A Context-Aware Framework for Integrating Ad Auctions and Recommendations

Anonymous Author(s)

ABSTRACT

Recently, many e-commerce platforms have favored presenting a mixed list of ads and organic content to users. The widely-used approach separately ranks ads and organic items, then sequentially inserts ads into the list of organic items. However, this method yields sub-optimal results. Firstly, it only ensures that each generated ad and organic item list achieves local optimality, while the predetermined insertion order fails to guarantee global optimality. Secondly, this approach overlooks the mutual effect between organic items and ads, resulting in an incomplete utilization of contextual information. Besides, it cannot prevent strategic behavior by advertisers. Therefore, we propose a context-aware integrated framework to address these issues. This framework applies automated mechanism design to integrated ad auctions for the first time. Specifically, it models ads and organic items simultaneously along with their contextual information and employs a learning-based approach to prevent advertisers from engaging in strategic behavior. Afterward, the framework directly generates a mixed list, enhancing the overall performance. We also propose **T**ransformer encoder-based **I**ntegrated **C**ontextual **Net**work (TICNet) to generate the optimal integrated contextual ad auction. Finally, we validate the effectiveness of TICNet on synthetic and real-world datasets. Our experimental results demonstrate that TICNet significantly outperforms baseline models across multiple metrics.

CCS CONCEPTS

• **Information systems → Computational advertising**; *Online advertising*; *Electronic commerce*.

KEYWORDS

Contextual Information, Integrated Ad Auction, Transformer, Automated Mechanism Design

ACM Reference Format:

Anonymous Author(s). 2025. A Context-Aware Framework for Integrating Ad Auctions and Recommendations. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 14 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Selling advertisements is one of the core revenue sources for e-commerce platforms. Platforms place ads in prominent positions to attract users' attention and increase their willingness to make purchases. As display space on the platform is limited, it is usually auctioned off to advertisers with marketing intentions to generate profits. Advertisers benefit from getting display resources to attract potential users and increase their revenue.

However, the content on e-commerce platforms not only includes advertisements but also organic items generated by recommendation systems. Although organic items do not directly bring profits to the platform, they often provide users with a better browsing experience due to the recommendation system's strong ability to understand user behavior. This, in turn, enhances the platform's user engagement and retention rates, helping the platform build a good reputation and ultimately have greater long-term value. Therefore, current e-commerce platforms generate a mixed list that includes both ad content and native content to ensure the platform's revenue while providing users with a high-quality experience.

So, what rules should be followed to mix advertisements and organic content? The mainstream approach is to consider the ad system and recommendation system separately. Each system generates ad lists and organic content lists from their respective candidate pools based on specific rules. The ads from the ad list are then sequentially inserted into the organic items list to create the final mixed list for display to users. However, this paradigm has the following issues:

- It can only achieve the suboptimal solutions for individual lists and cannot achieve the overall optimal solution. Since the order within the ad list and the order within the organic items list are fixed, only the relative sequence of ads and organic items can be changed when merging the lists, limiting the solution space and failing to attain the theoretically best solution.
- This method often does not consider the mutual influence between ads and organic items. The model does not incorporate contextual information from the other side when constructing ad lists or organic item lists, resulting in suboptimal items and sequences in each generated list, leading to suboptimal content presentation in the final combined list.
- It may cause advertisers to engage in strategic behavior. By designing reasonable allocation rule and payment scheme, the ad auction system can incentivize advertisers to bid truthfully, which is guaranteed to be their optimal strategy. However, when the ad list and organic content list are merged, the original ad positions change, causing truthful bidding by advertisers to no longer be their optimal strategy. This can lead to strategic behavior by advertisers, affecting the stability of the system, thereby harming long-term revenue.

1.1 Main Contributions

To address these issues, we propose a context-aware integrated framework. This framework characterizes both advertisements and organic items simultaneously in the context of global information

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

and takes platform revenue and user experience as core metrics of objective, forming a constrained optimization problem with DSIC and IR as constraints. We leverage the automatic mechanism design (AMD) methods to find the optimal desirable mechanisms by proposing the **T**ransformer encoder-based **I**ntegrated **C**ontextual **Net**work (TICNet). Additionally, we conduct a series of experiments on both synthetic and real-world datasets to showcase the efficiency of TICNet, compared to some baseline models. The main contributions of this paper are summarized as follows:

(1) **Unified context-aware mechanism to realize integration of ads and organic items.** Different from current methods of inserting ads into list of organic items, we model this process as a mechanism design problem, instead. We execute a mechanism directly on the candidate set of both advertisements and organic items, and decide the allocation of ads and organic items simultaneously. Additionally, necessary contextual information is taken into account to reflect the mutual influence between ads and organic items, while designing mechanisms.
(2) **Automatic mechanism design and sample complexity.** Due to the vast input space and high complexity of the solution space, it is much difficult to directly find the optimal mechanism. Therefore, we first exploit the concept of regret to characterize DSIC condition. Then we transform this mechanism design problem into a learning problem and adopt the techniques of AMD to seek the optimal mechanism. To demonstrate the efficiency of this transformation, we theoretically provide the sample complexity to bound the gap between the empirical and expected values on revenue, user experience and regret.
(3) **Transformer encoder-based Integrated Contextual Network.** We construct TICNet, a neural network framework based on transformer encoder structure, to generate the optimal mechanism. This framework can effectively capture implicit correlations among contextual information while also possessing properties such as permutation equivariance, allowing us to ignore input order.
(4) **Numerical experiments on synthetic and real data.** We conducted multiple experiments on synthetic and real datasets. In comparison with common-used and theoretical baseline mechanisms, we find that the mechanism generated by TICNet outperforms across multiple metrics.

1.2 Related Work

1.2.1 Integrating Recommendation System with Ad Auctions. For a long time, ad auctions have been the primary source of revenue for e-commerce platforms. Pioneered by the seminal work [28], a series of theoretical research focus on enhancing the revenue through designing auction mechanism. For instance, generalized second price auction and its variants [13, 18, 26, 33, 34] have been employed as the core mechanisms of multiple online platforms, due to their simplicity and scalability. On the other hand, some works [21, 25, 40] opt for leveraging advanced machine learning technique, like deep learning, to surpass revenue limitations.

Recently, platforms are aware of the mutual effect between ads and organic items and try to integrate ads into the list of organic items. A main focus is modeling this problem as a linear programming [4, 5, 38]. Furthermore, a line of works model the allocation of slots as a Markov decision process, employing reinforcement learning for resolution [22, 36, 37, 41, 42]. Nevertheless, these methods all view ads and organic items separately, not resulting in a global optimal mixed list. From theoretical perspective, there are a few of works taking ads and organic items into consideration simultaneously. Li et al. [20] introduce an optimally truthful mechanism aiming at balancing revenue and user experience. Li et al. [19] focus on the scenario where products can be selected as either ad or organic item, and propose two simple mechanisms within this setting. However, these methods do not consider the contextual information, ignoring the influence between ads and organic items.

1.2.2 Automated Mechanism Design. In our work, different from traditional methods, we adopt automated mechanism design to generate optimal integrated contextual ad auction mechanism. AMD can swiftly generate optimal auctions for different settings. Dütting et al. [12] propose the first neural network based structure, RegretNet, to handle multi-item auction, which can achieve high revenue and guarantee approximately DSIC. A series of works stemming from RegretNet has extended to various scenarios [8, 14, 15, 29–31, 39]. Additionally, based on the structure of Vickery-Clark-Groves (VCG) mechanism [6, 16, 35], another kind of approaches aim to design the optimal auction mechanisms that automatically adapt DSIC and IR in different settings [9, 10, 23, 24, 27, 32]. Notably, there are lots of works devoting to analyzing the sample complexity [1–3, 7], guaranteeing the performance of generated mechanisms. Since we mainly focus on the design of optimal integrated ad auction with contextual information in this paper, the relevant works are CITransNet [11], utilizing transformer to design optimal contextual auction and RegretFormer [17], taking advantage of attention mechanism to achieve the optimal revenue performance with restricted IC violation budget. However, these methods do not consider the externality from recommendation system, and cannot apply to the integrated scenarios. In contrast, our work proposes a brand new architecture to capture contextual information, suitable for organic items, and finally generate the optimal integrated context-aware mechanisms.

2 MODEL AND PRELIMINARIES

In this section, we mainly introduce the basic setup of integrated ad auctions with contextual information. Later on, we formalize the optimal mechanism design as a constrained optimization problem and transform it into a learning problem.

2.1 Contextual Integrated Ad Auction

On the internet platform, after a user initiating a search request, the platform will return a mixed list containing organic content and advertising content to the user. W.l.o.g, we consider a webpage containing K advertisement slots. For each slot $k \in [K] = \{1, \cdots, K\}$, we denote α_k by the click-through rate (CTR) of the slot k. For simplicity, we assume that the CTRs of the ad slots are in descending order, that is, $1 > \alpha_1 \geq \cdots \geq \alpha_K \geq 0$. Consider that there are m advertisers and n organic items competing for these K ad slots.

Every advertiser $i \in [m]$ possesses an ad-context denoted by $x_i \in \mathcal{X} \subset \mathbb{R}^{d_c}$. Similarly, each organic item $j \in [n]$ is associated with an organic-context represented as $y_j \in \mathcal{Y} \subset \mathbb{R}^{d_c}$, and every slot $k \in [K]$ comes with a slot-context denoted by $z_k \in \mathcal{Z} \subset \mathbb{R}^{d_z}$. Specifically, d_c represents the dimension of ad-context and organic-context information, while d_z represents that of slot-context information. Then we denote $\mathbf{x} = (x_1, \cdots, x_m)$ as the ad contexts, $\mathbf{y} = (y_1, \cdots, y_n)$ as the organic contexts and $\mathbf{z} = (z_1, \cdots, z_K)$ as the slot contexts. x, y and z are sampled from publicly known distributions $\mathcal{D}_x$, $\mathcal{D}_y$ and $\mathcal{D}_z$.

For each advertiser i, her private value for each click by user is denoted by v_i. Assume that the value v_i is sampled from a distribution $\mathcal{F}_{v_i|x_i,z}$, which depends on her ad-context and other slot-contexts information, over the domain set $\mathcal{V}_i$. Let $\mathbf{v} = \{v_1, \cdots, v_m\}$ represent the value profile and $\mathcal{V} = \mathcal{V}_1 \times \mathcal{V}_2 \times \cdots \mathcal{V}_m$ stand for the joint value domain set. Besides, we use $\mathbf{v}_{-i}$ to express the value profile of all advertisers except for i and $\mathcal{V}_{-i}$ to represent the joint value profile domain set except for i. Due to the reason that advertisers may submit their bids strategically, we utilize b_i as the bid of advertiser i, depending on her private value v_i. Similarly, we define $\mathbf{b}, \mathbf{b}_{-i}$ as related bid profiles accordingly.

Besides, we assume that showing ads or organic items to the user has a known effect on the user. Specifically, after the user receives a mixed list including ads or organic items from the platform, she will form distinct impressions about each item. These impressions are influenced by various factors such as the perceived value, creativity, and other attributes of the items, as well as the user's preferences for different items. We model this impression as user experience, which can reflect the user's click intentions and purchasing tendencies. For each advertiser $i \in [m]$ and organic item $j \in [n]$, we use g_i and g_j to denote the expected user experience, respectively. The user experience profile, denoted by $\mathbf{g} = \{g_1, \cdots, g_m, g_{m+1}, \cdots, g_{m+n}\}$, encompasses both advertisers and organic items. To be specific, the user experience profile g_i of advertiser i is sampled from a distribution $\mathcal{F}^q_{g_i|x_i,z}$ and g_j of organic item j is sampled from another distribution $\mathcal{F}^o_{g_j|y_j,z}$. Similar to the definition of joint value distribution, let $\mathcal{G} = \mathcal{G}_1 \times \mathcal{G}_2 \times \cdots \times \mathcal{G}_m \times \mathcal{G}_{m+1} \times \cdots \times \mathcal{G}_{m+n}$ be the joint user experience domain set of ads and organic items.

Now we formally introduce the integrated contextual ad auction, i.e., $\mathcal{M} = (\mathbf{a}, \mathbf{p})$, consisting of the allocation rule $\mathbf{a}(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = \{a_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\}_{i \in M \cup N}$ and the payment scheme $\mathbf{p}(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = \{p_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\}_{i \in M}$. To be specific, $a_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) : \mathcal{V} \times \mathcal{G} \times \mathcal{X}^m \times \mathcal{Y}^n \times \mathcal{Z}^K \rightarrow [0, 1]$ stands for the expected CTR that ad (or organic item) i can derive from showing to user. That is, $a_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = \sum_{k=1}^K a_{ik}(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\alpha_k$, where $a_{ik}(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \in \{0, 1\}$ indicates whether item i is allocated to slot k or not. As for the payment rule, $p_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})$ specifies the payment amount that advertiser i is required to charge when her ad is clicked.

In the setting of integrated contextual ad auction, each item only can obtain at most one slot and each slot is allocated to just one item. We will now formally outline the constraints of the allocation rule, called feasibility,

$$\sum_{k \in [K]} a_{ik}(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \leq 1, \quad \forall i \in [m] \cup [n],$$
$$\sum_{i \in [m] \cup [n]} a_{ik}(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = 1, \quad \forall k \in [K],$$
$$a_{ik}(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \in \{0, 1\}, \quad \forall i \in [m] \cup [n], k \in [K].$$

Under the setting of integrated contextual ad auction, each advertiser i pursues optimizing her quasi-linear utility, defined as

$$u_i(v_i; \mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = v_i \cdot a_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) - p_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})$$

for all $i \in [m], v_i \in \mathcal{V}_i, \mathbf{b} \in \mathcal{V}, \mathbf{g} \in \mathcal{G}, \mathbf{x} \in \mathcal{X}^m, \mathbf{y} \in \mathcal{Y}^n, \mathbf{z} \in \mathcal{Z}^K$.

In this paper, we focus on the mechanisms satisfying dominant strategy incentive compatibility (DSIC) and individual rationality (IR). DSIC guarantees that bidder's optimal strategy is to truthfully reveal her private value, and IR ensures that bidder always can obtain non-negative utility. Now we formally introduce these two economic properties:

Definition 1 (Dominant Strategy Incentive Compatibility). *An integrated contextual ad auction $(\mathbf{a}, \mathbf{p})$ is dominant strategy incentive compatible, if for any advertiser, her utility is maximized by truthfully telling no matter what the others report. Formally, it holds that, for all $i \in [m], v_i \in \mathcal{V}_i, b'_i \in \mathcal{V}_i, \mathbf{b}_{-i} \in \mathcal{V}_{-i}, \mathbf{g} \in \mathcal{G}, \mathbf{x} \in \mathcal{X}^m, \mathbf{y} \in \mathcal{Y}^n, \mathbf{z} \in \mathcal{Z}^K$,*

$$u_i(v_i; (v_i, \mathbf{b}_{-i}), \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \geq u_i(v_i; (b'_i, \mathbf{b}_{-i}), \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}).$$

Definition 2 (Individual Rationality). *An integrated contextual ad auction $(\mathbf{a}, \mathbf{p})$ is individual rational, if for any advertiser, her utility will be non-negative if she truthfully bids. Formally, it holds that, for all $i \in [m], v_i \in \mathcal{V}_i, b'_i \in \mathcal{V}_i, \mathbf{b}_{-i} \in \mathcal{V}_{-i}, \mathbf{g} \in \mathcal{G}, \mathbf{x} \in \mathcal{X}^m, \mathbf{y} \in \mathcal{Y}^n, \mathbf{z} \in \mathcal{Z}^K$,*

$$u_i(v_i; (v_i, \mathbf{b}_{-i}), \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \geq 0.$$

In an integrated contextual ad auction satisfying DSIC and IR, advertisers will truthfully bid. Denote $\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}$ by the joint distribution of value profile $\mathbf{v}$, user experience profile $\mathbf{g}$, ad contexts $\mathbf{x}$, organic contexts $\mathbf{y}$ and slot contexts $\mathbf{z}$. The platform mainly concentrates on two core metrics: revenue and user experience. The expected revenue of the platform is defined as:

$$\text{Rev} := \mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) \sim \mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}} \left[\sum_{i \in [m]} p_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \right],$$

while the expected user experience of the platform is defined as:

$$\text{UE} := \mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) \sim \mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}} \left[\sum_{i \in [m] \cup [n]} g_i \cdot a_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \right].$$

We aim to design an optimal contextual integrated ad auction which optimizes the blend of revenue and user experience while maintaining DSIC and IR. In order to balance the relationship of revenue and user experience, we introduce a hyperparameter $\gamma > 0$. Formally, our problem can be expressed as a constrained optimization problem:

$$\max_{(\mathbf{a},\mathbf{p})} \quad \mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) \sim \mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}} \left[\sum_{i \in [m]} p_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) + \gamma \cdot \sum_{i \in [m] \cup [n]} g_i \cdot a_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \right]$$
$$\text{s.t.} \quad \text{DSIC, IR, Feasibility.}$$

2.2 Integrated Contextual Ad Auction Design as a Learning Problem

Now we formulate the problem of designing the optimal integrated contextual ad auction as a learning problem. For the sake of achieving DSIC, we introduce the metric of *ex-post regret*. For advertiser i, her regret is the maximum utility she can gain through misreporting while keeping the bids of other advertisers fixed, i.e.,

$$rgt_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) = \max_{b_i' \in \mathcal{V}_i} \big[u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) - u_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big].$$

We can draw such conclusion that an integrated contextual ad auction achieves DSIC if and only if its regret equals to 0.

Then, we can reformulate the problem of designing the optimal integrated contextual ad auction as a constrained optimization:

$$\begin{aligned}\min_{(\mathbf{a},\mathbf{p})\in\mathcal{M}} \quad & -\mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\sim\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}}\Bigg[\sum_{i\in[m]} p_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) + \\ & \gamma\cdot\sum_{i\in[m]\cup[n]} g_i\cdot a_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Bigg] \\ \text{s.t.} \quad & \mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\sim\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}}\Bigg[\sum_{i\in[m]} rgt_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Bigg] = 0,\end{aligned}$$

where $\mathcal{M}$ is the set of all contextual integrated ad auctions that satisfy IR and feasibility. However, the problem remains intractable due to the high complexity of the constraints. To solve this problem, we parameterize the auction mechanism as $\mathcal{M}^w(\mathbf{a}^w,\mathbf{p}^w)\subseteq\mathcal{M}(\mathbf{a},\mathbf{p})$, where $w\in\mathbb{R}^{d_w}$ (with dimension d_w) are the parameters. Then we aim to find an contextual integrated ad auction $\mathcal{M}^w$ that minimizes the negated objective: $-\mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\sim\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}}[\sum_{i\in[m]} p_i^w(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) + \gamma\cdot\sum_{i\in[m]\cup[n]} g_i\cdot a_i^w(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})]$, while ensuring DSIC, IR and feasibility constraints through optimizing the parameters w.

Given a sample $\mathcal{L}$, containing L samples of $(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})$ drawn from the distribution $\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}$, we can estimate the ex-post regret under $\mathcal{M}^w(\mathbf{a}^w,\mathbf{p}^w)$ as:

$$\widehat{rgt}_i(w) = \frac{1}{L}\sum_{\ell=1}^{L} rgt_i(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)}).$$

To sum up, we can formulate our constrained optimization problem as a learning problem as follows:

$$\begin{aligned}\min_{w\in\mathbb{R}^{d_w}} \quad & -\frac{1}{L}\sum_{\ell=1}^{L}\Bigg[\sum_{i\in[m]} p_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)}) + \\ & \gamma\cdot\sum_{i\in[m]\cup[n]} g_i\cdot a_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})\Bigg] \\ \text{s.t.} \quad & \widehat{rgt}_i(w) = 0, \quad i = 1,\cdots,m.\end{aligned} \tag{1}$$

Additionally, we guarantee that our auction mechanism satisfies IR through the structure of our architecture. We will offer related details in Section 4.

3 SAMPLE COMPLEXITY

In practice, the observed results (such as revenue, user experience) often deviate from the predicted expectations. By analyzing the sample complexity, we can estimate the required scale of samples to keep these deviations within an acceptable range, thereby guaranteeing the stability and reliability of the mechanism's performance. Therefore, in order to assess to what size of samples used in the transformed learning problem can solve the original problem, in this section, we delve deeper into providing a sample complexity concerning the contextual integrated ad auction class $\mathcal{M}$ to bound three gaps. The first one is the gap between the empirical revenue and the expected revenue; The second one is the gap between the empirical user experience and the expected user experience; The last one is the gap between the empirical ex-post regret and the expected ex-post regret.

Basically, we measure the capacity of $\mathcal{M}$ through covering numbers. We define the $\ell_{\infty,1}$-distance between two auction mechanisms $(\mathbf{a},\mathbf{p}),(\mathbf{a}',\mathbf{p}')\in\mathcal{M}$ as

$$\max_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}} \sum_{i\in[m]\cup[n],k\in[K]} \Big|a_{ik}(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) - a_{ik}'(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big| + \sum_{i\in[m]}\Big|p_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) - p_i'(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|.$$

For any $r>0$, consider $\mathcal{N}_{\infty,1}(\mathcal{M},r)$ as the minimum number of balls with radius r that cover all the mechanisms in $\mathcal{M}$ under $\ell_{\infty,1}$-distance (referred to as the r-covering number of $\mathcal{M}$). We have the following result. The detailed proofs can be found in Appendix A.

THEOREM 1. *For each advertiser i, assume that the valuation function satisfies $v_i(S)\le 1$, and for each advertiser (or organic item) j, assume that the user experience function satisfies $g_j(S)\le 1, \forall S\subseteq M$. Fix $\delta,\epsilon\in(0,1)$, for any $(\mathbf{a}^w,\mathbf{p}^w)$, when*

$$L \ge \frac{9(m+n)^2}{2\epsilon^2}\left(\ln\frac{6}{\delta} + \ln\mathcal{N}_{\infty,1}(\mathcal{M},\frac{\epsilon}{6(m+n)})\right),$$

with probability at least $1-\delta$ over draw of training set S of L samples from $\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}$, we have following conclusions:

$$\left|\sum_{i=1}^{n}\left(\mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\sim\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}} p_i^w(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) - \sum_{\ell=1}^{L}\frac{p_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})}{L}\right)\right| \le \epsilon, \tag{2}$$

$$\left|\sum_{i=1}^{m+n}\left(\mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\sim\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}} g_i\cdot a_i^w(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}) - \sum_{\ell=1}^{L}\frac{g_i\cdot a_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})}{L}\right)\right| \le \epsilon, \tag{3}$$

and

$$\left|\mathbb{E}_{(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\sim\mathcal{F}_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}}\left[\sum_{i=1}^{n} rgt_i^w(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\right] - \sum_{i=1}^{n}\widehat{rgt}_i(w)\right| \le \epsilon. \tag{4}$$

4 TICNET

In this section, we formally introduce the **T**ransformer encoder-based **I**ntegrated **C**ontextual **Net**work (**TICNet**), to design the optimal integrated contextual ad auction.

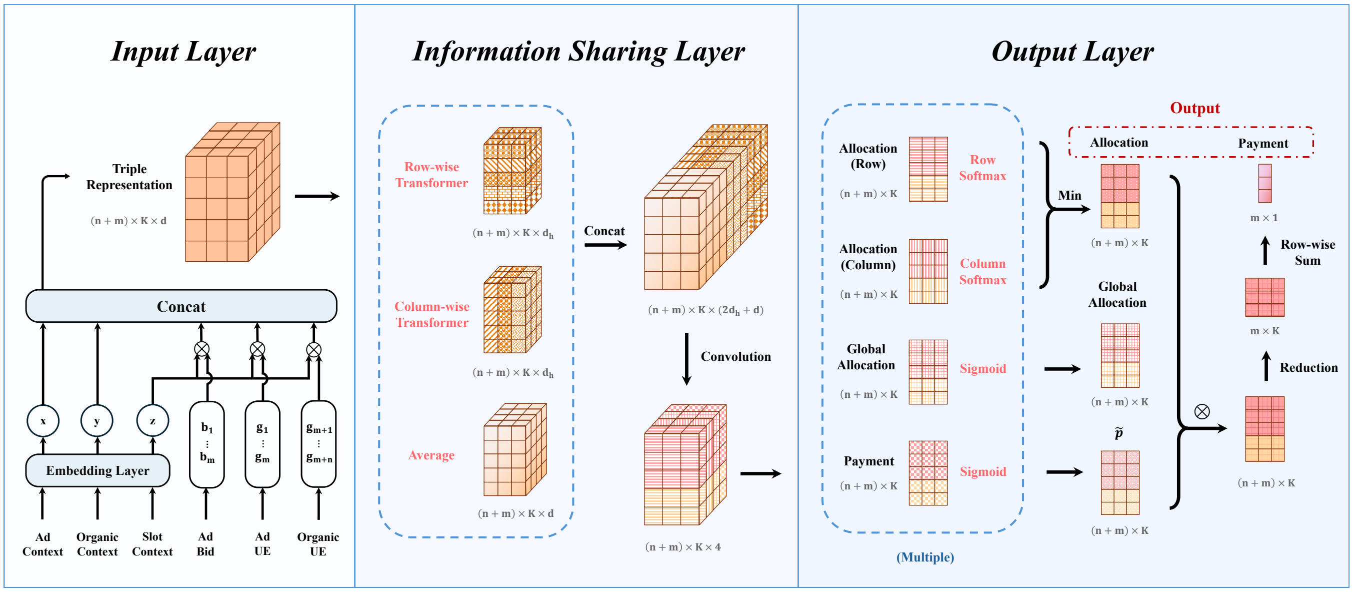

Figure 1: The Architecture of TICNet.

4.1 Architecture of TICNet

As illustrated in Figure 1, TICNet comprises three main components: the input layer, the information sharing layer, and the output layer. The input layer is responsible for receiving representations from ads, organic items, and slots, preprocessing and combining them to generate an aggregated matrix, which is subsequently fed into the information sharing layer. Information sharing layer performs a series of transformations on the aggregated matrix to model the implicit relationships among ads, organic items, and slots with contextual information. The outcome from the information sharing layer is then received by the output layer, which can calculate the allocation of ads and organic items, as well as the payment of each advertiser.

4.2 Input Layer

First of all, we apply embedding layers to obtain a representation $t_i \in \mathbb{R}^{d'_c}$ for each ad-context x_i, $t_j \in \mathbb{R}^{d'_c}$ for each organic-context y_j and $o_k \in \mathbb{R}^{d'_z}$ for every slot-context z_k from the raw state.

Afterward, as mentioned before, each advertiser i has bid profile and user experience profile, while each organic item j is only equipped with user experience profile. We assume that organic item j can also submit a bid to the platform, with its bid set to 0. Different from advertisers, every organic item j cannot misreport its bid intentionally. With the assumption above, we transfer bid profile and user experience profile into $\mathbf{e} = (e_{ik})_{i\in[m]\cup[n],k\in[K]}$ and $\mathbf{q} = (q_{jk})_{j\in[m]\cup[n],k\in[K]}$, where e_{ik} and q_{jk} are defined as:

$$\begin{cases} q_{jk} = g_j\alpha_k, & \forall j \in [m] \cup [n],\ \forall k \in [K]; \\ e_{ik} = b_i\alpha_k, & \forall i \in [m],\ \forall k \in [K]; \\ e_{ik} = 0, & \forall i \in [n],\ \forall k \in [K]. \end{cases}$$

Notably, CTRs of different slots is included in the slot context $\mathbf{z}$.

With the representation above, we then obtain the ad-organic-slot triple representation $R = (H_{i,k})_{i\in[m]\cup[n],k\in[K]}$, in which:

$$H_{i,k} = [e_{ik}||q_{ik}||t_i||o_k] \in \mathbb{R}^{2+d'_c+d'_z},$$

where "||" stands for concatenation.

In order to enhance the efficiency of TICNet, we utilize two linear layers with a ReLU activation to reduce the last dimension of H from $2 + d'_c + d'_z$ to $d - 2$, i.e.,

$$H' = \text{Linear}_2(\text{ReLU}(\text{Linear}_1(H))) \in \mathbb{R}^{(m+n)\times K\times(d-2)}$$

After concatenating H' with transferred bidding profile and user expeirence profile, we can obtain the output of the input layer:

$$O = [H'||\mathbf{e}||\mathbf{q}] \in \mathbb{R}^{(m+n)\times K\times d},$$

in which $O_{ij} \in O$ provide the joint representation of bid, volume and corresponding context information for ad (or organic item) taking place of slot j.

4.3 Information Sharing Layer

With the triple representation $O \in \mathbb{R}^{(m+n)\times K\times d}$ from input layer, we then focus on digging in the inner relationship among ads, organic items and slots. Basically, we adopt the transformer-based encoder to capture complex interactive information among different items and and slots.

To be specific, we employ a transformer-based encoder to model the interactions between ad (or organic item) i and different slots by operating row-wisely on O:

$$O^{\text{row}}_{i\cdot} = \text{Transformer}(O_{i\cdot}) \in \mathbb{R}^{(m+n)\times d_h}, \quad \forall i \in [m] \cup [n],$$

where d_h represents the dimension of the hidden layer of the MLP in the encoder. Similarly, we apply column-wise encoder on the k-th column of O to model the interactions between slot k and all ads and organic items:

$$O^{\text{column}}_{\cdot k} = \text{Transformer}(O_{\cdot k}) \in \mathbb{R}^{K \times d_h}, \quad \forall k \in [K].$$

In addition, we conduct averaging on O to obtain the global representation:

$$O^{\text{global}}_{ik} = \frac{1}{(m+n) \cdot K} \sum_{i \in [m] \cup [n]} \sum_{k \in [K]} O_{ij} \in \mathbb{R}^{d}.$$

Then we concatenate O^{row}, O^{column} and O^{global} together on the last dimension and obtain $O' = (O'_{ik})_{i \in [m] \cup [n], k \in [K]}$:

$$O'_{ik} = [O^{\text{row}}_{ik} || O^{\text{column}}_{ik} || O^{\text{global}}_{ik}] \in \mathbb{R}^{2 \cdot d_h + d}.$$

Afterward, we apply two linear layers with a ReLU activation to reduce the last dimension of O' from $2 \cdot d_h + d$ to d_{out}, i.e.,

$$S' = \text{Linear}_4(\text{ReLU}(\text{Linear}_3(O'))) \in \mathbb{R}^{(m+n) \times K \times d_{out}},$$

where P is the output of the information sharing block. We can stack multiple information sharing blocks together to enhance the representation capability of joint ad-organic-slot information.

4.4 Output Layer

Given the last information sharing block, we set $d_{out} = 4$ and obtain the joint ad-organic-slot representation $S = (S^r, S^c, S^a, S^p) \in \mathbb{R}^{(m+n) \times K \times 4}$, which will be used to calculate the allocation result and the corresponding payment for winning advertisers.

Now we move on to calculate the allocation output a^w_{ik}, representing the probability that ad (or organic item) i being allocated to slot k.

To ensure feasibility, we apply row-wise softmax on S^r to get $\hat{S}^r$, which satisfies each ad (or organic item) i will obtain no more than one slot. Similarly, we apply column-wise softmax on S^c to get $\hat{S}^c$, which guarantees that no slot is over-allocated. Formally, $\hat{S}^r$ and $\hat{S}^c$ are defined as:

$$\begin{cases} \hat{S}^r_{i\cdot} = \text{Softmax}(S^r_{i\cdot}), & \forall i \in [m] \cup [n] \\ \hat{S}^c_{\cdot k} = \text{Softmax}(S^c_{\cdot k}), & \forall k \in [K] \end{cases}$$

Then we obtain the original allocation probability by:

$$\hat{a}^w_{ik} = \min\left\{\hat{S}^r_{ik}, \hat{S}^c_{ik}\right\}, \quad \forall i \in [m] \cup [n], \forall k \in [K].$$

Afterward, due to some ads and organic items may fail to obtain any slot, however, they could still retain a certain portion of the original allocation result. It is important to reduce the impact of failed ads or organic items on the final allocation results. To address this, we employ S^a to adjust the original allocation probability. Essentially, we conduct sigmoid activation on S^a to derive $\hat{S}^a$, which is then multiplied by $\hat{a}^w_{ik}$ to yield the final allocation result a^w_{ik}:

$$\begin{cases} \hat{S}^a_{ik} & = \text{Sigmoid}(S^a_{ik}), \quad \forall i \in [m] \cup [n], \forall k \in [K] \\ a^w_{ik}(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) & = \hat{a}^w_{ik}(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \cdot \hat{S}^a_{ik}(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \end{cases}$$

Then we move on to calculate the payment. Initially, we calculate the auxiliary payment scalar based on S^p through sigmoid activation:

$$\tilde{p}^w_i = \text{Sigmoid}(\frac{1}{K} \sum_{k \in K} S^p), \quad \forall i \in [m],$$

Given the auxiliary payment scalar andfinal allocation result, we then obtain the final payment result as:

$$p^w_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = \tilde{p}^w_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \cdot \left[\sum_{k \in [K]} a^w_{ik}(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) \cdot e_{ik} \right].$$

Since $\tilde{p}^w_i \in (0, 1)$, the quasi-linear utility of advertiser i, denoted as $u^w_i = \sum_{k \in [K]} a^w_{ik} \cdot e_{ik} - p^w_i$ remains non-negative, thereby ensuring IR of our mechanism.

4.5 Permutation Equivariant

Till now, we have introduced the structure of each component of TICNet. Reviewing the overall structure, TICNet meets important property, i.e., *permutation-equivariant*. Thus, TICNet offers several benefits. Firstly, it ensures that the allocation and payment outcomes for the same ad (organic item) remain unaffected regardless of how these candidates are ordered in the input. Secondly, it reduces *ex-post regret* while maintaining the platform's revenue. Moreover, TICNet requires fewer samples to achieve high performance because it exploit the symmetry in the integrated ad auction setting, reducing the hypothesis space and making it easier for the model to learn the underlying structure of the problem, thereby enhancing its generalization capability.

Theorem 2 (Permutation Equivariant). *With bid profile, user experience profile, ad contexts, organic contexts and slot contexts as input, TICNet keeps permutation-equivariant.*

Due to the space limitation, the detailed proof of Theorem 2 is provided in Appendix B.

4.6 Training Procedure

As for the training process of TICNet, we have optimized the Equation (1) using the augmented Lagrangian method, and corresponding formula is:

$$\begin{aligned} C_\rho(w; \boldsymbol{\lambda}) = &- \frac{1}{L} \sum_{\ell=1}^{L} \Bigg[\sum_{i \in [m]} p^w_i(\mathbf{v}^{(\ell)}, \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) + \\ & \gamma \cdot \sum_{j \in [m] \cup [n]} g_j \cdot a^w_j(\mathbf{v}^{(\ell)}, \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) \Bigg] \\ & + \sum_{i=1}^{m} \lambda_i \cdot \widehat{rgt}_i(w) + \frac{\rho}{2} \sum_{i=1}^{m} (\widehat{rgt}_i(w))^2, \end{aligned}$$

During optimization, we first calculate the optimal misreports and then alternately update model parameters and Lagrange multipliers. We have provided the detailed description in Appendix C, containing training settings in Appendix C.1 and corresponding Algorithm 1 in Appendix C.2.

5 EXPERIMENTS

In this section, we conduct a series of experiments on synthetic and real-world datasets to validate the superiority of TICNet.

Method	A: $2\times3\times2$				B: $3\times2\times2$				C: $3\times5\times2$			
	SW	Rev	UE	Rev+γUE	SW	Rev	UE	Rev+γUE	SW	Rev	UE	Rev+γUE
GSP and Fixed Positions	**0.590**	0.233	0.441	0.453	0.595	0.302	0.484	0.545	**0.642**	0.351	0.449	0.576
IAS	0.513	**0.422**	0.362	0.603	**0.601**	**0.517**	0.300	0.667	0.611	**0.505**	0.280	0.645
TICNet	0.467	0.376	**0.517**	**0.634**†	0.516	0.417	**0.520**	**0.677**†	0.596	0.484	**0.466**	**0.717**†

Table 1: The notation $m\times n\times K$ represents a setting where m advertisers and n organic items compete for K slots. In the metric Rev+γUE, the hyperparameter γ is set to 0.5 across different settings. For the GSP and Fixed Positions mechanism, the first slot is reserved for advertisements. The regret of mechanism generated by the TICNet is less than 0.001. The best performance is highlighted in bold. "†" indicates a statistically significant improvement in a paired t-test at $p < 0.05$ level.

5.1 Experimental Settings

5.1.1 Baseline Methods. We compare the TICNet with the following two representative mechanisms that can be applied to integrate advertising auction and recommendation system:

- **GSP** [13] with **Fixed Positions**, a well-known method, which is implemented sequentially. First, we determine the number and specific positions of ad slots. Next, organic items are ranked based on their user experience profiles. The order of ads is then established using the GSP mechanism. Finally, the winning ads are placed in the predetermined slots and charged accordingly.
- **IAS** [20], a Myerson-based mechanism, ranks ads and organic items using a ranking score of $\phi_i(v_i)+\gamma g_i$, where $\phi_i(v_i)=v_i-(1-F_i(v_i))/f_i(v_i)$ is the virtual value. Payments are determined according to the Myerson payment rule [28].

5.1.2 Evaluation Metrics. To assess the performance of TICNet and other baselines, we evaluate the empirical social welfare: SW $=\frac{1}{L}\sum_{\ell=1}^{L}\sum_{i=1}^{m} v_i^{(\ell)}\cdot a_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})$, the empirical revenue: Rev $=\frac{1}{L}\sum_{\ell=1}^{L}\sum_{i=1}^{m} p_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})$, the empirical user experience: UE $=\frac{1}{L}\sum_{\ell=1}^{L}\sum_{i=1}^{m+n} g_i\cdot a_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})$ and the empirical combination of revenue and user experience with coefficient γ: Rev+γUE $=\frac{1}{L}\sum_{\ell=1}^{L}[\sum_{i=1}^{m} p_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})+\gamma\sum_{i=1}^{m+n} g_i\cdot a_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})]$.

5.2 Synthetic Data

For the synthetic data, we generate datasets under various settings, each consisting of a training sample of 100,000 profiles and a testing sample of 5,000 profiles.

Contexts can be categorized into two types based on their features: continuous and discrete. For continuous features, we sample contextual information from a multi-dimensional uniform distribution. For discrete features, we preset a finite number of discrete contextual types. We generate the bid profiles of advertisers and the user experience profiles of advertisers and organic items based on their corresponding contextual information.

We then conduct several experiments comparing TICNet with baseline methods in the following settings:

(A) 2 ads, 3 organic items and 2 slots with CTR $\alpha=(0.7,\,0.3)$, continuous ad-contexts, organic-contexts and slot-contexts, where $\mathcal{X}=[-1,1]^5$, $\mathcal{Y}=[0,2]^5$ and $\mathcal{Z}=[-2,2]^5$. The raw value of each ad is independently drawn from $U[0,1]$. The raw user experience of these 2 ads and 3 organic items is independently drawn from $U[0,0.5]$ and $U[0.5,1]$, respectively. Given contextual information $x_i\in\mathcal{X}$, $y_j\in\mathcal{Y}$ and $z_k\in\mathcal{Z}$, slot-wise value profile v'_{ik} is derived from raw value profile v_i according to $v_i\cdot\text{Sigmoid}(x_i^T z_k)$, slot-wise user experience profile g'_{jk} is obtained from raw user experience profile g_j according to $g_j\cdot\text{Sigmoid}(y_j^T z_k)$.

(B) 3 ads, 2 organic items and 2 slots with CTRs $\alpha=(0.7,0.3)$. The context information, raw profiles and corresponding transferred profiles are drawn similarly to Setting A.

(C) 3 ads, 5 organic items and 2 slots with CTRs $\alpha=(0.7,0.3)$. The context information, raw profiles and corresponding transferred profiles are drawn similarly to Setting A.

The results of Settings A, B and C are outlined in Table 1. Notably, the integrated mechanism generated by TICNet consistently attains a notably higher value for the Rev+γUE metric compared to other mechanisms, while maintaining a regret of less than 0.001. This demonstrates the effectiveness of the TICNet architecture in blending ads and organic items.

To further validate the superiority of TICNet, we carry out experiments across various settings. These include different value distributions, different hyperparameter γ-values, different numbers of slots, and different ratios of candidate ads to organic items. The results for the first two experiments are presented in the main body of the paper, while the results for the latter two are provided in Appendix D. Below are the detailed experimental setups.

5.2.1 Different Value Distributions. To showcase the generalization capability of TICNet across various value distributions, we select Setting B and perform three sets of experiments with different distributions. We sample the raw values from three distinct distributions: a uniform distribution $U[0,1]$; a normal distribution $N(0.5,0.5)$ truncated to the $[0,1]$ interval, and a lognormal distribution $LN(0.2,1.69)$ also truncated to the $[0,1]$ interval. The contextual information and corresponding transferred profiles are obtained in a similar manner as outlined in Setting B.

The experimental results in Table 2 indicate that TICNet consistently achieves the highest Rev+γUE across three different value distributions. This performance significantly exceeds that of the other two mechanisms, highlighting TICNet's stability and robustness in handling various value distributions.

5.2.2 Hyper-parameter Analysis. To assess the impact of the hyperparameter γ on the experimental outcomes, we adjust γ within

Method	Uniform				Normal				Lognormal			
	SW	Rev	UE	Rev+γUE	SW	Rev	UE	Rev+γUE	SW	Rev	UE	Rev+γUE
GSP and Fixed Positions	0.595	0.302	0.485	0.545	**0.596**	0.303	**0.484**	0.546	0.387	0.224	0.488	0.467
IAS	**0.601**	**0.518**	0.299	0.668	0.582	**0.485**	0.246	0.608	0.394	**0.332**	0.284	0.474
TICNet	0.516	0.416	**0.519**	**0.675**†	0.527	0.402	0.439	**0.621**†	**0.433**	0.319	**0.512**	**0.575**†

Table 2: The results of experiments for different value distributions. The setting is 3 ads and 2 organic items with 2 slots. The regret of mechanism generated by the TICNet is less than 0.001. The best performance is highlighted in bold. "†" indicates a statistically significant improvement in a paired t-test at $p < 0.05$ level.

the interval [0.2, 2.0] based on Setting A. We plot the Pareto-curves for various mechanisms, with the results illustrated in Figure 2.

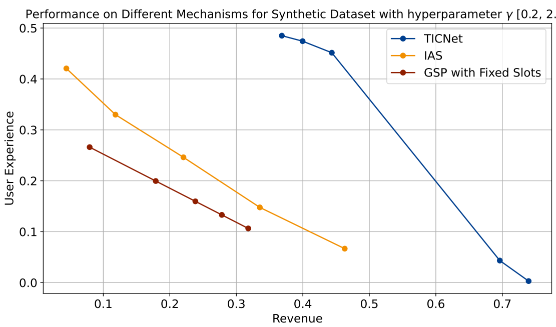

Figure 2: The Pareto curves of TICNet and other baseline mechanisms on synthetic dataset.

Figure 2 shows that the TICNet curve is positioned in the upper right compared to the other curves. This indicates that, across a wide range of γ values, TICNet effectively balances revenue and user experience, resulting in a more desirable blended list than the baseline methods.

5.3 Real-world Dataset

In addition to the synthetic dataset, we also conducted experiments using real log data from an online platform in August 2024. In the environment of feed, the system returns a ranked list containing ads and organic items. We split the dataset into training and test sets with a 9:1 ratio. Additionally, since the number of advertisers, organic items, and slots varies across different auction logs, we trimmed the dataset to ensure appropriate settings.

We present the experimental results on the test set of the real-world dataset in Table 3. Compared to the two baseline methods, the mechanism generated by TICNet achieves a significantly higher Rev+γUE value, with paired t-tests at the $p < 0.05$ level. Additionally, TICNet's mechanism demonstrates approximate DSIC, with a regret of less than 0.001. These results highlight the effectiveness of TICNet in real-world auction scenarios.

We also conduct hyperparameter analysis using a real-world dataset by adjusting γ within the interval [0.2, 2.0] interval. The resulting Pareto curves for different mechanisms are plotted in Figure 3. As seen in the figure, the Pareto curve of TICNet is positioned

Method	SW	Rev	UE	Rev+γUE
GSP and Fixed Positions	1.184	0.728	1.294	1.375
IAS	1.196	0.922	1.201	1.522
TICNet	**1.282**	**1.018**	**1.384**	**1.710**†

Table 3: The results of experiments for real-world dataset. The best performance is highlighted in bold. The regret of mechanism generated by the TICNet is less than 0.001. Symbol "†" indicates a statistically significant improvement in a paired t-test at $p < 0.05$ level.

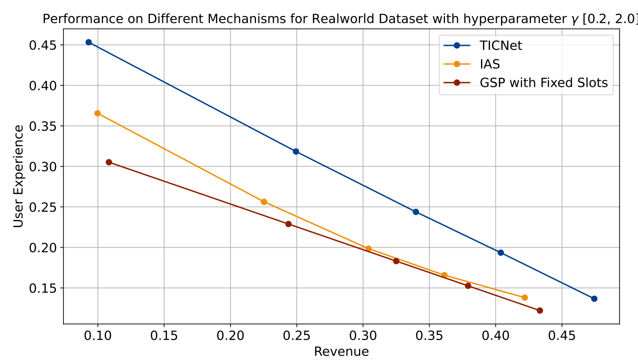

Figure 3: The Pareto curves of TICNet and other baseline mechanisms on real-world dataset.

above those of the other mechanisms, indicating that TICNet effectively balances revenue and GMV, thereby generating the optimal mixed list.

6 CONCLUSION

In this paper, we focus on designing the optimal integrated contextual mechanism, which blends sponsored advertisements and organic items into a mixed list. To address this problem, we propose TICNet, a transformer encoder-based neural network that maintains permutation equivariance and exhibits strong generalization performance, supported by theoretical proof. We then conduct a series of experiments to demonstrate TICNet's effectiveness compared to baseline mechanisms. In the future, it would be delightful to explore other mechanisms in the realm of integrated mechanisms.

REFERENCES

[1] Michael Albert, Vincent Conitzer, and Giuseppe Lopomo. 2015. Assessing the robustness of Cremer-McLean with automated mechanism design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29.

[2] Maria-Florina F Balcan, Tuomas Sandholm, and Ellen Vitercik. 2016. Sample complexity of automated mechanism design. *Advances in Neural Information Processing Systems* 29 (2016).

[3] Yang Cai and Constantinos Daskalakis. 2017. Learning multi-item auctions with (or without) samples. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 516–527.

[4] Carlos Carrion, Zenan Wang, Harikesh Nair, Xianghong Luo, Yulin Lei, Peiqin Gu, Xiliang Lin, Wenlong Chen, Junsheng Jin, Fanan Zhu, et al. 2023. Blending advertising with organic content in e-commerce via virtual bids. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 15476–15484.

[5] Dagui Chen, Qi Yan, Chunjie Chen, Zhenzhe Zheng, Yangsu Liu, Zhenjia Ma, Chuan Yu, Jian Xu, and Bo Zheng. 2022. Hierarchically constrained adaptive ad exposure in feeds. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 3003–3012.

[6] Edward H Clarke. 1971. Multipart pricing of public goods. *Public choice* (1971), 17–33.

[7] Richard Cole and Tim Roughgarden. 2014. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. 243–252.

[8] Michael Curry, Ping-Yeh Chiang, Tom Goldstein, and John Dickerson. 2020. Certifying strategyproof auction networks. *Advances in Neural Information Processing Systems* 33 (2020), 4987–4998.

[9] Michael Curry, Tuomas Sandholm, and John Dickerson. 2023. Differentiable economics for randomized affine maximizer auctions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. 2633–2641.

[10] Zhijian Duan, Haoran Sun, Yurong Chen, and Xiaotie Deng. 2024. A scalable neural network for DSIC affine maximizer auction design. *Advances in Neural Information Processing Systems* 36 (2024).

[11] Zhijian Duan, Jingwu Tang, Yutong Yin, Zhe Feng, Xiang Yan, Manzil Zaheer, and Xiaotie Deng. 2022. A context-integrated transformer-based neural network for auction design. In *International Conference on Machine Learning*. PMLR, 5609–5626.

[12] Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C Parkes, and Sai Srivatsa Ravindranath. 2024. Optimal auctions through deep learning: Advances in differentiable economics. *J. ACM* 71, 1 (2024), 1–53.

[13] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.

[14] Zhe Feng, Harikrishna Narasimhan, and David C Parkes. 2018. Deep learning for revenue-optimal auctions with budgets. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems*. 354–362.

[15] Noah Golowich, Harikrishna Narasimhan, and David C Parkes. 2018. Deep Learning for Multi-Facility Location Mechanism Design.. In *IJCAI*. 261–267.

[16] Theodore Groves. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* (1973), 617–631.

[17] Dmitry Ivanov, Iskander Safiulin, Igor Filippov, and Ksenia Balabaeva. 2022. Optimal-er auctions through attention. *Advances in Neural Information Processing Systems* 35 (2022), 34734–34747.

[18] Sébastien Lahaie and David M Pennock. 2007. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM Conference on Electronic Commerce*. 50–56.

[19] Ningyuan Li, Yunxuan Ma, Yang Zhao, Qian Wang, Zhilin Zhang, Chuan Yu, Jian Xu, Bo Zheng, and Xiaotie Deng. 2024. Ad vs Organic: Revisiting Incentive Compatible Mechanism Design in E-commerce Platforms. In *Proceedings of the ACM on Web Conference 2024*. 235–244.

[20] Weian Li, Qi Qi, Changjun Wang, and Changyuan Yu. 2023. Optimally integrating ad auction into e-commerce platforms. *Theoretical Computer Science* 976 (2023), 114141.

[21] Guogang Liao, Xuejian Li, Ze Wang, Fan Yang, Muzhi Guan, Bingqi Zhu, Yongkang Wang, Xingxing Wang, and Dong Wang. 2022. NMA: neural multi-slot auctions with externalities for online advertising. *arXiv preprint arXiv:2205.10018* (2022).

[22] Guogang Liao, Ze Wang, Xiaoxu Wu, Xiaowen Shi, Chuheng Zhang, Yongkang Wang, Xingxing Wang, and Dong Wang. 2022. Cross dqn: Cross deep q network for ads allocation in feed. In *Proceedings of the ACM Web Conference 2022*. 401–409.

[23] Anton Likhodedov and Tuomas Sandholm. 2004. Methods for boosting revenue in combinatorial auctions. In *AAAI*. 232–237.

[24] Anton Likhodedov, Tuomas Sandholm, et al. 2005. Approximating revenue-maximizing combinatorial auctions. In *AAAI*, Vol. 5. 267–274.

[25] Xiangyu Liu, Chuan Yu, Zhilin Zhang, Zhenzhe Zheng, Yu Rong, Hongtao Lv, Da Huo, Yiqing Wang, Dagui Chen, Jian Xu, et al. 2021. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3354–3364.

[26] Brendan Lucier, Renato Paes Leme, and Éva Tardos. 2012. On revenue in the generalized second price auction. In *Proceedings of the 21st international conference on World Wide Web*. 361–370.

[27] Yuchao Ma, Weian Li, Wanzhi Zhang, Yahui Lei, Zhicheng Zhang, Qi Qi, Qiang Liu, and Xingxing Wang. 2024. Joint Bidding in Ad Auctions. In *Annual Conference on Theory and Applications of Models of Computation*. Springer, 344–354.

[28] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.

[29] Neehar Peri, Michael Curry, Samuel Dooley, and John Dickerson. 2021. Preferencenet: Encoding human preferences in auction design with deep learning. *Advances in Neural Information Processing Systems* 34 (2021), 17532–17542.

[30] Jad Rahme, Samy Jelassi, Joan Bruna, and S Matthew Weinberg. 2021. A permutation-equivariant neural network architecture for auction design. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 5664–5672.

[31] Jad Rahme, Samy Jelassi, and S Matthew Weinberg. 2020. Auction learning as a two-player game. *arXiv preprint arXiv:2006.05684* (2020).

[32] Tuomas Sandholm and Anton Likhodedov. 2015. Automated design of revenue-maximizing combinatorial auctions. *Operations Research* 63, 5 (2015), 1000–1025.

[33] David RM Thompson and Kevin Leyton-Brown. 2013. Revenue optimization in the generalized second-price auction. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 837–852.

[34] Hal R Varian. 2007. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.

[35] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.

[36] Weixun Wang, Junqi Jin, Jianye Hao, Chunjie Chen, Chuan Yu, Weinan Zhang, Jun Wang, Xiaotian Hao, Yixi Wang, Han Li, et al. 2019. Learning adaptive display exposure for real-time advertising. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2595–2603.

[37] Ruobing Xie, Shaoliang Zhang, Rui Wang, Feng Xia, and Leyu Lin. 2021. Hierarchical reinforcement learning for integrated recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 4521–4528.

[38] Jinyun Yan, Zhiyuan Xu, Birjodh Tiwana, and Shaunak Chatterjee. 2020. Ads allocation in feed via constrained optimization. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 3386–3394.

[39] Zhen Zhang, Weian Li, Yahui Lei, Bingzhe Wang, Zhicheng Zhang, Qi Qi, Qiang Liu, and Xingxing Wang. 2024. Joint Auction in the Online Advertising Market. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4362–4373.

[40] Zhilin Zhang, Xiangyu Liu, Zhenzhe Zheng, Chenrui Zhang, Miao Xu, Junwei Pan, Chuan Yu, Fan Wu, Jian Xu, and Kun Gai. 2021. Optimizing multiple performance metrics with deep GSP auctions for e-commerce advertising. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 993–1001.

[41] Xiangyu Zhao, Changsheng Gu, Haoshenglun Zhang, Xiwang Yang, Xiaobing Liu, Jiliang Tang, and Hui Liu. 2021. Dear: Deep reinforcement learning for online advertising impression in recommender systems. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 750–758.

[42] Xiangyu Zhao, Xudong Zheng, Xiwang Yang, Xiaobing Liu, and Jiliang Tang. 2020. Jointly learning to recommend and advertise. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 3319–3327.

A PROOF OF THEOREM 1

The proof of Theorem 1 combines covering numbers with a concentration inequality from [11].

A.1 Basic Definitions

DEFINITION 3 ($l_{\infty,1}$-DISTANCE). *Let Q be a feature space and $\mathcal{W}$ be a space of functions from Q to $\mathbb{R}^n$. The $l_{\infty,1}$-distance on the space $\mathcal{F}$ is defined as $l_{\infty,1}(f,c) = \max_{q\in Q}(\sum_{i=1}^n |f_i(x) - c_i(x)|)$.*

DEFINITION 4 (l_∞-DISTANCE). *Let Q be a feature space and $\mathcal{W}$ be a space of functions from Q to $\mathbb{R}^n$. The l_∞-distance on the space $\mathcal{F}$ is defined as $l_\infty(f,c) = \max_{q\in Q}|f_i(x) - c_i(x)|$.*

DEFINITION 5 (COVERING NUMBER). *Let $\mathcal{N}_{\infty,1}(\mathcal{W}, r)$ be the minimum number of balls with radius r that can cover $\mathcal{W}$ under $l_{\infty,1}$-distance, and $\mathcal{N}_\infty(\mathcal{W}, r)$ be the minimum number of balls with radius r that can cover $\mathcal{W}$ under l_∞-distance.*

Let $\mathcal{U}_i$ be the class of utility functions for advertiser $i \in [m]$, i.e.,

$$\mathcal{U}_i = \Big\{u_i : \mathcal{V}_i \times \mathcal{V} \times \mathcal{G} \times \mathcal{X}^m \times \mathcal{Y}^n \times \mathcal{Z}^K \to \mathbb{R} \Big| \\ u_i(v_i; \mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = v_i \cdot a_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) - p_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\Big\}.$$

Similarly, let $\mathcal{U}$ be the class of utility profiles of all advertisers over $\mathcal{M}$. Define the $\ell_{\infty,1}$-distance between two utility profiles u and u' according to Definition 3 and $\mathcal{N}_{\infty,1}(\mathcal{U}, r)$ as the covering number under such $\ell_{\infty,1}$-distance. We also define the ℓ_∞-distance and $\mathcal{N}_\infty(\mathcal{U}_i, r)$ as the r-covering number of $\mathcal{U}_i$ under ℓ_∞-distance.

Let $\mathrm{RGT}_i \circ \mathcal{U}_i$ be the class of all regret functions for advertiser $i \in [m]$, i.e.,

$$\mathrm{RGT}_i \circ \mathcal{U}_i = \Big\{rgt_i : \mathcal{V}_i \times \mathcal{V} \times \mathcal{G} \times \mathcal{X}^m \times \mathcal{Y}^n \times \mathcal{Z}^K \to \mathbb{R} \Big| \\ rgt_i(v_i; \mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = \max_{b_i' \in \mathcal{V}_i} u_i(v_i; (b_i', \mathbf{b}_{-\mathbf{i}}), \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) - u_i(v_i, \mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\Big\}.$$

Similarly, we denote $\mathrm{RGT} \circ \mathcal{U}$ by the class of regret functions of all advertisers over $\mathcal{M}$. Define the $\ell_{\infty,1}$-distance between two regret profiles rgt and rgt' according to Definition 3 and $\mathcal{N}_{\infty,1}(\mathrm{RGT}\circ\mathcal{U}, r)$ as the covering number under such $\ell_{\infty,1}$-distance.

Let $\mathcal{P}$ be the class of all the profiles of payment functions for advertiser $i \in [m]$, i.e.,

$$\mathcal{P}_i = \{p_i : \mathcal{V} \times \mathcal{G} \times \mathcal{X}^m \times \mathcal{Y}^n \times \mathcal{Z}^K \to \mathbb{R}_{\geq 0} | p \in \mathcal{P}\}$$

Similarly, we denote $\mathcal{P}$ by the class of payment functions of all advertisers over $\mathcal{M}$. Then let $\mathcal{N}_{\infty,1}(\mathcal{P}, r)$ be the r-covering number of $\mathcal{P}$ under the $\ell_{\infty,1}$-distance and $\mathcal{N}_\infty(\mathcal{P}_i, r)$ be the r-covering number of $\mathcal{P}_i$ under the ℓ_∞-distance.

Let $\mathcal{G}$ be the class of all the profiles of user experience for winning advertiser (or organic item) $i \in [m] \cup [n]$, i.e.,

$$\mathcal{G}_i = \{\hat{g}_i : \mathcal{V} \times \mathcal{G} \times \mathcal{X}^m \times \mathcal{Y}^n \times \mathcal{Z}^K \to \mathbb{R}_{\geq 0} | \\ \hat{g}_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) = g_i \cdot a_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\}$$

Similarly, we denote $\mathcal{G}$ by the class of winning user experience functions of all advertisers and organic items over $\mathcal{M}$. Then let $\mathcal{N}_{\infty,1}(\mathcal{G}, r)$ be the r-covering number of $\mathcal{G}$ under the $\ell_{\infty,1}$-distance and $\mathcal{G}_\infty(\mathcal{G}_i, r)$ be the r-covering number of $\mathcal{G}_i$ under the ℓ_∞-distance.

A.2 Technical Lemmas

With the definitions above, we then introduce a few of important lemmas. Notably, the Lemma 1 has been proved by [11]. We recall it here to make our paper completed.

LEMMA 1 ([11]). *Let $\mathcal{S} = \{z_1, \cdots, z_L\}$ be a set of i.i.d. samples drawn from the distribution $\mathcal{D}$ over $\mathcal{Z}$. Assume that $\mathcal{F}$ 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 $z \in \mathcal{Z}$. Then we have*

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

PROOF. Let $\mathcal{F}_r$ be the minimum function class that r-covers $\mathcal{F}$. Then for any $f \in \mathcal{F}$, there exists $f_r \in \mathcal{F}_r$ such that $|f(z) - f_r(z)| \leq r$. For the simplicity, we denote $\frac{1}{L}\sum_{i=1}^L f(z_i)$ by $\mathbb{E}_\mathcal{S}[f(z)]$. Then for all $\epsilon > 0$, set $r = \frac{\epsilon}{3}$, we can get

$$\begin{aligned}
&\mathbb{P}\Big[\exists f \in \mathcal{F}, \Big|\mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}[f(z)]\Big| > \epsilon\Big] \\
=&\mathbb{P}\Big[\exists f \in \mathcal{F}, \Big|\mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}_\mathcal{S}[f_r(z)] + \mathbb{E}_\mathcal{S}[f_r(z)] - \mathbb{E}[f_r(z)] + \\
&\mathbb{E}[f_r(z)] - \mathbb{E}[f(z)]\Big| > \epsilon\Big] \\
\leq&\mathbb{P}\Big[\exists f \in \mathcal{F}, \Big|\mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}_\mathcal{S}[f_r(z)]\Big| + \Big|\mathbb{E}_\mathcal{S}[f_r(z)] - \mathbb{E}[f_r(z)]\Big| + \\
&\Big|\mathbb{E}[f_r(z)] - \mathbb{E}[f(z)]\Big| > \epsilon\Big] \\
\leq&\mathbb{P}\Big[\exists f \in \mathcal{F}, r + \Big|\mathbb{E}_\mathcal{S}[f_r(z)] - \mathbb{E}[f_r(z)]\Big| + r > \epsilon\Big] \\
\leq&\mathbb{P}\Big[\exists f_r \in \mathcal{F}_{\frac{\epsilon}{3}}, \Big|\mathbb{E}_\mathcal{S}[f_r(z)] - \mathbb{E}[f_r(z)]\Big| > \frac{1}{3}\epsilon\Big], \\
\leq&\mathcal{N}_\infty(\mathcal{F}, \frac{\epsilon}{3})\mathbb{P}\Big[\Big|\mathbb{E}_\mathcal{S}[f(z)] - \mathbb{E}[f(z)]\Big| > \frac{1}{3}\epsilon\Big] \\
\leq&2\mathcal{N}_\infty(\mathcal{F}, \frac{\epsilon}{3}) \exp\left(-\frac{2L\epsilon^2}{9(b-a)^2}\right),
\end{aligned} \tag{5}$$

where the last inequality is supported by Hoeffding Inequality. □

The following three lemmas provide the covering numbers bound for payment, regret and user experience.

LEMMA 2. *$\mathcal{N}_{\infty,1}(\mathcal{P}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \epsilon)$.*

PROOF. By the definition of the covering number for the auction class $\mathcal{M}$, there exists a cover $\hat{\mathcal{M}}$ for $\mathcal{M}$ of size $|\hat{\mathcal{M}}| \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \epsilon)$ such that for any $(\mathbf{a}, \mathbf{p}) \in \mathcal{M}$, there is a $(\hat{\mathbf{a}}, \hat{\mathbf{p}}) \in \hat{\mathcal{M}}$ for all $\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}$,

$$\sum_{i=1}^{m+n}\Big|\mathbf{a}_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) - \hat{\mathbf{a}}_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\Big| + \\ \sum_{i=1}^{m}\Big|\mathbf{p}_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) - \hat{\mathbf{p}}_i(\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\Big| \leq \epsilon.$$

As a result, we can have $\hat{\mathcal{P}} = \{\hat{\mathbf{p}} | (\hat{\mathbf{a}}, \hat{\mathbf{p}}) \in \hat{\mathcal{M}}\}$, then for any $\mathbf{p} \in \mathcal{P}$, there exist a $\hat{\mathbf{p}} \in \hat{\mathcal{P}}$, for all $\mathbf{b}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}$,

$$\sum_i \Big|\mathbf{p}_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z}) - \hat{\mathbf{p}}_i(\mathbf{v}, \mathbf{g}, \mathbf{x}, \mathbf{y}, \mathbf{z})\Big| \leq \epsilon$$

Therefore, we have $\mathcal{N}_{\infty,1}(\mathcal{P}, \epsilon) \leq \mathcal{N}_{\infty,1}(\mathcal{M}, \epsilon)$. □

Lemma 3. $\mathcal{N}_{\infty,1}(\mathrm{RGT}\circ\mathcal{U},\epsilon) \le \mathcal{N}_{\infty,1}(\mathcal{M},\frac{\epsilon}{2m})$.

Proof. By the definition of covering number $\mathcal{N}_{\infty,1}(\mathcal{U},r)$, there exists a cover $\hat{\mathcal{U}}$ with size at most $\mathcal{N}_{\infty,1}(\mathcal{U},\epsilon/2)$ such that for any $u\in\mathcal{U}$, there is a $\hat{u}\in\hat{\mathcal{U}}$ with

$$\max_{\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}}\sum_{i=1}^{n}\left|u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\right|\le\frac{\epsilon}{2}.$$

For any $u\in\mathcal{U}$, taking $\hat{u}\in\hat{\mathcal{U}}$ satisfying the above condition, then for any $\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}$, we have

$$\begin{aligned}
&\Big|\max_{b_i'\in\mathcal{V}_i}\big(u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big)-\\
&\quad\max_{\bar{b}_i\in\mathcal{V}_i}\big(\hat{u}_i(v_i;(\bar{b}_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big)\Big|\\
\le&\Big|\max_{b_i'\in\mathcal{V}_i}\big(u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\max_{\bar{b}_i\in\mathcal{V}_i}\big(\hat{u}_i(v_i;(\bar{b}_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\\
&\quad+\hat{u}_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
\le&\Big|\max_{b_i'\in\mathcal{V}_i}\big(u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\max_{\bar{b}_i\in\mathcal{V}_i}\big(\hat{u}_i(v_i;(\bar{b}_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
&\quad+\Big|\hat{u}_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
\le&\Big|\max_{b_i'\in\mathcal{V}_i}\big(u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\max_{\bar{b}_i\in\mathcal{V}_i}\big(\hat{u}_i(v_i;(\bar{b}_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
&\quad+\max_{b_i'\in\mathcal{V}_i}\Big|\hat{u}_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|.
\end{aligned}$$

Then we denote $b_i^*\in\arg\max_{b_i'}u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})$ and $\hat{b}_i^*\in\arg\max_{\bar{v}_i}\hat{u}_i(v_i,(\bar{v}_i,v_{-i}),x,y)$, then

$$\begin{aligned}
&\max_{b_i'\in\mathcal{V}_i}\big(u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\max_{\bar{b}_i\in\mathcal{V}_i}\big(\hat{u}_i(v_i;(\bar{b}_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\\
=&u_i(v_i;(b_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i;(\bar{b}_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\\
\le&u_i(v_i;(b_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i;(b_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\\
\le&\max_{b_i'\in\mathcal{V}_i}\Big|u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
&\max_{\bar{b}_i\in\mathcal{V}_i}\big(\hat{u}_i(v_i;(\bar{b}_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\max_{b_i'\in\mathcal{V}_i}\big(u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\\
=&\hat{u}_i(v_i;(\bar{b}_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(b_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}))\\
\le&\hat{u}_i(v_i;(b_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(b_i^*,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\\
\le&\max_{b_i'\in\mathcal{V}_i}\Big|u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|.
\end{aligned}$$

Thus,

$$\begin{aligned}
&\max_{b_i'\in\mathcal{V}_i}\big(u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big)\\
&-\max_{\bar{b}_i\in\mathcal{V}_i}\big(\hat{u}_i(v_i;(\bar{b}_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i;(v_i,\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big)\\
\le&2\max_{b_i'\in\mathcal{V}_i}\Big|\hat{u}_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-u_i(v_i;(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|.
\end{aligned}$$

Summing the inequalities by i, we can conclude that $\mathcal{N}_{\infty,1}(\mathrm{RGT}\circ\mathcal{U},\epsilon)\le\mathcal{N}_{\infty,1}(\mathcal{U},\frac{\epsilon}{2})$.

Then we prove that $\mathcal{N}_{\infty,1}(\mathcal{U},\epsilon)\le\mathcal{N}_{\infty,1}(\mathcal{M},\frac{\epsilon}{n})$.

Then by the definition of the covering number for the auction class $\mathcal{M}$, there exists a cover $\hat{\mathcal{M}}$ for $\mathcal{M}$ of size $|\hat{\mathcal{M}}|\le\mathcal{N}_{\infty,1}(\mathcal{M},\frac{\epsilon}{n})$ such that for any $(\mathbf{a},\mathbf{p})\in\mathcal{M}$, there is a $(\hat{\mathbf{a}},\hat{\mathbf{p}})\in\hat{\mathcal{M}}$ for all $\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}$,

$$\begin{aligned}
&\sum_i\Big[\big|\mathbf{a}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{\mathbf{a}}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big|+\\
&\qquad\big|\mathbf{p}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{\mathbf{p}}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big|\Big]\le\frac{\epsilon}{n}.
\end{aligned}$$

For all $\mathbf{v}\in\mathcal{V}$, $b_i'\in\mathcal{V}_i$,

$$\begin{aligned}
&\Big|u_i(v_i,(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i,(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
=&\Big|v_i\cdot a_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-p_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\\
&-v_i\cdot\hat{a}_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})+\hat{p}_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
\le&\Big|v_i\cdot[a_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{a}_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})]\Big|\\
&+\Big|p_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{p}_i((b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\\
\le&\frac{\epsilon}{m}.
\end{aligned}$$

Thus,

$$\sum_{i=1}^{m}\Big|u_i(v_i,(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{u}_i(v_i,(b_i',\mathbf{b}_{-i}),\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big|\le m\cdot\frac{\epsilon}{m}=\epsilon.$$

This completes the proof that $\mathcal{N}_{\infty,1}(\mathcal{U},\epsilon)\le\mathcal{N}_{\infty,1}(\mathcal{M},\frac{\epsilon}{m})$.

Therefore,

$$\mathcal{N}_{\infty,1}(\mathrm{RGT}\circ\mathcal{U},\epsilon)\le\mathcal{N}_{\infty,1}(\mathcal{U},\frac{\epsilon}{2})\le\mathcal{N}_{\infty,1}(\mathcal{M},\frac{\epsilon}{2m}).$$

This completes the proof of Lemma 3. □

Lemma 4. $\mathcal{N}_{\infty,1}(\mathrm{RGT}\circ\mathcal{U},\epsilon)\le\mathcal{N}_{\infty,1}(\mathcal{M},\frac{\epsilon}{2m})$.

Proof. Similar to the proof of Lemma 2, there exists a cover $\hat{\mathcal{M}}$ for $\mathcal{M}$ of size $|\hat{\mathcal{M}}|\le\mathcal{N}_{\infty,1}(\mathcal{M},\epsilon)$ such that for any $(\mathbf{a},\mathbf{p})\in\mathcal{M}$, there is a $(\hat{\mathbf{a}},\hat{\mathbf{p}})\in\hat{\mathcal{M}}$ for all $\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}$,

$$\begin{aligned}
&\sum_{i=1}^{m+n}\big|\mathbf{a}_i(\mathbf{b},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{\mathbf{a}}_i(\mathbf{b},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big|+\\
&\sum_{i=1}^{m}\big|\mathbf{p}_i(\mathbf{b},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{\mathbf{p}}_i(\mathbf{b},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\big|\le\epsilon.
\end{aligned}$$

As a result, we can have $\hat{\mathcal{G}}=\{\hat{\mathbf{p}}|(\hat{\mathbf{a}},\hat{\mathbf{p}})\in\hat{\mathcal{M}}\}$, then for any $\mathbf{g}\in\mathcal{G}$, there exist a $\hat{\mathbf{g}}\in\hat{\mathcal{G}}$, for all $\mathbf{b},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z}$,

$$\begin{aligned}
&\sum_i^{m+n}\big|[\mathbf{a}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{\mathbf{a}}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})]g_i\big|\\
\le&\sum_i^{m+n}\big|[\mathbf{a}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})-\hat{\mathbf{a}}_i(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})]\big|\\
\le&\epsilon
\end{aligned}$$

Therefore, we have $\mathcal{N}_{\infty,1}(\mathcal{G},\epsilon)\le\mathcal{N}_{\infty,1}(\mathcal{M},\epsilon)$. □

A.3 Proof of Theorem 1

Proof of Theorem 1. For all $\epsilon, \delta \in (0,1)$, when

$$L \geq \frac{9n^2}{2\epsilon^2}\left(\ln\frac{4}{\delta} + \ln \mathcal{N}_{\infty,1}(\mathcal{M}, \frac{\epsilon}{6n})\right),$$

Combining Lemma 1 and Lemma 2 together, we get

$$\begin{aligned}&\mathbb{P}\Big[\exists(g^w, p^w) \in \mathcal{M}, \Big|\mathbb{E}_{(\mathrm{v,g,x,y,z})\sim\mathcal{D}_{\mathrm{v,g,x,y,z}}}\Big[\sum_{i=1}^{m} p_i^w(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big] - \\ &\qquad \frac{1}{L}\sum_{\ell=1}^{L}\sum_{i=1}^{m} p_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})\Big| > \epsilon\Big] \\ &\leq 2\mathcal{N}_\infty(\mathcal{P}, \frac{\epsilon}{3})\exp(-\frac{2L\epsilon^2}{9m^2}) \\ &\leq 2\mathcal{N}_\infty(\mathcal{M}, \frac{\epsilon}{3})\exp(-\frac{2L\epsilon^2}{9m^2}) \\ &\leq 2\mathcal{N}_\infty(\mathcal{M}, \frac{\epsilon}{6m})\exp(-\frac{2L\epsilon^2}{9m^2}) \\ &\leq 2\mathcal{N}_\infty(\mathcal{M}, \frac{\epsilon}{6(m+n)})\exp(-\frac{2L\epsilon^2}{9(m+n)^2}) \\ &\leq \frac{\delta}{3}.\end{aligned} \tag{6}$$

Similarly, combining Lemma 1 and Lemma 3 together, we have

$$\begin{aligned}&\mathbb{P}\Big[\exists(g^w, p^w) \in \mathcal{M}, \Big|\mathbb{E}_{(\mathrm{v,g,x,y,z})\sim\mathcal{D}_{\mathrm{v,g,x,y,z}}}\Big[\sum_{i=1}^{m} rgt_i(w)\Big] - \\ &\qquad \sum_{i=1}^{m} \widehat{rgt}_i(w)\Big| > \epsilon\Big] \\ &\leq 2\mathcal{N}_\infty(\mathrm{RGT}\circ\mathcal{U}, \frac{\epsilon}{3})\exp(-\frac{2L\epsilon^2}{9m^2}) \\ &\leq 2\mathcal{N}_\infty(\mathcal{M}, \frac{\epsilon}{6m})\exp(-\frac{2L\epsilon^2}{9m^2}) \\ &\leq 2\mathcal{N}_\infty(\mathcal{M}, \frac{\epsilon}{6(m+n)})\exp(-\frac{2L\epsilon^2}{9(m+n)^2}) \\ &\leq \frac{\delta}{3}.\end{aligned} \tag{7}$$

Combining Lemma 1 and Lemma 4 together, we get

$$\begin{aligned}&\mathbb{P}\Big[\exists(\mathbf{a}^w, \mathbf{p}^w) \in \mathcal{M}, \Big|\mathbb{E}_{(\mathrm{v,g,x,y,z})\sim\mathcal{D}_{\mathrm{v,g,x,y,z}}}\Big[\sum_{i=1}^{m+n} g_i \cdot a_i^w(\mathbf{v},\mathbf{g},\mathbf{x},\mathbf{y},\mathbf{z})\Big] - \\ &\qquad \frac{1}{L}\sum_{\ell=1}^{L}\sum_{i=1}^{m+n} g_i \cdot a_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})\Big| > \epsilon\Big] \\ &\leq 2\mathcal{N}_\infty(\mathcal{G}, \frac{\epsilon}{3})\exp(-\frac{2L\epsilon^2}{9(m+n)^2}) \\ &\leq 2\mathcal{N}_\infty(\mathcal{M}, \frac{\epsilon}{3})\exp(-\frac{2L\epsilon^2}{9(m+n)^2}) \\ &\leq 2\mathcal{N}_\infty(\mathcal{M}, \frac{\epsilon}{6(m+n)})\exp(-\frac{2L\epsilon^2}{9(m+n)^2}) \\ &\leq \frac{\delta}{3}.\end{aligned} \tag{8}$$

Combining Equation (6), Equation (7), Equation (8) and the Union Bound, with probability at most $\frac{\delta}{3} + \frac{\delta}{3} + \frac{\delta}{3} = \delta$, one of the three events of Equation (6), Equation (7) and Equation (8) happens. Therefore, with probability at least $1-\delta$, Equation (2), Equation (3) and Equation (4) hold. We complete the proof of Theorem 1. □

B PROOF OF THEOREM 2

Definition 6 (Permutation-equivariant [30]). *An auction mechanism $(\mathbf{a}^w, \mathbf{p}^w)$ is permutation-equivariant if for any two permutation matrices $\Pi_n \in \{0,1\}^{n\times n}$ and $\Pi_m \in \{0,1\}^{m\times m}$, and any input V, both $\mathbf{a}^w(\Pi_n V \Pi_m) = \Pi_n \mathbf{a}(V)\Pi_m$ and $\mathbf{p}^w(\Pi_n V \Pi_m) = \Pi_n \mathbf{p}(V)$ hold.*

Proof of Theorem 2. With the bidding profile $\mathbf{b} \in \mathbb{R}^m$, the TICNet's input layer in Section 4.2 transfers bids $\mathbf{b} \in \mathbb{R}^m$ into $\mathbf{e} \in \mathbb{R}^{(m+n)\times K}$ through:

$$\begin{cases} e_{ik} = b_i\alpha_k, & \forall i \in [m],\ \forall k \in [K]; \\ e_{ik} = 0, & \forall i \in [n],\ \forall k \in [K]. \end{cases}$$

Similarly, with the user experience profile $\mathbf{g} \in \mathbb{R}^{m+n}$, the TICNet's input layer transfers user experience $\mathbf{g} \in \mathbb{R}^{m+n}$ into $\mathbf{q} \in \mathbb{R}^{(m+n)\times K}$ through:

$$q_{jk} = g_j\alpha_k, \quad \forall j \in [m]\cup[n],\ \forall k \in [K]$$

Notably, matrices $\mathbf{e}$ and $\mathbf{q}$ contain information about both advertisers and organic items. Correspondingly, we combine ad-contexts $\mathbf{x} \in \mathbb{R}^{m\times d_c}$ and organic-contexts $\mathbf{y} \in \mathbb{R}^{n\times d_c}$ along the first dimension into participant-contexts $\mathbf{s} \in \mathbb{R}^{(m+n)\times d_c}$.

Then with above generated matrices and slot-contexts $\mathbf{z} \in \mathbb{R}^{K\times d_z}$, we have

$$\mathbf{a}^w(\Pi_{m+n}\mathbf{e}\Pi_K, \Pi_{m+n}\mathbf{q}\Pi_K, \Pi_{m+n}\mathbf{s}, \Pi_K^T\mathbf{z}) = \Pi_{m+n}\mathbf{a}^w(\mathbf{e},\mathbf{q},\mathbf{s},\mathbf{z})\Pi_K,$$

and

$$\mathbf{p}^w(\Pi_{m+n}\mathbf{e}\Pi_K, \Pi_{m+n}\mathbf{q}\Pi_K, \Pi_{m+n}\mathbf{s}, \Pi_K^T\mathbf{z}) = \Pi_{m+n}\mathbf{a}^w(\mathbf{e},\mathbf{q},\mathbf{s},\mathbf{z}).$$

In addition, we adopt the transformer encoder architecture, excluding the positional encoding layer. It consists of multiple layers, each of which contains multi-head attention and a feed-forward neural network. These two components are not affected by the order of input, thus maintains permutation-equivariant.

To sum up, TICNet is permutation-equivariant. □

C DETAILED TRAINING PROCEDURE

C.1 Traning Settings

We adopt the augmented Lagrangian algorithm to convert original constrained optimization problem into an unconstrained optimization problem within the space of parameter $w \in d_w$:

$$\begin{aligned}C_\rho(w;\boldsymbol{\lambda}) = &-\frac{1}{L}\sum_{\ell=1}^{L}\Big[\sum_{i\in[m]} p_i^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)}) + \\ &\gamma\cdot\sum_{j\in[m]\cup[n]} g_j\cdot a_j^w(\mathbf{v}^{(\ell)},\mathbf{g}^{(\ell)},\mathbf{x}^{(\ell)},\mathbf{y}^{(\ell)},\mathbf{z}^{(\ell)})\Big] \\ &+\sum_{i=1}^{m}\lambda_i\cdot\widehat{rgt}_i(w) + \frac{\rho}{2}\sum_{i=1}^{m}(\widehat{rgt}_i(w))^2,\end{aligned}$$

Algorithm 1: TICNet Training

Input: Minibatches $\mathcal{L}_1, ..., \mathcal{L}_T$ of size B;
Parameters:
$\forall t \in \{1, \ldots, T\}, \rho_t > 0, \varphi > 0, \eta > 0, \Gamma \in \mathbb{N}, T \in \mathbb{N}, H \in \mathbb{N}$;
Initialize: $w^0 \in \mathbb{R}^d, \lambda^0 \in \mathbb{R}^m$;
for $t = 0$ **to** T **do**
  Receive minibatch $\mathcal{L}_t = \{(\mathbf{v}^{(1)}, \mathbf{g}^{(1)}, \mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \mathbf{z}^{(1)}), \cdots, (\mathbf{v}^{(B)}, \mathbf{g}^{(B)}, \mathbf{x}^{(B)}, \mathbf{y}^{(B)}, \mathbf{z}^{(B)})\}$;
  Initialize misreport $b_i'^{(\ell)} \in \mathcal{V}_i, \forall \ell \in \{1, \ldots, B\}, i \in [m]$;
  for $r = 0$ *to* Γ **do**
    $\forall \ell \in \{1, \ldots, B\}, i \in [m]$;
    $v_i'^{(\ell)} = v_i'^{(\ell)} + \varphi \nabla_{v_i'} \left[u_i^w \left(v_i^{(\ell)}; \left(b_i', \mathbf{v}_{-i}^{(\ell)} \right), \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)} \right) \right] \Big|_{v_i' = v_i'^{(\ell)}}$;
  end
  Compute regret gradient:$\forall \ell \in [B], i \in M$:
  $h_{\ell,i} = \nabla_w \Big[\max_{b_i' \in V_i} u_i^w (v_i^{(\ell)}; (b_i', \mathbf{v}_{-i}^{(\ell)}), \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) - u_i^w (v_i^{(\ell)}; \mathbf{v}^{(\ell)}, \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) \Big]$;
  Compute Lagrangian gradient and update w^t:
  $w^{t+1} \leftarrow w^t - \eta \nabla_w C_{\rho_t} \left(w^t, \lambda^t \right)$;
  Update Lagrange multipliers once in H iterations:
  if *t is a multiple of H* **then**
    $\lambda_i^{t+1} \leftarrow \lambda_i^t + \rho_t \widetilde{rgt}_i \left(w^{t+1} \right), \quad \forall i \in M$
  else
    $\boldsymbol{\lambda}^{t+1} \leftarrow \boldsymbol{\lambda}^t$
  end
end

where $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \cdots, \lambda_n) \in \mathbb{R}^n$ stands for the Lagarangian multipliers and $\rho > 0$ is the hyperparameter that control the quadratic term.

Afterward, we move on to divide the dataset $\mathcal{L}$ into minibatches of size B. We denote the total number of epoches by T. For each epoch $t \in \{1, 2, \cdots, T\}$, we randomly select a minibatch $\mathcal{L}_t$, i.e., $\mathcal{L}_t = \{(\mathbf{v}^{(1)}, \mathbf{g}^{(1)}, \mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \mathbf{z}^{(1)}), \cdots, \mathbf{b}^{(B)}, \mathbf{g}^{(B)}, \mathbf{x}^{(B)}, \mathbf{y}^{(B)}, \mathbf{z}^{(B)})\}$ and feed it to TICNet to train, until all the minibatches have been trained in this epoch. We will then redivide the dataset $\mathcal{L}$ into minibatches of size B and repeat training procedure until T.

For each minibatch, we first calculate the optimal misreports through gradient ascent. It is worth noting that since organic items do not engage in strategic behavior, we only take advertisers into account when computing the optimal misreports. Specifically, we compute the optimal misreport, for each advertiser i and corresponding profile ℓ, by taking Γ updates from a randomly initialized valuation, each update of the form

$$b_i'^{(\ell)} = b_i'^{(\ell)} + \mu \nabla_{b_i'} \left[u_i^w \left(v_i^{(\ell)}; \left(b_i', \mathbf{v}_{-i}^{(\ell)} \right), \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)} \right) \right] \Big|_{b_i' = b_i'^{(\ell)}},$$

with hyperparameter $\mu > 0$. We then update the Lagrange multipliers and model parameters alternately:

$$\begin{cases} w^{\text{new}} \in \arg\min_w C_\rho(w^{\text{old}}, \boldsymbol{\rho}^{\text{old}}) \\ \lambda_i^{\text{new}} = \lambda_i^{\text{old}} + \rho \widehat{rgt}_i(w^{\text{new}}), \quad i \in [n] \end{cases}$$

We denote $\widetilde{rgt}_i$ as the empirical regret value calculated on minibatch $\mathcal{L}_t$. Then for fixed $\boldsymbol{\lambda}^t$, the gradient of C_ρ w.r.t. w is defined as:

$$\begin{aligned} &\nabla_w C_\rho(w; \boldsymbol{\lambda}^t) \\ &= -\frac{1}{B} \sum_{\ell=1}^{B} \Bigg[\sum_{i \in [m]} \nabla_w p_i^w(\mathbf{v}^{(\ell)}, \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) + \\ &\gamma \cdot \sum_{j \in [m] \cup [n]} g_j \nabla_w a_j^w(\mathbf{v}^{(\ell)}, \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) \Bigg] \\ &+ \sum_{\ell=1}^{B} [\sum_{i=1}^{m} \lambda_i^t h_{\ell,i} + \rho \sum_{i=1}^{m} \widetilde{rgt}_i(w) h_{\ell,i}], \end{aligned}$$

where

$$\begin{aligned} h_{\ell,i} = &\nabla_w \Bigg[\max_{b_i' \in V_i} u_i^w(v_i^{(\ell)}; (b_i', \mathbf{v}_{-i}^{(\ell)}), \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) - \\ &u_i^w(v_i^{(\ell)}; \mathbf{v}^{(\ell)}, \mathbf{g}^{(\ell)}, \mathbf{x}^{(\ell)}, \mathbf{y}^{(\ell)}, \mathbf{z}^{(\ell)}) \Bigg]. \end{aligned}$$

C.2 Training Algorithm of TICNet

D EXTENSIVE EXPERIMENTS ON SYNTHETIC DATASET

For the synthetic data, we have conducted one extensive experiment to validate the performance of TICNet and other baselines. The details of the setting is:

(D) 1 ad, 4 organic items and 3 slots with CTRs $\beta = (0.6, 0.3, 0.1)$, continuous ad-contexts, organic-contexts and slot-contexts, in which $\mathcal{X} = [-1, 1]^5$, $\mathcal{Y} = [0, 2]^5$ and $\mathcal{Z} = [-2, 2]^5$. The raw value profile of each ad is independently drawn from $U[0, 1]$. The raw user experience profiles of these 2 ads and 3 organic items are independently drawn from $U[0, 0.5]$ and $U[0.5, 1]$, respectively. Given contextual information $x_i \in \mathcal{X}, y_j \in \mathcal{Y}$ and $z_k \in \mathcal{Z}$, slot-wise value profile v_{ik}' is transferred from raw value profile v_i according to $v_i \cdot \text{Sigmoid}(x_i^T z_k)$, slot-wise user experience profile g_{jk}' is obtained from raw user experience profile g_j according to $g_j \cdot \text{Sigmoid}(y_j^T z_k)$.

The experimental results are shown in Table 4. We can conclude that the integrated mechanism generated by TICNet are capable of achieving higher value for the Rev+γUE metric compared to other baseline mechanisms, while keeping the regret value less than 0.001. These results demonstrate that TICNet is competent to integrate ads and organic items effectively.

To further validate the superiority of TICNet, we conduct experiments across various settings. These include different numbers of slots, different ratios of candidate ads to organic items and large scale. The detailed settings and corresponding results for these three experiments are outlined below:

Method	D: $1 \times 4 \times 3$			
	SW	Rev	UE	Rev+γUE
GSP and Fixed Positions	**0.582**	0	0.498	0.299
IAS	0.270	0.215	0.581	0.506
TICNet	0.303	**0.222**	**0.630**	**0.537**†

Table 4: The notation $1 \times 4 \times 3$ represents a setting where 1 advertiser and 4 organic items compete for 3 slots. In the metric Rev+γUE, the hyperparameter γ is set to 0.5 across different settings. For the GSP and Fixed Positions mechanisms, the first slot is reserved for candidate ads. The regret of mechanism generated by the TICNet is less than 0.001. The best performance is highlighted in bold. "†" indicates a statistically significant improvement in a paired t-test at $p < 0.05$ level.

D.1 Different Number of Slots

To demonstrate the performance of TICNet compared to other baseline methods with different slots, we use Setting C as the basic setting and adjust the number of slots and their corresponding CTR values. We conduct a total of four experiments for validation. The specific setup is as follows: Setting C1. two slots with CTRs α = (0.7, 0.3); Setting C2. three slots with CTRs α = (0.6, 0.3, 0.1); Setting C3. four slots with CTRs α = (0.5, 0.3, 0.15, 0.05); Setting C4. five slots with CTRs α = (0.4, 0.3, 0.15, 0.1, 0.05). The other settings keep unchanged compared with initial Setting C. The curves of the combination of final revenue and user experience as the number of slots varying are presented in Figure 4.

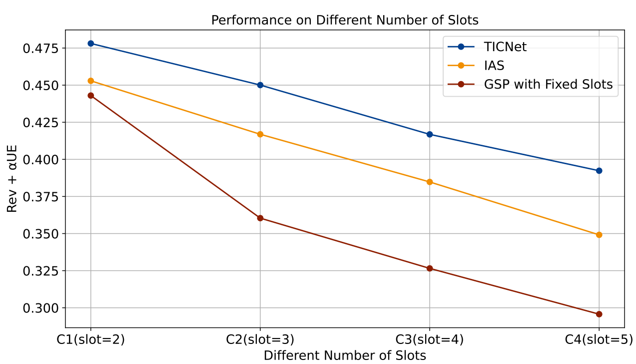

Figure 4: The performance of TICNet and other baseline mechanisms across different numbers of slots.

From Figure 4, we observe that TICNet outperforms all other mechanisms, indicating that TICNet maintains strong adaptability even when the number of slots changes.

D.2 Different Ratios of Candidate Ads to Organic Items

Based on Setting (D), we keep the total number of candidates constant while varying the ratio of ads to organic items to observe the changes in revenue and user experience. We conduct four experiments for validation, with the experimental setups as follows: Setting D1. one ad and four organic items; Setting D2. two ads and three organic items; Setting D3. three ads and two organic items; Setting D4. four ads and one organic item. The raw user experience profiles of these ads and organic items are independently drawn from $U[0, 0.5]$ and $U[0.5, 1]$, respectively. The other settings remain unchanged compared with initial Setting C. In all these four settings, we set γ =0.5. The curves showing the final revenue and user experience as functions of the ratio of ads to organic items are presented in Figure 5.

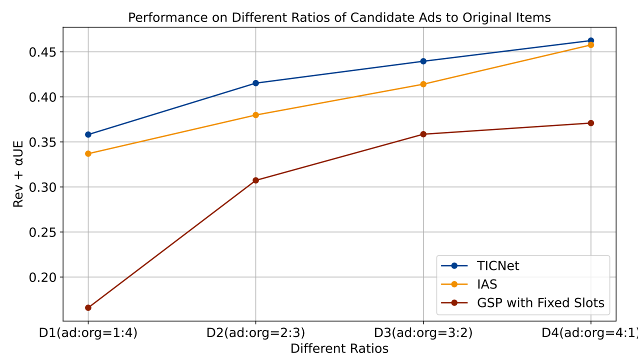

Figure 5: The performance of TICNet and other baseline mechanisms across different ratios of candidate ads to organic items.

From Figure 5, we can observe that TICNet outperforms other baseline mechanisms in terms of the Rev+γUE metric, demonstrating its superiority to generate mixed list.