training, to constrain $W^{(\ell)}$ to be equivariant. In contrast, RegretNet-PE employs orbit averaging to be permutation-equivariant. Specifically, RegretNet-PE adopts a weight matrix $I_K \otimes W^{(\ell)} (\rho_{g_1}^T \dots \rho_{g_K}^T)^T$ in the first layer, weight matrices $I_K \otimes W^{(\ell)}$ in the following layers, and multiplies a matrix $(\rho_{g_1}^{-1}, \dots, \rho_{g_K}^{-1})$ to the output layer, where K is the scale of the group $G = \{g_1, \dots, g_K\}$, ρ_{g_k} represents the permutation operator on bidders and items, I_K is an identity matrix, and \otimes is the Kronecker product. It is worth noting that RegretNet-PE is only designed for verifying our theory.

Auction settings. We first adopt the two-bidder single-item, two-bidder single-item, three-bidder single-item, and five-bidder single-item settings in the experiments that compare the learned mechanisms with theoretical optimal mechanisms. The optimal auction mechanism for any single-item auction is known [23]. We thus compare the mechanisms learned by our method with the optimal auction mechanisms in the single-item settings. Also, we compare RegretNet-PE and EquivariantNet in the one-bidder, two-item setting, and the two-bidder, two-item setting. Besides, we employ a multivariate uniform distribution $U[0, 1]^m$ to model the bidder valuation profiles. In all settings, we sample 640,000 data points for training and 5,000 points for test. Due to the space limitation, we place the results of two-bidder five-item and five-bidder three-item settings in Appendix B.2.

Model training. We optimize the auction mechanism model via solving the following optimization problem, following the standard settings [12, 29, 10, 17],

$$\mathcal{L}_\rho(\omega, \lambda) = -\frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) + \sum_{i=1}^n \lambda_i \widehat{reg}_i(\omega) + \frac{\rho}{2} \left(\sum_{i=1}^n \widehat{reg}_i(\omega) \right)^2,$$

where $\lambda \in \mathbb{R}^n$ is the Lagrange multiplier and $\rho > 0$ is the factor of the quadratic regularization term. During the training process, the objective function $\mathcal{L}_\rho(\omega, \lambda)$ is minimized via Adam with a learning rate of 0.001 with respect to the model parameter ω and the Lagrange multiplier λ is updated once in every 100 iterations, until the ex-post regret is smaller than 0.001. The regularization factor ρ is set to 1.0 initially and gradually increased along the training process. In calculating the best bid profile v_i' of every bidder i , we first randomly initialize the bid profiles once in training and 1,000 times in test, optimize each of them individually via Adam with the same settings, and take the best one as the approximated best bidding.

Evaluation. We leverage three metrics to evaluate the performance of the auction mechanism, which are: (1) the empirical revenue \widehat{rev} , (2) the empirical ex-post regret averaging across all bidders, *i.e.*, $\widehat{reg} = \frac{1}{n} \sum_{i=1}^n \widehat{reg}_i$, and (3) the generalization error defined on top of regrets, *i.e.*, $GE = |\widehat{reg}_{test} - \widehat{reg}_{train}|$, where \widehat{reg}_{test} and \widehat{reg}_{train} are the empirical ex-post regrets during test and training, respectively.

Computing resource. The experiments are conducted on 1 GPU (NVIDIA[®] Tesla[®] V100 16GB) and 10 CPU cores (Intel[®] Xeon[®] Processor E5-2650 v4 @ 2.20GHz).

Experimental results. We train a RegretNet and a RegretNet-PE on the training data. The well-trained RegretNet is then projected to be permutation-equivariant, denoted as “RegretNet-Test”. The results are collected in Tables 1 and 2.

From Tables 1 and 2, we observe that (1) compared to RegretNet, RegretNet-PE has a significantly higher revenue with a lower ex-post regret, and narrows the generalization gap between the training ex-post regret and its test counterpart; (2) compared to RegretNet, RegretNet-Test receives the same revenue with a significantly lower ex-post regret; and (3) under comparable ex-post regrets, RegretNet-PE has considerably higher revenue than EquivariantNet, while all permutation-equivariant models (RegretNet-Test and RegretNet-PE) can outperform RegretNet. These results show significant benefits of permutation-equivariance on revenue, ex-post regret, and generalizability, which fully supports our theoretical findings in Theorems 3.1, 3.4, and 3.5.

5 Conclusion and future works

In this paper, under *additive valuation* and *symmetric valuation* setting, we study the benefits of permutation-equivariance in auction mechanisms in two aspects: a better performance and a better generalization. First, we prove a smaller expected ex-post regret and the same expected revenue when projecting a mechanism to be permutation-equivariant. Next, we propose the permutation-equivariant mechanism space has a smaller covering number, which promises the permutation-equivariant models a better generalization. Extensive experiments are conducted to verify our theoretical results. Our results help understand the optimal auction mechanisms’ characterization and the learning processes difference between non-equivariant models and equivariant models.

Beyond the additive valuation setting, an interesting direction is to extend our results to other conditions, including the combinatorial and the unit-demand auctions. Meanwhile, the understanding of the difference in the aspect of the training process between non-equivariant models and equivariant models is still elusive.

Social impact. Our results can help understand and design optimal auction mechanisms for symmetric valuation distribution. As a result, our work could inspire more near-optimal auction mechanisms and promote economic growth. No potential negative social impact is identified.

Acknowledgments and Disclosure of Funding

This work was supported in part by the Major Science and Technology Innovation 2030 “New Generation Artificial Intelligence” Key Project (No. 2021ZD0111700), the National Natural Science Foundation of China (No. 62006137), and the Beijing Outstanding Young Scientist Program (No. BJJWZYJH012019100020098).

We sincerely appreciate Hang Yu, Xiaowen Wei, Kaifan Yang, Shaopeng Fu, and Qingsong Zhang for the valuable comments and the anonymous NeurIPS reviewers for the helpful feedback.

References

- [1] S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *The 16th Annual Symposium on Foundations of Computer Science*, 2011.
- [2] S. Alaei, H. Fu, N. Haghpanah, and J. Hartline. The simple economics of approximately optimal auctions. *The 16th Annual Symposium on Foundations of Computer Science*, 2013.
- [3] M.-F. F. Balcan, T. Sandholm, and E. Vitercik. Sample complexity of automated mechanism design. *Neural Information Processing Systems (NeurIPS)*, 2016.

- [4] L. Chen, L. Wang, and Y. Lan. Auction models with resource pooling in modern supply chain management. *Modern Supply Chain Research and Applications*, 2019.
- [5] R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In *ACM symposium on Theory of computing*, 2014.
- [6] V. Conitzer and T. Sandholm. Complexity of mechanism design. *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2002.
- [7] V. Conitzer and T. Sandholm. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *ACM Conference on Electronic Commerce*, 2004.
- [8] M. Curry, P.-Y. Chiang, T. Goldstein, and J. Dickerson. Certifying strategyproof auction networks. *Neural Information Processing Systems (NeurIPS)*, 2020.
- [9] C. Daskalakis, A. Deckelbaum, and C. Tzamos. Strong duality for a multiple-good monopolist. *Econometrica*, 2017.
- [10] Z. Duan, J. Tang, Y. Yin, Z. Feng, X. Yan, M. Zaheer, and X. Deng. A context-integrated transformer-based neural network for auction design. *arXiv preprint arXiv:2201.12489*, 2022.
- [11] S. Dughmi, L. Han, and N. Nisan. Sampling and representation complexity of revenue maximization. *Web and Internet Economic*, 2014.
- [12] P. Dütting, Z. Feng, H. Narasimhan, D. Parkes, and S. S. Ravindranath. Optimal auctions through deep learning. In *International Conference on Machine Learning (ICML)*, 2019.
- [13] B. Elesedy and S. Zaidi. Provably strict generalisation benefit for equivariant models. In *International Conference on Machine Learning (ICML)*, 2021.
- [14] Z. Feng, H. Narasimhan, and D. C. Parkes. Deep learning for revenue-optimal auctions with budgets. In *the 17th International Conference on Autonomous Agents and Multiagent Systems*, 2018.
- [15] F. He and D. Tao. Recent advances in deep learning theory. *arXiv preprint arXiv:2012.10931*, 2020.
- [16] J. Huang, Z. Han, M. Chiang, and H. V. Poor. Auction-based resource allocation for cooperative communications. *IEEE Journal on Selected Areas in Communications*, 2008.
- [17] D. Ivanov, I. Safulin, K. Balabaeva, and I. Filippov. Optimal-er auctions through attention. *arXiv preprint arXiv:2202.13110*, 2022.
- [18] B. J. Jansen and T. Mullen. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business*, 2008.
- [19] C. Lyle, M. van der Wilk, M. Kwiatkowska, Y. Gal, and B. Bloem-Reddy. On the benefits of invariance in neural networks. *arXiv preprint arXiv:2005.00178*, 2020.
- [20] A. M. Manelli and D. R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 2006.
- [21] M. Mohri and A. M. Medina. Learning theory and algorithms for revenue optimization in second-price auctions with reserve. *International Conference on Machine Learning (ICML)*, 2013.
- [22] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [23] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 1981.
- [24] H. Narasimhan and D. C. Parkes. A general statistical framework for designing strategy-proof assignment mechanisms. *Uncertainty in Artificial Intelligence (UAI)*, 2016.
- [25] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007.

- [26] G. Pavlov. Optimal mechanism for selling two goods. *The BE Journal of Theoretical Economics*, 2011.
- [27] G. Popescu. An algorithmic characterization of multi-dimensional mechanisms. *Computing Reviews*, 2012.
- [28] O. Punny, M. Atzmon, H. Ben-Hamu, E. J. Smith, I. Misra, A. Grover, and Y. Lipman. Frame averaging for invariant and equivariant network design. *International Conference on Learning Representations (ICLR)*, 2022.
- [29] J. Rahme, S. Jelassi, J. Bruna, and S. M. Weinberg. A permutation-equivariant neural network architecture for auction design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- [30] T. Roughgarden. Twenty lectures on algorithmic game theory. *Cambridge Books*, 2016.
- [31] T. Sandholm and A. Likhodedov. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 2015.
- [32] V. Syrgkanis. A sample complexity measure with applications to learning optimal auctions. *Neural Information Processing Systems (NeurIPS)*, 2017.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]** Our contributions are stated in the 4th-8th sentences of the abstract, and the 4th-6th paragraphs of the introduction. The scope is described in the related work section.
 - (b) Did you describe the limitations of your work? **[Yes]** We describes the problem setting we considered in the 3rd paragraph of the introduction. Our contributions are for this problem setting.
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** Our work is committed to learn the optimal auction mechanism. No negative societal impact is identified.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]** We have read the ethics review guidelines. We ensure that our paper conforms all of them.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** We state all of our assumptions in the third paragraph of introduction, the first paragraph of preliminary, Theorem 3.1, and Theorem 3.2. Wherever we use an assumption, we always state it again.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** Full proofs are given in the supplemental material.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[No]** We are associated in a industrial organization. The code is currently under review of our organization.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** We mostly follow the setting of the RegretNet. Other details are given in Sections 4 and B.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[No]**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** They are given in Section 4.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]** They are given in Section 4.
 - (b) Did you mention the license of the assets? **[Yes]** They are given in Section 4.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[No]** We are associated in a industrial organization. The code is currently under review of our organization.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[Yes]** The data is synthetic in our data.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[Yes]** The data is synthetic in our data. No personally identifiable information or offensive content is contained.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

A Proofs

This appendix collects all proofs omitted from the main text due to space limitation.

A.1 Proofs about Orbit Averaging

In this section, we prove Proposition 2.3 and the feasibility of the projected mechanisms. For simplicity, we use the notations $\sigma_n(i)$ and $\sigma_m(j)$ to present the ranks of the bidder i and the item j after bidder-permutation σ_n and item-permutation σ_m , respectively.

Proof of Proposition 2.3. By the definition of orbit averaging \mathcal{Q} , we have

$$\mathcal{Q}f \circ \rho_g(x) = \frac{1}{|G|} \sum_{h \in G} \psi_h^{-1} f(\rho_h \rho_g x) = \psi_g \circ \frac{1}{|G|} \sum_{h \in G} \psi_{hg}^{-1} f(\rho_{hg} x) = \psi_g \circ \mathcal{Q}f(x).$$

In addition, if f is equivariant, then we have

$$\mathcal{Q}f = \frac{1}{|G|} \sum_{g \in G} \psi_g^{-1} \circ f \circ \rho_g = \frac{1}{|G|} \sum_{g \in G} \psi_g^{-1} \circ \psi_g \circ f = \frac{1}{|G|} \sum_{g \in G} f = f.$$

Thus, orbit averaging is a projection to equivariant function space and fixes all equivariant functions. In addition, orbit averaging fixes all equivariant functions. That means, every equivariant function can be obtained by orbit averaging. In this sense, every equivariant models are contained in the orbit averaging framework.

Proof of the feasibility of projected mechanisms. We verify all feasibility conditions for the projected mechanisms as follows.

Firstly, for the allocation rule, we have

$$\begin{aligned} \sum_{i=1}^n (\mathcal{Q}_1 g)_{ij}(v, x, y) &= \frac{1}{n!} \sum_{i=1}^n \sum_{\sigma_n \in S_n} g_{\sigma_n^{-1}(i)j}(\sigma_n v, \sigma_n x, y) \\ &= \frac{1}{n!} \sum_{\sigma_n \in S_n} \left[\sum_{i=1}^n g_{\sigma_n^{-1}(i)j}(\sigma_n v, \sigma_n x, y) \right] \leq \frac{1}{n!} \sum_{\sigma_n \in S_n} 1 = 1, \end{aligned}$$

$$\begin{aligned} \sum_{i=1}^n (\mathcal{Q}_2 g)_{ij}(v, x, y) &= \frac{1}{m!} \sum_{i=1}^n \sum_{\sigma_m \in S_m} g_{i\sigma_m^{-1}(j)}(v\sigma_m, x, y\sigma_m) \\ &= \frac{1}{m!} \sum_{\sigma_m \in S_m} \left[\sum_{i=1}^n g_{i\sigma_m^{-1}(j)}(v\sigma_m, x, y\sigma_m) \right] \leq \frac{1}{m!} \sum_{\sigma_m \in S_m} 1 = 1, \end{aligned}$$

and

$$\begin{aligned} \sum_{i=1}^n (\mathcal{Q}_3 g)_{ij}(v, x, y) &= \frac{1}{n!m!} \sum_{i=1}^n \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} g_{\sigma_n^{-1}(i)\sigma_m^{-1}(j)}(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m) \\ &= \frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} \left[\sum_{i=1}^n g_{\sigma_n^{-1}(i)\sigma_m^{-1}(j)}(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m) \right] \leq \frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} 1 = 1, \end{aligned}$$

thus we know the projected allocation rule is also feasible, which will never allocate one item more than once.

In addition, for the payment rule, we have

$$\begin{aligned} (\mathcal{Q}_1 p)_i(v, x, y) &= \frac{1}{n!} \sum_{\sigma_n \in S_n} p_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y) \\ &\leq \frac{1}{n!} \sum_{\sigma_n \in S_n} \left[\sum_{j=1}^m g_{\sigma_n^{-1}(i)j}(\sigma_n v, \sigma_n x, y)(\sigma_n v)_{\sigma_n^{-1}(i)j} \right] \\ &= \sum_{j=1}^m \left[\frac{1}{n!} \sum_{\sigma_n \in S_n} g_{\sigma_n^{-1}(i)j}(\sigma_n v, \sigma_n x, y)v_{ij} \right] \leq \sum_{j=1}^m (\mathcal{Q}_1 g)_{ij}(v, x, y)v_{ij}, \end{aligned}$$

$$\begin{aligned} (\mathcal{Q}_2 p)_i(v, x, y) &= \frac{1}{m!} \sum_{\sigma_m \in S_m} p_i(v\sigma_m, x, y\sigma_m) \\ &\leq \frac{1}{m!} \sum_{\sigma_m \in S_m} \left[\sum_{j=1}^m g_{ij}(v\sigma_m, x, y\sigma_m)(v\sigma_m)_{ij} \right] \\ &= \sum_{j=1}^m \left[\frac{1}{m!} \sum_{\sigma_m \in S_m} g_{i\sigma_m^{-1}(j)}(v\sigma_m, x, y\sigma_m)v_{ij} \right] \leq \sum_{j=1}^m (\mathcal{Q}_2 g)_{ij}(v, x, y)v_{ij}, \end{aligned}$$

and

$$\begin{aligned} (\mathcal{Q}_3 p)_i(v, x, y) &= \frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} p_{\sigma_n^{-1}(i)}(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m) \\ &\leq \frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} \left[\sum_{j=1}^m g_{\sigma_n^{-1}(i)j}(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m)(\sigma_n v\sigma_m)_{\sigma_n^{-1}(i)j} \right] \\ &= \sum_{j=1}^m \left[\frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} g_{\sigma_n^{-1}(i)\sigma_m^{-1}(j)}(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m)v_{ij} \right] \leq \sum_{j=1}^m (\mathcal{Q}_3 g)_{ij}(v, x, y)v_{ij}, \end{aligned}$$

thus, we have completed this proof.

A.2 Proof of Theorem 3.1

In this section, we prove Theorem 3.1.

Proof of Theorem 3.1. We first study the part in the condition of bidder-symmetry; *i.e.*, the orbit averaging \mathcal{Q} . is the bidder averaging \mathcal{Q}_1 , acting on the allocation rule g and the payment rule p as below,

$$\mathcal{Q}_1 g(v, x, y) = \frac{1}{n!} \sum_{\sigma_n \in S_n} \sigma_n^{-1} g(\sigma_n v, \sigma_n x, y),$$

and

$$\mathcal{Q}_1 p(v, x, y) = \frac{1}{n!} \sum_{\sigma_n \in S_n} \sigma_n^{-1} p(\sigma_n v, \sigma_n x, y).$$

Step 1: We first prove that the auction mechanism has the same expected revenue after projection, *i.e.*,

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_1 p]_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right].$$

Given that for all permutation π ,

$$\sum_{i=1}^n p_i = \sum_{i=1}^n p_{\pi(i)},$$

we have the following equation,

$$\begin{aligned} \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \mathcal{Q}_1 p_i(v, x, y) \right] &= \mathbb{E}_{(v,x,y)} \left[\frac{1}{n!} \sum_{i=1}^n \sum_{\sigma_n \in S_n} p_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\frac{1}{n!} \sum_{i=1}^n \sum_{\sigma_n \in S_n} p_i(\sigma_n v, \sigma_n x, y) \right] = \frac{1}{n!} \sum_{i=1}^n \sum_{\sigma_n \in S_n} \mathbb{E}_{(v,x,y)} \left[p_i(\sigma_n v, \sigma_n x, y) \right] \\ &= \frac{1}{n!} \sum_{i=1}^n \sum_{\sigma_n \in S_n} \mathbb{E}_{(v,x,y)} \left[p_i(v, x, y) \right] = \sum_{i=1}^n \mathbb{E}_{(v,x,y)} \left[p_i(v, x, y) \right]. \end{aligned}$$

Thus, we complete the first step.

Step 2: We then prove that the sum of all bidders' utilities remains the same after projection, *i.e.*,

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_1 u]_i(v_i, v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n u_i(v_i, v, x, y) \right].$$

Given that

$$\sum_{i=1}^n g_{\pi(i)j} v_{\pi(i)j} = \sum_{i=1}^n g_{ij} v_{ij},$$

we have the following equation,

$$\begin{aligned}
& \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_1 u]_i(v_i, v, x, y) \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \left[\sum_{j=1}^m [\mathcal{Q}_1 g]_{ij}(v, x, y) v_{ij} - [\mathcal{Q}_1 p]_i(v, x, y) \right] \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m [\mathcal{Q}_1 g]_{ij}(v, x, y) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\frac{1}{n!} \sum_{i=1}^n \sum_{j=1}^m \sum_{\sigma_n \in S_n} g_{\sigma_n^{-1}(i)j}(\sigma_n v, \sigma_n x, y) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\
&= \frac{1}{n!} \sum_{\sigma_n \in S_n} \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{\sigma_n^{-1}(i)j}(\sigma_n v, \sigma_n x, y) [\sigma_n v]_{\sigma_n^{-1}(i)j} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\
&= \frac{1}{n!} \sum_{\sigma_n \in S_n} \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{ij}(\sigma_n v, \sigma_n x, y) [\sigma_n v]_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\
&= \frac{1}{n!} \sum_{\sigma_n \in S_n} \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{ij}(v, x, y) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{ij}(v, x, y) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n u_i(v_i, v, x, y) \right].
\end{aligned}$$

Thus, we have completed the second step.

Step 3: We lastly prove that the auction mechanism has a smaller ex-post regret after projection, *i.e.*,

$$\mathbb{E} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n [\mathcal{Q}_1 u]_i(v_i, (v'_i, v_{-i}), x, y) \right] \leq \mathbb{E} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n u(v_i, (v'_i, v_{-i}), x, y) \right].$$

For the simplicity, we denote $\sum_{j=1}^m g_{ij} v_{ij}$ by $\langle g_i, v_i \rangle$. Then, we have,

$$\begin{aligned}
& \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n [\mathcal{Q}_1 u]_i(v_i, (v'_i, v_{-i}), x, y) \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \sum_{j=1}^m [\mathcal{Q}_1 g]_{ij}((v'_i, v_{-i}), x, y) v_{ij} - [\mathcal{Q}_1 p]_i((v'_i, v_{-i}), x, y) \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \langle [\mathcal{Q}_1 g]_i((v'_i, v_{-i}), x, y), v_i \rangle - [\mathcal{Q}_1 p]_i((v'_i, v_{-i}), x, y) \right] \\
&= \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \frac{1}{n!} \sum_{\sigma_n \in S_n} \langle g_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y), v_i \rangle - p_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y) \right].
\end{aligned}$$

Combining the following inequality,

$$\max_z \sum_{k=1}^K f_k(z) \leq \sum_{k=1}^K \max_z f_k(z),$$

we have that

$$\begin{aligned}
& \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \frac{1}{n!} \sum_{\sigma_n \in S_n} \langle g_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y), v_i \rangle - p_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y) \right] \\
& \leq \mathbb{E}_{(v,x,y)} \left[\frac{1}{n!} \sum_{\sigma_n \in S_n} \max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \langle g_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y), v_i \rangle - p_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y) \right] \\
& = \frac{1}{n!} \sum_{\sigma_n \in S_n} \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \langle g_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y), v_i \rangle - p_{\sigma_n^{-1}(i)}(\sigma_n(v'_i, v_{-i}), \sigma_n x, y) \right] \\
& = \frac{1}{n!} \sum_{\sigma_n \in S_n} \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n u_i(v_i, (v'_i, v_{-i}), x, y) \right] \\
& = \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n u_i(v_i, (v'_i, v_{-i}), x, y) \right].
\end{aligned}$$

We thus have completed the proof of eqs. (1) and (2) when the orbit averaging \mathcal{Q}_1 is the condition of bidder-symmetry.

Then, we prove this theorem in the condition of item-symmetry; *i.e.*, the orbit averaging \mathcal{Q}_1 is the item averaging \mathcal{Q}_2 , acting on the allocation rule g and the payment rule p as shown below,

$$\mathcal{Q}_2 g(v, x, y) = \frac{1}{m!} \sum_{\sigma_m \in S_m} g(v\sigma_m, x, y\sigma_m)\sigma_m^{-1},$$

and,

$$\mathcal{Q}_2 p(v, x, y) = \frac{1}{m!} \sum_{\sigma_m \in S_m} p(v\sigma_m, x, y\sigma_m).$$

Step 1: We first prove that the auction mechanism has the same expected revenue after projection, *i.e.*,

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_2 p]_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right].$$

Since the valuation joint distribution is invariant under bidder permutation, we have $\mathbb{E}_{(v,x,y)} [f(v\sigma_m, x, y\sigma_m)] = \mathbb{E}_{(v,x,y)} [f(v, x, y)]$. Then, we have

$$\begin{aligned}
& \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_2 p]_i(v, x, y) \right] \\
& = \mathbb{E}_{(v,x,y)} \left[\frac{1}{m!} \sum_{i=1}^n \sum_{\sigma_m \in S_m} p_i(v\sigma_m, x, y\sigma_m) \right] \\
& = \frac{1}{m!} \sum_{i=1}^n \sum_{\sigma_m \in S_m} \mathbb{E}_{(v,x,y)} \left[p_i(v\sigma_m, x, y\sigma_m) \right] \\
& = \frac{1}{m!} \sum_{i=1}^n \sum_{\sigma_m \in S_m} \mathbb{E}_{(v,x,y)} \left[p_i(v, x, y) \right] = \sum_{i=1}^n \mathbb{E}_{(v,x,y)} \left[p_i(v, x, y) \right].
\end{aligned}$$

We have thus completed the first step.

Step 2: We then prove that the sum of all bidders' utilities remains same after projection, *i.e.*,

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_2 u]_i(v_i, v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n u_i(v_i, v, x, y) \right].$$

Given that for all permutation π ,

$$\sum_{j=1}^m g_{i\pi(j)} v_{i\pi(j)} = \sum_{j=1}^m g_{ij} v_{ij},$$

we have the following equation,

$$\begin{aligned} & \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_2 u]_i(v_i, v, x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \left[\sum_{j=1}^m [\mathcal{Q}_2 g]_{ij}(v, x, y) v_{ij} - [\mathcal{Q}_2 p]_i(v, x, y) \right] \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m [\mathcal{Q}_2 g]_{ij}(v, x, y) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\frac{1}{m!} \sum_{i=1}^n \sum_{j=1}^m \sum_{\sigma_m \in S_m} g_{i\sigma_m^{-1}(j)}(v\sigma_m, x, y\sigma_m) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\ &= \frac{1}{m!} \sum_{\sigma_m \in S_m} \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{i\sigma_m^{-1}(j)}(v\sigma_m, x, y\sigma_m) [v\sigma_m]_{i\sigma_m^{-1}(j)} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\ &= \frac{1}{m!} \sum_{\sigma_m \in S_m} \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{ij}(v\sigma_m, x, y\sigma_m) [v\sigma_m]_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\ &= \frac{1}{m!} \sum_{\sigma_m \in S_m} \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{ij}(v, x, y) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \sum_{j=1}^m g_{ij}(v, x, y) v_{ij} \right] - \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n u_i(v_i, v, x, y) \right]. \end{aligned}$$

Thus, we have completed the second step.

Step 3: We lastly prove that the auction mechanism have a smaller ex-post regret after projection, *i.e.*,

$$\mathbb{E} \left[\max_{v' \in \mathcal{V}^n \times m} \sum_{i=1}^n [\mathcal{Q}_2 u]_i(v_i, (v'_i, v_{-i}), x, y) \right] \leq \mathbb{E} \left[\max_{v' \in \mathcal{V}^n \times m} \sum_{i=1}^n u(v_i, (v'_i, v_{-i}), x, y) \right].$$

By the definition of the utility u and the item averaging \mathcal{Q}_2 , we have

$$\begin{aligned} & \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^n \times m} \sum_{i=1}^n [\mathcal{Q}_2 u]_i(v_i, (v'_i, v_{-i}), x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^n \times m} \sum_{i=1}^n \sum_{j=1}^m [\mathcal{Q}_2 g]_{ij}((v'_i, v_{-i}), x, y) v_{ij} - [\mathcal{Q}_2 p]_i((v'_i, v_{-i}), x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^n \times m} \sum_{i=1}^n \langle [\mathcal{Q}_2 g]_i((v'_i, v_{-i}), x, y), v_i \rangle - [\mathcal{Q}_2 p]_i((v'_i, v_{-i}), x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^n \times m} \sum_{i=1}^n \frac{1}{m!} \sum_{\sigma_m \in S_m} \langle g_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m)\sigma_m^{-1}, v_i \rangle - p_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m) \right]. \end{aligned}$$

Combining the following inequality

$$\max_z \sum_{k=1}^K f_k(z) \leq \sum_{k=1}^K \max_z f_k(z),$$

we have

$$\begin{aligned} & \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \frac{1}{m!} \sum_{\sigma_m \in S_m} \langle g_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m)\sigma_m^{-1}, v_i \rangle - p_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m) \right] \\ & \leq \mathbb{E}_{(v,x,y)} \left[\frac{1}{m!} \sum_{\sigma_m \in S_m} \sum_{i=1}^n \max_{v' \in \mathcal{V}^{n \times m}} \langle g_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m)\sigma_m^{-1}, v_i \rangle - p_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m) \right] \\ & = \frac{1}{m!} \sum_{\sigma_m \in S_m} \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n \langle g_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m), v_i\sigma_m \rangle - p_i((v'_i, v_{-i})\sigma_m, x, y\sigma_m) \right] \\ & = \frac{1}{m!} \sum_{\sigma_m \in S_m} \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n u_i(v_i, (v'_i, v_{-i}), x, y) \right] \\ & = \mathbb{E}_{(v,x,y)} \left[\max_{v' \in \mathcal{V}^{n \times m}} \sum_{i=1}^n u_i(v_i, (v'_i, v_{-i}), x, y) \right]. \end{aligned}$$

We thus have proved this theorem in the condition of item-symmetry.

The proofs are completed. \square

A.3 Proof of Lemma 3.3

In this section, we present the proof of Lemma 3.3.

Proof of Lemma 3.3. Both the bidder averaging and the item averaging are linear. Thus, we have the following results,

$$\begin{aligned} & \mathcal{Q}_1 \circ \mathcal{Q}_2 f(v, x, y) \\ & = \mathcal{Q}_1 \left[\frac{1}{m!} \sum_{\sigma_m \in S_m} f(v\sigma_m, x, y\sigma_m)\sigma_m^{-1} \right] \\ & = \frac{1}{n!} \sum_{\sigma_n \in S_n} \sigma_n^{-1} \left[\frac{1}{m!} \sum_{\sigma_m \in S_m} f(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m)\sigma_m^{-1} \right] \\ & = \frac{1}{n!m!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} \sigma_n^{-1} f(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m)\sigma_m^{-1} = \mathcal{Q}_3 f(v, x, y), \end{aligned}$$

and

$$\begin{aligned} & \mathcal{Q}_2 \circ \mathcal{Q}_1 f(v, x, y) \\ & = \mathcal{Q}_2 \left[\frac{1}{n!} \sum_{\sigma_n \in S_n} \sigma_n^{-1} f(\sigma_n v, \sigma_n x, y) \right] \\ & = \frac{1}{m!} \sum_{\sigma_m \in S_m} \left[\frac{1}{n!} \sum_{\sigma_n \in S_n} \sigma_n^{-1} f(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m) \right] \sigma_m^{-1} \\ & = \frac{1}{m!n!} \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} \sigma_n^{-1} f(\sigma_n v\sigma_m, \sigma_n x, y\sigma_m)\sigma_m^{-1} = \mathcal{Q}_3 f(v, x, y). \end{aligned}$$

The above two equations hold for any f . Then, we may prove that

$$\mathcal{Q}_3 = \mathcal{Q}_1 \circ \mathcal{Q}_2 = \mathcal{Q}_2 \circ \mathcal{Q}_1,$$

which is exactly the claim of this theorem.

The proof is completed. \square

A.4 Proof of Theorem 3.4

In this section, we apply our Lemma 3.3 and Theorem 3.1 to prove Theorem 3.4.

Proof of Theorem 3.4. For the simplicity, we rewrite \mathcal{Q}_3p and \mathcal{Q}_3reg as $\mathcal{Q}_2(\mathcal{Q}_1p)$ and $\mathcal{Q}_2(\mathcal{Q}_1reg)$, respectively. Then, for a payment rule p , we have that,

$$\begin{aligned} & \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_3p]_i(v, x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_2(\mathcal{Q}_1p)]_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_1p]_i(v, x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right]. \end{aligned}$$

Also, we have the following result,

$$\begin{aligned} & \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_3reg]_i(v, x, y) \right] \\ &= \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_2(\mathcal{Q}_1reg)]_i(v, x, y) \right] \leq \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}_1reg]_i(v, x, y) \right] \\ &\leq \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n reg_i(v, x, y) \right]. \end{aligned}$$

Moreover, we have the following result on the regret gap Δ_3 ,

$$\mathbb{E}_{(v,x,y)} [\Delta_3(g, p; v, x, y)] = \mathbb{E}_{(v,x,y)} [\Delta_1(g, p; v, x, y)] + \mathbb{E}_{(v,x,y)} [\Delta_2(\mathcal{Q}_1g, \mathcal{Q}_1p; v, x, y)] \geq 0.$$

This proof is completed. \square

A.5 Proof of Theorem 3.5

In this section, we present the proof of Theorem 3.5.

We start with the definitions or notations necessary for our proof. We define the allocation rule space and the payment rule space as follows,

$$\mathcal{G} = \{g^\omega : \omega \in \Omega\} \text{ and } \mathcal{P} = \{p^\omega : \omega \in \Omega\},$$

where ω is the auction mechanism parameter and Ω is the set of all feasible parameters. We then define the induced utility and ex-post regret spaces as follows,

$$\mathcal{U} = \left\{ u^\omega : u_i^\omega(v'_i, v, x, y) = \sum_{j=1}^m g_{ij}^\omega(v, x, y) v'_{ij} - p_i^\omega(v, x, y) \right\},$$

and

$$\mathcal{R} = \left\{ reg^\omega : reg_i^\omega(v, x, y) = \max_{v'_i} u^\omega(v_i, (v'_i, v_{-i}), x, y) - u^\omega(v_i, v, x, y) \right\}.$$

Then, the $l_{\infty,1}$ -distance on \mathcal{U} and \mathcal{P} is defined as below,

$$l_{\infty,1}(u, u') = \max_{(v, v'_i, x, y)} \left(\sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, (v'_i, v_{-i}), x, y)| \right),$$

and

$$l_{\infty,1}(p, p') = \max_{(v, x, y)} \left(\sum_{i=1}^n |p_i(v, x, y) - p'_i(v, x, y)| \right).$$

We now present the proof of Theorem 3.5.

Proof of Theorem 3.5. We prove Theorem 3.5 in two steps: (1) we first prove that the distance between any two mechanisms is smaller when we project them to be permutation-equivariant; (2) then, we prove that the smaller distance implies a smaller covering number.

Step 1: We prove that the distance between two mechanisms becomes smaller after projection, *i.e.*,

$$l_{\infty,1}(\mathcal{Q}.p, \mathcal{Q}.p') \leq l_{\infty,1}(p, p'),$$

and

$$l_{\infty,1}(\mathcal{Q}.u, \mathcal{Q}.u') \leq l_{\infty,1}(u, u'),$$

where $u, u' \in \mathcal{U}$, $p, p' \in \mathcal{P}$, and $\mathcal{Q}. = \mathcal{Q}_1$ or \mathcal{Q}_2 .

When \mathcal{Q} . is \mathcal{Q}_1 , we have that

$$\begin{aligned} & l_{\infty,1}(\mathcal{Q}_1 p, \mathcal{Q}_1 p') \\ &= \max_{(v,x,y)} \sum_{i=1}^n |\mathcal{Q}_1 p_i(v, x, y) - \mathcal{Q}_1 p'_i(v, x, y)| \\ &= \max_{(v,x,y)} \sum_{i=1}^n \left| \frac{1}{n!} \sum_{\sigma_n \in S_n} [p_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y) - p'_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y)] \right| \\ &\leq \max_{(v,x,y)} \sum_{i=1}^n \frac{1}{n!} \sum_{\sigma_n \in S_n} |p_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y) - p'_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y)| \\ &\leq \sum_{\sigma_n \in S_n} \frac{1}{n!} \max_{(v,x,y)} \sum_{i=1}^n |p_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y) - p'_{\sigma_n^{-1}(i)}(\sigma_n v, \sigma_n x, y)| \\ &= \sum_{\sigma_n \in S_n} \frac{1}{n!} \max_{(v,x,y)} \sum_{i=1}^n |p_i(\sigma_n v, \sigma_n x, y) - p'_i(\sigma_n v, \sigma_n x, y)| \\ &= \sum_{\sigma_n \in S_n} \frac{1}{n!} \max_{(v,x,y)} \sum_{i=1}^n |p_i(v, x, y) - p'_i(v, x, y)| \\ &= \max_{(v,x,y)} \sum_{i=1}^n |p_i(v, x, y) - p'_i(v, x, y)| \\ &= l_{\infty,1}(p, p'), \end{aligned}$$

and

$$\begin{aligned} & l_{\infty,1}(\mathcal{Q}_1 u, \mathcal{Q}_1 u') \\ &= \max_{v,v',x,y} \sum_{i=1}^n |[\mathcal{Q}_1 u]_i(v_i, (v'_i, v_{-i}), x, y) - [\mathcal{Q}_1 u']_i(v_i, (v'_i, v_{-i}), x, y)| \\ &= \max_{v,v',x,y} \sum_{i=1}^n \left| \frac{1}{n!} \sum_{\sigma_n \in S_n} u_{\sigma_n^{-1}(i)}(v_i, \sigma_n(v'_i, v_{-i}), \sigma_n x, y) - u'_{\sigma_n^{-1}(i)}(v_i, \sigma_n(v'_i, v_{-i}), \sigma_n x, y) \right| \\ &\leq \max_{v,v',x,y} \sum_{i=1}^n \frac{1}{n!} \sum_{\sigma_n \in S_n} |u_{\sigma_n^{-1}(i)}(v_i, \sigma_n(v'_i, v_{-i}), \sigma_n x, y) - u'_{\sigma_n^{-1}(i)}(v_i, \sigma_n(v'_i, v_{-i}), \sigma_n x, y)| \\ &\leq \frac{1}{n!} \sum_{\sigma_n \in S_n} \max_{v,v',x,y} \sum_{i=1}^n |u_{\sigma_n^{-1}(i)}(v_i, \sigma_n(v'_i, v_{-i}), \sigma_n x, y) - u'_{\sigma_n^{-1}(i)}(v_i, \sigma_n(v'_i, v_{-i}), \sigma_n x, y)| \\ &= \frac{1}{n!} \sum_{\sigma_n \in S_n} \max_{v,v',x,y} \sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, (v'_i, v_{-i}), x, y)| \\ &= \max_{v,v',x,y} \sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, (v'_i, v_{-i}), x, y)| \\ &= l_{\infty,1}(u, u'). \end{aligned}$$

Then, when \mathcal{Q} is \mathcal{Q}_2 , we prove the result as below,

$$\begin{aligned}
& l_{\infty,1}(\mathcal{Q}_2 p, \mathcal{Q}_2 p') \\
&= \max_{(v,x,y)} \sum_{i=1}^n |\mathcal{Q}_2 p_i(v, x, y) - \mathcal{Q}_2 p'_i(v, x, y)| \\
&= \max_{(v,x,y)} \sum_{i=1}^n \left| \frac{1}{m!} \sum_{\sigma_m \in S_m} [p_i(v\sigma_m, x, y\sigma_m) - p'_i(v\sigma_m, x, y\sigma_m)] \right| \\
&\leq \max_{(v,x,y)} \sum_{i=1}^n \frac{1}{m!} \sum_{\sigma_m \in S_m} |p_i(v\sigma_m, x, y\sigma_m) - p'_i(v\sigma_m, x, y\sigma_m)| \\
&\leq \sum_{\sigma_m \in S_m} \frac{1}{m!} \max_{(v,x,y)} \sum_{i=1}^n |p_i(v\sigma_m, x, y\sigma_m) - p'_i(v\sigma_m, x, y\sigma_m)| \\
&= \sum_{\sigma_m \in S_m} \frac{1}{m!} \max_{(v,x,y)} \sum_{i=1}^n |p_i(v, x, y) - p'_i(v, x, y)| \\
&= \max_{(v,x,y)} \sum_{i=1}^n |p_i(v, x, y) - p'_i(v, x, y)| \\
&= l_{\infty,1}(p, p'),
\end{aligned}$$

and

$$\begin{aligned}
& l_{\infty,1}(\mathcal{Q}_2 u, \mathcal{Q}_2 u') \\
&= \max_{v,v',x,y} \sum_{i=1}^n |[\mathcal{Q}_2 u]_i(v_i, (v'_i, v_{-i}), x, y) - [\mathcal{Q}_2 u']_i(v_i, (v'_i, v_{-i}), x, y)| \\
&= \max_{v,v',x,y} \sum_{i=1}^n \left| \frac{1}{m!} \sum_{\sigma_m \in S_m} |u_i(v_i\sigma_m, (v'_i, v_{-i})\sigma_m, x, y\sigma_m) - u'_i(v_i\sigma_m, (v'_i, v_{-i})\sigma_m, x, y\sigma_m)| \right| \\
&\leq \max_{v,v',x,y} \sum_{i=1}^n \frac{1}{m!} \sum_{\sigma_m \in S_m} |u_i(v_i\sigma_m, (v'_i, v_{-i})\sigma_m, x, y\sigma_m) - u'_i(v_i\sigma_m, (v'_i, v_{-i})\sigma_m, x, y\sigma_m)| \\
&\leq \frac{1}{m!} \sum_{\sigma_m \in S_m} \max_{v,v',x,y} \sum_{i=1}^n |u_i(v_i\sigma_m, (v'_i, v_{-i})\sigma_m, x, y\sigma_m) - u'_i(v_i\sigma_m, (v'_i, v_{-i})\sigma_m, x, y\sigma_m)| \\
&= \frac{1}{m!} \sum_{\sigma_m \in S_m} \max_{v,v',x,y} \sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, (v'_i, v_{-i}), x, y)| \\
&= \max_{v,v',x,y} \sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, (v'_i, v_{-i}), x, y)| \\
&= l_{\infty,1}(u, u').
\end{aligned}$$

Thus, we have completed Step 1.

Step 2: We prove that a smaller distance implies a smaller covering number.

Let \mathcal{X} and \mathcal{Y} be two metric spaces with two different distances l_1 and l_2 , respectively. There exists a surjective mapping f from \mathcal{Y} to \mathcal{X} , such that $l_1(f(x), f(y)) \leq l_2(x, y)$ for all $x, y \in \mathcal{Y}$. The covering numbers $\mathcal{N}_1(\mathcal{X}, r)$ and $\mathcal{N}_2(\mathcal{Y}, r)$ are defined as the minimum numbers of balls with radius r that can cover \mathcal{X} and \mathcal{Y} under l_1 and l_2 , respectively.

By the definition of the covering number $\mathcal{N}_2(\mathcal{Y}, r)$, there exists a set \mathcal{A} of scale $\mathcal{N}_2(\mathcal{Y}, r)$, such that

$$l_2(x, \mathcal{A}) = \inf_{y \in \mathcal{A}} l_2(x, y) < r, \forall x \in \mathcal{Y}.$$

Then, $f(\mathcal{A})$ is also a r -cover for \mathcal{X} under distance l_1 , i.e., for any $x \in \mathcal{X}$, we have

$$l_1(f(x), f(\mathcal{A})) = \inf_{y \in \mathcal{A}} l_1(f(x), f(y)) \leq \inf_{y \in \mathcal{A}} l_2(x, y) = l_2(x, \mathcal{A}) < r.$$

Because f is surjective, for any $x' \in \mathcal{X}$, there exists an $x \in \mathcal{Y}$, such that $x' = f(x)$. Then, for any $x' \in \mathcal{X}$, we have that

$$l_1(x', f(\mathcal{A})) = l_1(f(x), f(\mathcal{A})) < r.$$

By the definition of $\mathcal{N}_1(\mathcal{X}, r)$, we have

$$\mathcal{N}_1(\mathcal{X}, r) \leq |f(\mathcal{A})| \leq |\mathcal{A}| = \mathcal{N}_2(\mathcal{Y}, r).$$

Eventually, combining the results in Step 1 and in Step 2, we have that

$$\mathcal{N}_{\infty,1}(\mathcal{Q}\mathcal{U}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, r),$$

and

$$\mathcal{N}_{\infty,1}(\mathcal{Q}\mathcal{P}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{P}, r),$$

for both the bidder averaging \mathcal{Q}_1 and the item averaging \mathcal{Q}_2 . \square

A.6 Proof of Theorem 3.6 and Corollary 3.7

We first introduce two lemmas. The first lemma gives a concentration inequality via the covering number. This result can be used to bound the gap between expected revenue/ex-post regret and empirical revenue/ex-post regret. The second lemma bounds the covering number $\mathcal{N}_{\infty,1}(\mathcal{R}, 2r)$ by the covering number $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$. Both lemmas has been proved by [10]. We recall them here to make our paper completed.

Lemma A.1 (cf. Lemma E.1, [10]). *Let $\mathcal{S} = \{z_1, \dots, z_L\}$ be a set of i.i.d. sample points drawn from a distribution \mathcal{D} over \mathcal{Z} . Suppose \mathcal{F} is a set of functions from \mathcal{Z} to \mathbb{R} such that $f(z) \in [a, b]$ for all $f \in \mathcal{F}$ and $z \in \mathcal{Z}$. We define l_∞ on \mathcal{F} as*

$$l_\infty(f, f') = \max_{z \in \mathcal{Z}} |f(z) - f'(z)|,$$

and $\mathcal{N}_\infty(\mathcal{F}, r)$ as the minimum number of balls with radius r that can cover \mathcal{F} under l_∞ -distance. Then, we have the following concentration inequality,

$$\mathbb{P}\left[\exists f \in \mathcal{F} : \left| \frac{1}{L} \sum_{i=1}^L f(z_i) - \mathbb{E}[f(z)] \right| > \epsilon\right] \leq 2\mathcal{N}_\infty\left(\mathcal{F}, \frac{\epsilon}{3}\right) \exp\left(-\frac{2L\epsilon^2}{9(b-a)^2}\right).$$

Proof. By the definition of $\mathcal{N}_\infty(\mathcal{F}, r)$, for any $f \in \mathcal{F}$, there exists an $f_r \in \mathcal{F}_r$ such that \mathcal{F}_r is an r -cover for \mathcal{F} and $l_\infty(f, f_r) < r$. Denote $\frac{1}{L} \sum_{i=1}^L f(z_i)$ by $\mathbb{E}_\mathcal{S}[f(z)]$. Then, we have

$$\begin{aligned} & \mathbb{P}\left[\exists f \in \mathcal{F} : \left| \mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}[f(z)] \right| > \epsilon\right] \\ &= \mathbb{P}\left[\exists f \in \mathcal{F} : \left| \mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}_\mathcal{S}[f_r(z)] + \mathbb{E}_\mathcal{S}[f_r(z)] - \mathbb{E}[f_r(z)] + \mathbb{E}[f_r(z)] - \mathbb{E}[f(z)] \right| > \epsilon\right] \\ &\leq \mathbb{P}\left[\exists f \in \mathcal{F} : \left| \mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}_\mathcal{S}[f_r(z)] \right| + \left| \mathbb{E}_\mathcal{S}[f_r(z)] - \mathbb{E}[f_r(z)] \right| + \left| \mathbb{E}[f_r(z)] - \mathbb{E}[f(z)] \right| > \epsilon\right] \\ &\leq \mathbb{P}\left[\exists f_r \in \mathcal{F}_r : \left| \mathbb{E}_\mathcal{S}[f_r(z)] - \mathbb{E}[f_r(z)] \right| > \frac{\epsilon}{3}\right] \\ &\leq \mathcal{N}_\infty\left(\mathcal{F}, \frac{\epsilon}{3}\right) \mathbb{P}\left[\left| \mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}[f(z)] \right| > \frac{\epsilon}{3}\right] \\ &\leq 2\mathcal{N}_\infty\left(\mathcal{F}, \frac{\epsilon}{3}\right) \exp\left(-\frac{2L\epsilon^2}{9(b-a)^2}\right). \end{aligned}$$

The third inequality follows from the fact that when $r = \frac{\epsilon}{3}$, we have

$$|f(z) - f_r(z)| < \frac{\epsilon}{3},$$

for all $z \in \mathcal{Z}$ and $f \in \mathcal{F}$. Then, from the Hoeffding's inequality, we have

$$\left| \frac{1}{L} \sum_{i=1}^L f(z_i) - \frac{1}{L} \sum_{i=1}^L f_r(z_i) \right| < \frac{\epsilon}{3} \quad \text{and} \quad \left| \mathbb{E}[f(z)] - \mathbb{E}[f_r(z)] \right| < \frac{\epsilon}{3}.$$

The proof is completed. \square

The following lemma bounds the covering number $\mathcal{N}_{\infty,1}(\mathcal{R}, 2r)$ by the covering number $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$. Then the gap between expected ex-post regret and empirical ex-post regret can be bounded by the covering number $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$.

Lemma A.2 (cf. Lemma E.3, [10]). *We define $\mathcal{N}_{\infty,1}(\mathcal{R}, r)$ and $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$ as the minimum numbers of balls with radius r that can cover spaces \mathcal{R} and \mathcal{U} under distance $l_{\infty,1}$, respectively. Then, we have that*

$$\mathcal{N}_{\infty,1}(\mathcal{R}, 2r) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, r)$$

Proof. By the definition of $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$, there exists an r -cover \mathcal{U}_r for \mathcal{U} , such that $|\mathcal{U}_r| = \mathcal{N}_{\infty,1}(\mathcal{U}, r)$ and for any $u \in \mathcal{U}$,

$$l_{\infty,1}(u, \mathcal{U}_r) = \inf_{u' \in \mathcal{U}_r} l_{\infty,1}(u, u') < r.$$

We define \mathcal{R}_r as

$$\{reg \in \mathcal{R} : reg_i(v, x, y) = \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - u_i(v_i, v, x, y) \text{ for some } u_i \in \mathcal{U}_r\}.$$

Then, we can prove that \mathcal{R}_r is a $2r$ -cover for the space \mathcal{R} , i.e.,

$$\begin{aligned} & l_{\infty,1}(reg, \mathcal{R}_r) \\ &= \inf_{reg' \in \mathcal{R}_r} l_{\infty,1}(reg, reg') \\ &= \inf_{reg' \in \mathcal{R}_r} \max_{v, x, y} \sum_{i=1}^n |reg_i(v, x, y) - reg'_i(v, x, y)| \\ &= \inf_{u' \in \mathcal{U}_r} \max_{v, x, y} \sum_{i=1}^n \left| \left[\max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - u_i(v_i, v, x, y) \right] \right. \\ &\quad \left. - \left[\max_{v'_i} u'_i(v_i, (v'_i, v_{-i}), x, y) - u'_i(v_i, v, x, y) \right] \right| \\ &\leq \inf_{u' \in \mathcal{U}_r} \max_{v, x, y} \left[\sum_{i=1}^n \left| \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{v'_i} u'_i(v_i, (v'_i, v_{-i}), x, y) \right| \right. \\ &\quad \left. + \left| u_i(v_i, v, x, y) - u'_i(v_i, v, x, y) \right| \right] \\ &\leq \max_{v, x, y} \left[\sum_{i=1}^n \left| \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{v'_i} u_i^*(v_i, (v'_i, v_{-i}), x, y) \right| \right. \\ &\quad \left. + \left| u_i(v_i, v, x, y) - u_i^*(v_i, v, x, y) \right| \right] \quad (\text{where } l_{\infty,1}(u, u^*) < r) \\ &\leq \max_{v, x, y} \sum_{i=1}^n \left| \max_{v'_i} u_i(v_i, (v'_i, v_{-i}), x, y) - \max_{v'_i} u_i^*(v_i, (v'_i, v_{-i}), x, y) \right| + r \\ &\leq \max_{v, x, y} \sum_{i=1}^n \max_{v'_i} |u_i(v_i, (v'_i, v_{-i}), x, y) - u_i^*(v_i, (v'_i, v_{-i}), x, y)| + r \\ &= \max_{v, v'_i, x, y} \sum_{i=1}^n |u_i(v_i, (v'_i, v_{-i}), x, y) - u_i^*(v_i, (v'_i, v_{-i}), x, y)| + r < 2r. \end{aligned}$$

Eventually, we have

$$\mathcal{N}_{\infty,1}(\mathcal{R}, 2r) \leq |\mathcal{R}_r| \leq |\mathcal{U}_r| = \mathcal{N}_{\infty,1}(\mathcal{U}, r).$$

The proof is completed. \square

We now prove Theorem 3.6 and Corollary 3.7.

Proof of Theorem 3.6 and Corollary 3.7. Applying Lemma A.1 to the spaces \mathcal{P} and \mathcal{U} , we have that

$$\begin{aligned} & \mathbb{P} \left[\exists \omega \in \Omega : \left| \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i^\omega(v, x, y) \right] - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) \right| > \epsilon \right] \\ & \leq 2\mathcal{N}_{\infty,1} \left(\mathcal{P}, \frac{\epsilon}{3} \right) \exp \left(-\frac{2L\epsilon^2}{9n^2} \right), \end{aligned}$$

and

$$\begin{aligned} & \mathbb{P} \left[\exists \omega \in \Omega : \left| \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \text{reg}_i^\omega(v, x, y) \right] - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n \text{reg}_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) \right| > \epsilon \right] \\ & \leq 2\mathcal{N}_{\infty,1} \left(\mathcal{R}, \frac{\epsilon}{3} \right) \exp \left(-\frac{2L\epsilon^2}{9n^2} \right) \\ & \leq 2\mathcal{N}_{\infty,1} \left(\mathcal{U}, \frac{\epsilon}{6} \right) \exp \left(-\frac{2L\epsilon^2}{9n^2} \right), \end{aligned}$$

where the last inequality follows from Lemma A.2.

Further, we assume that

$$\epsilon \geq \sqrt{\frac{9n^2}{2L} \left(\log \frac{4}{\delta} + \max \left\{ \log \mathcal{N}_{\infty,1} \left(\mathcal{P}, \frac{\epsilon}{3} \right), \log \mathcal{N}_{\infty,1} \left(\mathcal{U}, \frac{\epsilon}{6} \right) \right\} \right)}.$$

Then, we have the following inequalities,

$$\mathbb{P} \left[\exists \omega \in \Omega : \left| \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i^\omega(v, x, y) \right] - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) \right| > \epsilon \right] \leq \frac{\delta}{2},$$

and

$$\mathbb{P} \left[\exists \omega \in \Omega : \left| \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \text{reg}_i^\omega(v, x, y) \right] - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n \text{reg}_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) \right| > \epsilon \right] \leq \frac{\delta}{2}.$$

Thus, with probability at least $1 - \delta$, for any $\omega \in \Omega$, we have that

$$\left| \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i^\omega(v, x, y) \right] - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) \right| < \epsilon, \quad (5)$$

and

$$\left| \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n \text{reg}_i^\omega(v, x, y) \right] - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n \text{reg}_i^\omega(v^{(l)}, x^{(l)}, y^{(l)}) \right| < \epsilon. \quad (6)$$

Equivalently, when the number of samples L is large enough, *i.e.*,

$$L \geq \frac{9n^2}{2\epsilon^2} \left(\log \frac{4}{\delta} + \max \left\{ \log \mathcal{N}_{\infty,1} \left(\mathcal{P}, \frac{\epsilon}{3} \right), \log \mathcal{N}_{\infty,1} \left(\mathcal{U}, \frac{\epsilon}{6} \right) \right\} \right),$$

then, the eqs. (5) and (6) both hold with probability at least $1 - \delta$.

The proof is completed. \square

A.7 Proof of the Generalization Bound for Myerson Auctions

Denote $\text{rev}(v, x, y)$ as $\sum_{i=1}^n p_i(v, x, y)$, then we have the following theorem,

Theorem A.3. *Assume the item valuation for each bidder is not larger than 1. When the sample complexity satisfies $L \geq \frac{1}{2\epsilon^2} \log \frac{2}{\delta}$, with probability at least $1 - \delta$, we have*

$$\left| \frac{1}{L} \sum_{\ell=1}^L \text{rev}(v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) - \mathbb{E}_{(v,x,y)} [\text{rev}(v, x, y)] \right| \leq \epsilon.$$

Proof. Since $v_i \leq 1$, we have

$$rev(v, x, y) = \sum_{i=1}^n p_i(v, x, y) \leq \sum_{i=1}^n b_i(v, x, y)v_i \leq \sum_{i=1}^n b_i(v, x, y) \leq 1.$$

According to the Hoeffding's inequality, we have

$$\mathbb{P} \left[\left| \frac{1}{L} \sum_{\ell=1}^L rev(v^{(\ell)}, x^{(\ell)}, y^{(\ell)}) - \mathbb{E}_{(v,x,y)} [rev(v, x, y)] \right| \geq \epsilon \right] \leq 2 \exp(-2L\epsilon^2).$$

Let $2 \exp(-2L\epsilon^2) \leq \delta$, then we obtain what we need. The proof is completed. \square

A.8 Orbit Averaging over Subsets of Bidders/Items

In addition, we can extended our theory to orbit averaging over the subset of the bidders/items.

Theorem A.4. *Let \mathcal{Q} be the orbit averaging over any subset of bidders and items, and (g, p) be any mechanism. Then we have*

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}p]_i(v, x, y) \right] = \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n p_i(v, x, y) \right],$$

and

$$\mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n reg_i(v, x, y) \right] \geq \mathbb{E}_{(v,x,y)} \left[\sum_{i=1}^n [\mathcal{Q}reg]_i(v, x, y) \right],$$

where reg_i is the ex-post regret of bidder i .

Proof. Without loss of generality, we assume that \mathcal{Q}_1 takes average over the first \tilde{n} bidders, \mathcal{Q}_2 takes average over the first \tilde{m} items and \mathcal{Q}_3 takes average over the first \tilde{n} bidders and \tilde{m} items. Denote Z_1 as $(v_{i_j}, x_i : i > \tilde{n}, j \in [m])$ and Z_2 as $(v_{i_j}, y_j : i \in [n], j > \tilde{m})$. Following Theorem 3.1, we have

$$\mathbb{E} \left[\sum_{i=1}^n [\mathcal{Q}_1 p]_i(v, x, y) \middle| Z_1 \right] = \mathbb{E} \left[\sum_{i=1}^n p_i(v, x, y) \middle| Z_1 \right],$$

and

$$\mathbb{E} \left[\sum_{i=1}^n reg_i(v, x, y) \middle| Z_1 \right] \geq \mathbb{E} \left[\sum_{i=1}^n [\mathcal{Q}_1 reg]_i(v, x, y) \middle| Z_1 \right].$$

Then, combining the fact that $\mathbb{E}[\mathbb{E}[X|Y]] = \mathbb{E}[X]$, we have

$$\mathbb{E} \left[\sum_{i=1}^n [\mathcal{Q}_1 p]_i(v, x, y) \right] = \mathbb{E} \left[\sum_{i=1}^n p_i(v, x, y) \right],$$

and

$$\mathbb{E} \left[\sum_{i=1}^n reg_i(v, x, y) \right] \geq \mathbb{E} \left[\sum_{i=1}^n [\mathcal{Q}_1 reg]_i(v, x, y) \right].$$

Similarly, replace Z_1 by Z_2 , we can obtain the equations all hold for \mathcal{Q}_2 .

Finally, we prove that $\mathcal{Q}_3 = \mathcal{Q}_1 \mathcal{Q}_2 = \mathcal{Q}_2 \mathcal{Q}_1$. The proof is same with the proof of Lemma 3.3. Only replace n and m by \tilde{n} and \tilde{m} respectively, and we obtain the result.

The proof is completed. \square

Theorem A.5. *Let \mathcal{Q} be the orbit averaging over any subset of bidders and items, and $\mathcal{U} = \{u^\omega : \omega \in \Omega\}$ and $\mathcal{P} = \{p^\omega : \omega \in \Omega\}$ the sets of all possible utilities and payment rules. Then we have*

$$\mathcal{N}_{\infty,1}(\mathcal{Q}\mathcal{U}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{U}, r) \text{ and } \mathcal{N}_{\infty,1}(\mathcal{Q}\mathcal{P}, r) \leq \mathcal{N}_{\infty,1}(\mathcal{P}, r),$$

where $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$ and $\mathcal{N}_{\infty,1}(\mathcal{P}, r)$ are the minimum numbers of balls with radius r that can cover \mathcal{U} and \mathcal{P} under $l_{\infty,1}$ -distance, respectively.

Proof. Without loss of generality, we assume that \mathcal{Q}_1 takes average over the first \tilde{n} bidders, \mathcal{Q}_2 takes average over the first \tilde{m} items and \mathcal{Q}_3 takes average over the first \tilde{n} bidders and \tilde{m} items.

We can prove the distance between two mechanisms becomes smaller after orbit averaging, *i.e.*,

$$l_{\infty,1}(\mathcal{Q}p, \mathcal{Q}p') \leq l_{\infty,1}(p, p') \quad \text{and} \quad l_{\infty,1}(\mathcal{Q}u, \mathcal{Q}u') \leq l_{\infty,1}(u, u').$$

Only replace n and m by \tilde{n} and \tilde{m} in the proof of Theorem 3.5, and we obtain the results.

The proof is completed. \square

A.9 Average over Subgroups

It is worth noting that our proof only relies on one assumption that the valuation joint distribution is invariant under the bidder/item permutation. Consequently, we may adopt orbit averaging over any subgroup of $S_n \times S_m$, while the benefits on revenue and ex-post regret still hold. Hence, there is a trade-off between the auction mechanism performance (revenue and ex-post regret) and the computational complexity: better performance requires more computation. In addition, the choice of the subgroup can also depend on the input feature x [28], which could be more flexible.

B Additional Experimental Details

This section presents additional experimental details and results omitted from the main text due to space limitation.

B.1 Additional Experimental Settings

In this section, we present detailed experimental settings.

B.1.1 Network Architectures

We first describe the RegretNet’s architecture [12]. A RegretNet consists of two parts: the allocation network $g^\omega : \mathbb{R}^{nm} \rightarrow [0, 1]^{nm}$ and the payment network $p^\omega : \mathbb{R}^{nm} \rightarrow \mathbb{R}_{\geq 0}^n$, both of which are modeled as three-layer fully-connected networks with *tanh* activations. Every layer in the two networks includes 100 nodes.

For each item j , the payment network outputs a probability vector $(g_{1j}^\omega(b), \dots, g_{nj}^\omega(b))^T$, where $g_{ij}^\omega(b)$ is the probability of allocating the item j to the bidder i . To avoid allocating one item over once, a feasible allocation network needs to satisfy $\sum_{i=1}^n g_{ij}^\omega(b) \leq 1$ for all $j \in [m]$, $\omega \in \Omega$, and $b \in \mathcal{V}^{nm}$. Therefore, we compute the allocation via a *softmax activation function*. In addition, to present the probability that the item is reserved, an extra dummy node is included in the softmax computation.

To ensure the *individual rational* condition, the payment network p^ω is required to output a payment vector $p^\omega(b)$, such that $p_i^\omega(b) \leq \sum_{j=1}^m g_{ij}^\omega(b)b_{ij}$ for all $i \in [n]$. Therefore, the payment network first computes a fractional payment $\bar{p}_i^\omega(b) \in [0, 1]$ for each bidder i using a sigmoidal unit. Then, the final payment of the bidder i is

$$p_i^\omega(b) = \bar{p}_i^\omega(b) \sum_{j=1}^m g_{ij}^\omega(b)b_{ij} \leq \sum_{j=1}^m g_{ij}^\omega(b)b_{ij}.$$

An overview of the RegretNet’s architecture is illustrated in the following Figure 1.

RegretNet-PE is designed by modifying RegretNet. We adopt the allocation rule as $\tilde{g}^\omega = \mathcal{Q}_3 g^\omega$ and the payment rule as $\tilde{p}^\omega = \mathcal{Q}_3 p^\omega$, respectively. In this way, we may guarantee that in RegretNet-PE, the allocation is feasible and the mechanism is *individual rational*, *i.e.*, $\sum_{i=1}^n \tilde{g}_{ij}^\omega(b) \leq 1$, and $\tilde{p}_i^\omega(b) \leq \sum_{j=1}^m \tilde{g}_{ij}^\omega(b)b_{ij}$. We may also show that the RegretNet-PE is always permutation-equivariant and has the same number of coefficients as the RegretNet. The proof can be found in Appendix A.1.

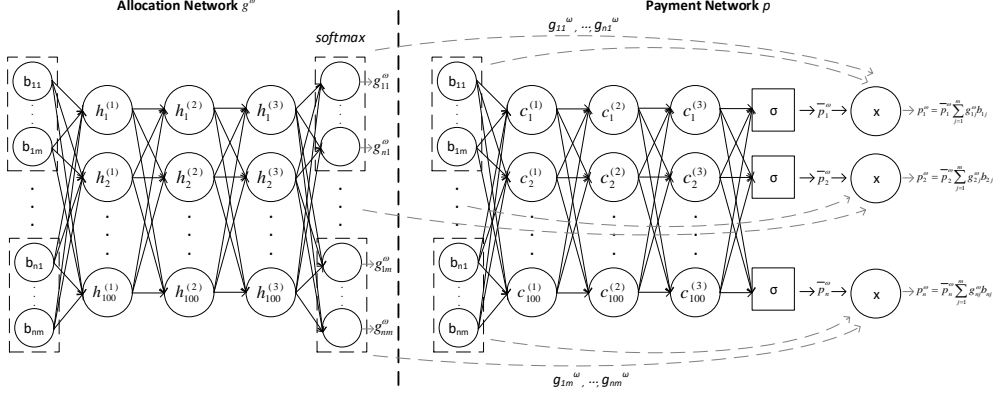


Figure 1: The allocation network g^ω and the payment network p^ω of the RegretNet with the overall mechanism parameter $\omega = (\omega_g, \omega_p)$.

B.1.2 Training Procedures

We adopt the augmented Lagrangian method to minimize the following object function with a quadratic penalty term for violating the constraints,

$$\mathcal{L}_\rho(\omega, \lambda) = -\frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^\omega(v^{(l)}) + \sum_{i=1}^n \lambda_i \widehat{reg}_i(\omega) + \frac{\rho}{2} \left(\sum_{i=1}^n \widehat{reg}_i \right)^2,$$

where L is the number of samples, λ is a vector of Lagrange multipliers, and $\rho > 0$ is a parameter to control the weight of the quadratic penalty. We alternately update the overall mechanism parameter ω and the Lagrange multiplier λ as follows:

- (a) $\omega^{new} \in \arg \min_{\omega} \mathcal{L}_\rho(\omega^{old}, \lambda^{old})$ (every iteration);
- (b) $\lambda_i^{new} = \lambda_i^{old} + \rho \cdot \widehat{reg}_i(\omega^{new})$, $\forall i \in [n]$ (every T_λ iterations).

The training procedure is described in the following Algorithm 1.

We divide all training samples \mathcal{S} into T batches $\mathcal{S}_1, \dots, \mathcal{S}_T$ of size B . At iteration t , we use the batch $\mathcal{S}_t = \{v^{(1)}, \dots, v^{(B)}\}$.

The update (a) is computed via Adam. The gradient of \mathcal{L}_ρ w.r.t. ω for a fixed λ^t is as below,

$$\nabla_{\omega} \mathcal{L}_\rho(\omega, \lambda^t) = -\frac{1}{B} \sum_{l=1}^B \sum_{i=1}^n \nabla_{\omega} p_i^\omega(v^{(l)}) + \sum_{l=1}^B \sum_{i=1}^n \lambda_i^t g_{l,i}^t + \rho \sum_{l=1}^B \sum_{i=1}^n \widehat{reg}_i(\omega) g_{l,i}^t,$$

where

$$\widehat{reg}_i(\omega) = \frac{1}{B} \sum_{l=1}^B \max_{v'_i \in \mathcal{V}^m} u_i^\omega(v_i^{(l)}, (v'_i, v_{-i}^{(l)})) - u_i^\omega(v_i^{(l)}, v^{(l)}),$$

and

$$g_{l,i}^t = \nabla_{\omega} \left[\max_{v'_i \in \mathcal{V}^m} u_i^\omega(v_i^{(l)}, (v'_i, v_{-i}^{(l)})) - u_i^\omega(v_i^{(l)}, v^{(l)}) \right] \Big|_{\omega=\omega^t}.$$

Because $\widehat{reg}_i(\omega)$ and $g_{l,i}^t$ both contain a ‘‘max’’ over misreports², we use another Adam to compute the approximated best biddings $v'^{(l)}$. In each update on ω^t , we perform R updates to compute a best bidding $v'^{(l)}$ for each $i \in [n]$. In particular, we maintain the misreports $v'^{(l)}$ for each sample l as the initial value in the next iteration. Then, we use these biddings $v'^{(l)}$ to compute the gradient $\nabla_{\omega} \mathcal{L}_\rho(\omega, \lambda^t)$ and then, update ω^t as $\omega^{t+1} = \omega^t - \eta \nabla_{\omega} \mathcal{L}_\rho(\omega^t, \lambda^t)$. After every T_λ iterations, we update λ^t as $\lambda_i^{t+1} = \lambda_i^t + \rho \widehat{reg}_i(\omega^{t+1})$. In addition, we increase the value of ρ every a certain number of iterations, where we set the value of ρ_t in each iteration t prior to training.

²The misreport refers to an arbitrary bid, rather than restricted to be a truthful bid [12].

Algorithm 1: RegretNet and RegretNet-PE Training

Input: Batches $\mathcal{S}_1, \dots, \mathcal{S}_T$ of size B **Parameters:** $\forall t \in [T], \rho_t > 0, \gamma > 0, \eta > 0, T \in \mathbb{N}, R \in \mathbb{N}, T_\lambda \in \mathbb{N}$ **Initialize:** $\omega^0 \in \mathbb{R}^d, \lambda^0 \in \mathbb{R}^n$ **for** $t = 0$ **to** T **do** Receive batch $S_t = \{v^{(1)}, \dots, v^{(B)}\}$ Initialize misreports $v_i^{(l)} \in \mathcal{V}^m, \forall l \in [B], i \in [n]$ **for** $r = 0$ **to** R **do** $\forall l \in [B], i \in [n]:$

$v_i^{(l)} \leftarrow v_i^{(l)} + \gamma \nabla_{v_i'} u_i^\omega(v_i^{(l)}, (v_i^{(l)}, v_{-i}^{(l)}))$

end Compute ex-post regret gradient : $\forall l \in [B], i \in [n]:$

$g_{i,l}^t \leftarrow \nabla_\omega \left[u_i^\omega(v_i^{(l)}, (v_i^{(l)}, v_{-i}^{(l)})) - u_i^\omega(v_i^{(l)}, v^{(l)}) \right] \Big|_{\omega=\omega^t}$

 Compute Lagrangian gradient using Equation 4 and update ω^t :

$\omega^{t+1} \leftarrow \omega^t - \eta \nabla_\omega \mathcal{L}_{\rho_t}(\omega^t, \lambda^t)$

 Update Lagrange multipliers λ once in T_λ iterations: **if** t is a multiple of T_λ **then**

$\lambda_i^{t+1} = \lambda_i^t + \rho_t \widehat{reg}_i(\omega^{t+1}), \forall i \in [n]$

else

$\lambda^{t+1} = \lambda^t$

end**end**

Table 3: Additional experimental results. " $n \times m$ Normal" refers that there are n bidders and m items, and the valuation is drawn from the truncated normal distribution $\mathcal{N}(0.3, 0.1)$ in $[0,1]$. The true values of the ex-post regret and the generalization error (GE) are the products of the values in the table and a factor of 10^{-5} .

Method	2 × 1 Normal			3 × 1 Normal		
	Revenue	Regret	GE	Revenue	Regret	GE
Optimal	0.304	0	-	0.391	0	-
RegretNet	0.275	97.0	8.50	0.321	84.0	45.5
RegretNet-Test	0.275	95.2	-	0.321	75.0	-
RegretNet-PE	0.276	85.4	8.40	0.382	69.7	27.6
Method	2 × 2 Normal			5 × 3 Normal		
	Revenue	Regret	GE	Revenue	Regret	GE
RegretNet	0.577	343	246	1.05	114	77.0
RegretNet-Test	0.577	327	-	1.05	32.0	-
RegretNet-PE	0.577	318	77.0	1.09	75.0	70.0

B.1.3 Test Settings

To verify our Theorem 3.1 and Theorem 3.4, we first train a RegretNet and then project the will-trained RegretNet to be permutation-equivariant through bidder-item aggregated averaging \mathcal{Q}_3 , denoted as "RegretNet-Test". To meet the *symmetric valuation* condition in Theorem 3.1 and Theorem 3.4, we sample a set of valuations from the distribution, which is denoted by \mathcal{S} , and then, induce a set of symmetric samples $\tilde{\mathcal{S}} = \{\sigma_n v \sigma_m : \sigma_n \in S_n, v \in \mathcal{S}, \sigma_m \in S_m\}$ for test.

Table 4: Additional experimental result, where " $n \times m$ Compound" refers that there are n bidders and m items, and the valuations are i.i.d. sampled from the compound distributions.

Method	3×1 Compound		5×1 Compound	
	Revenue	Regret	Revenue	Regret
RegretNet	0.516	< 0.001	0.329	< 0.001
EquivariantNet	0.498	< 0.001	0.311	< 0.001
RegretNet-PE	0.539	< 0.001	0.356	< 0.001

Table 5: Additional experimental result, where " $n \times m$ Uniform" refers that there are n bidders and m items, and the valuations are i.i.d. sampled from the uniform distribution $U[0, 1]$.

Method	2×5 Uniform			5×3 Uniform		
	Revenue	Regret	GE	Revenue	Regret	GE
RegretNet	2.24	104	86.4	1.56	28.4	19.4
RegretNet-Test	2.24	74.0	-	1.56	8.60	-
RegretNet-PE	2.38	89.9	24.9	1.85	20.1	11.8

For a RegretNet-PE, there is no difference between test on $\tilde{\mathcal{S}}$ and on \mathcal{S} , because

$$\frac{1}{n!m!L} \sum_{l=1}^L \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} f(\sigma_n v_i \sigma_m) = \frac{1}{n!m!L} \sum_{l=1}^L \sum_{\sigma_n \in S_n} \sum_{\sigma_m \in S_m} f(v_i) = \frac{1}{L} \sum_{l=1}^L f(v_i),$$

for any permutation-equivariant function f and a RegretNet-PE is always permutation-equivariant.

To compute the best bidding v'_i for each bidder i , we first randomly initialize 1,000 misreports in all settings, and then, perform 2,000 updates on each misreport via Adam with the same settings. Finally, we choose the best one (which induces a maximal utility of the bidder i) as the approximated best bidding v'_i .

B.1.4 Implementation Details

We train the models (RegretNet and RegretNet-PE) for up to 150 epochs with a batch size of 128 ($B = 128$) and report the early-stop results for RegretNet-PE to obtain a comparable ex-post regret. The terminal iteration numbers for RegretNet-PE are 10,000 in the 2×1 setting, 17,000 in the 3×1 setting, 18,000 in the 5×1 setting, 300,000 in the 1×2 setting, and 600,000 in the 2×2 setting. Our insight is that the larger terminal iteration number required in the 2×2 setting is because of the small model size, *i.e.*, where the networks have three layers (each of 100 nodes). The value of ρ is initialized as 1.0 and increased by 5 every 200 batches. For each update on ω^t , we initialize one misreport and update the misreport by Adam for each bidder with 25 steps ($R = 25$) and learning rate 0.1 ($\gamma = 0.1$). The final optimal misreports will be used to initialize the misreports for the same batch in the next epoch. We update ω^t via Adam for every batch with a learning rate of 0.001. Besides, we update λ^t every 200 batches.

B.2 Additional Experiment Results

We present additional experimental results. Each valuation v_{ij} is sampled independently from (1) a truncated normal distribution $\mathcal{N}(0.3, 0.1)$ in $[0, 1]$; (2) a compound distribution $\mathcal{N}(\frac{x_i}{6}, 0.1)$ truncated in $[0, 1]$, where x_i is sampled independently and uniformly from $\{1, 2, 3, 4, 5\}$ (cf. Setting A, [10]); and (3) a compound distribution $U[0, \text{Sigmoid}(x_i^T y_j)]$, where x_i and y_j are sampled independently and uniformly from $[-1, 1]$ (Setting C, [10]). All results are shown in Table 3 and Table 4. The revenue and ex-post regret of RegretNet and EquivariantNet in Table 4 come from the previous work [10]. In Table 4, we report the ex-post regret as " < 0.001 " following the previous works.

Moreover, we extend our experiments to more complex settings, including two-bidder five-item and five-bidder three-item settings. Due to the computation limitations, we sample $\{3840, 1280\}$ data




Figure 2: Allocation rule learned by RegretNet (up) and RegretNet-PE (down) for two-bidder and one-item setting. The solid regions describe the probability to allocating the item to bidder 1 (left) and bidder 2 (right). The optimal auction mechanism is described by the regions separated by the dashed black lines, where the number 0 or 1 is the probability of optimal allocation rule in the region.

points and initialize $\{150, 120\}$ misreports for test. Each valuation is sampled from the uniform distribution $U[0, 1]$ and the truncated normal distribution $\mathcal{N}(0.3, 0.1)$ in $[0, 1]$. The results are shown in Tables 3 and 5.

B.3 Allocation Rules Learned by RegretNet and RegretNet-PE

In this section, we show the allocation rules learned by RegretNet and RegretNet-PE in two-bidder, one-item setting and one-bidder, two-item setting, where the valuation is drawn from the uniform distribution $U[0, 1]$. The optimal auction mechanisms are both known.

B.3.1 Two-bidder and One-item Setting

For the two-bidder, one-item setting, the optimal mechanism is well-known as Myerson auction [23], which allocates the item to the highest bidder with receiving a payment of the maximum of the second price and the reserve price, if the highest bid is higher than the reserve price. The allocation rules learned by RegretNet and RegretNet-PE are shown in Figure 2. From Figure 2. We can find that the two learned allocation rules are both almost the same as the optimal mechanism.




Figure 3: Allocation rule learned by RegretNet (up) and RegretNet-PE (down) for one-bidder and two-item setting. The solid regions describe the probability of allocating the first item (left) and the second item (right). The optimal auction mechanism is described by the regions separated by the dashed black lines, where the number 0 or 1 is the probability of optimal allocation rule in the region.

B.3.2 One-bidder and Two-item Setting

The optimal mechanism is given by [20]. Same with the above, we show the allocation rules learned by RegretNet and that learned by RegretNet-PE in Figure 3. The improvement is significant. when one item's valuation is close to 0 and another item's is close to 1, the mechanism learned by RegretNet has a positive probability to allocate the item with the lower valuation to the bidder, while RegretNet-PE and the optimal mechanisms would not.