

Contextual Generative Auction with Permutation-level Externalities for Online Advertising

Ruitao Zhu
Shanghai Jiao Tong University
Shanghai, China
sjtu_zrt@sjtu.edu.cn

Zhenjia Ma
Alibaba Group
Beijing, China
mazhenjia.mzj@alibaba-inc.com

Jie Zhang
Alibaba Group
Beijing, China
kongpan.zj@taobao.com

Yangsuo Liu
Alibaba Group
Beijing, China
liuyangsuo.lys@alibaba-inc.com

Chufeng Shi
Shanghai Jiao Tong University
Shanghai, China
cfshi99@sjtu.edu.cn

Jian Xu
Alibaba Group
Beijing, China
xiyu.xj@alibaba-inc.com

Fan Wu
Shanghai Jiao Tong University
Shanghai, China
fwu@cs.sjtu.edu.cn

Dagui Chen
Alibaba Group
Beijing, China
dagui.cdg@taobao.com

Zhenzhe Zheng*
Shanghai Jiao Tong University
Shanghai, China
zhengzhenzhe@sjtu.edu.cn

Bo Zheng
Alibaba Group
Beijing, China
bozheng@alibaba-inc.com

Abstract

Online advertising has become a core revenue driver for internet industry, with ad auctions playing a crucial role in ensuring platform revenue and advertiser incentives. Classical auction mechanisms, such as GSP, rely on the independent CTR assumption and fail to account for the interplay among the displayed items, also called as externalities in economics. Recent advancements in learning-based auctions enable the encoding of high-dimensional contextual features. However, existing methods are limited by the “prediction-before-allocation” design paradigm, which models set-level externalities within candidate ads and fails to consider the context of the final allocation, leading to suboptimal results. In this work, we introduce Contextual Generative Auction (CGA), a novel framework that incorporates permutation-level externalities in multi-slot ad auctions. Built on the structure of our theoretically derived optimal auction, CGA decouples the optimization of allocation and payment. We construct an autoregressive generative model for allocation, and reformulate incentive compatibility (IC) constraint into minimizing ex-post regret that supports gradient computation, enabling end-to-end learning of the optimal payment rule. Extensive offline and online experiments demonstrate that

CGA significantly enhances platform revenue and CTR compared to existing methods, and effectively approximates the optimal auction with nearly maximal revenue and minimal regret.

CCS Concepts

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

Keywords

Learning-based Auction Design; Externalities; Generative Auction

ACM Reference Format:

Ruitao Zhu, Yangsu Liu, Dagui Chen, Zhenjia Ma, Chufeng Shi, Zhenzhe Zheng, Jie Zhang, Jian Xu, Bo Zheng, and Fan Wu. 2025. Contextual Generative Auction with Permutation-level Externalities for Online Advertising. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '25)*, August 3–7, 2025, Toronto, ON, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3690624.3709313>

1 Introduction

Online advertising serves as a cost-efficient and precise channel for advertisers to promote contents to millions of online users, which has become the main revenue source for internet industry [12, 19]. Upon receiving a user request, online platforms conduct real-time auctions to determine ad allocation across multiple slots on a webpage, and compute the payment for each advertiser obtaining an ad slot. The optimal ad auction aims to maximize platform revenue while satisfying economic properties, such as Incentive Compatibility (IC) and Individual Rationality (IR)¹, along with computational complexity constraints for online deployment.

¹Please refer to the definitions of IC and IR in Section 2.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '25, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1245-6/25/08
<https://doi.org/10.1145/3690624.3709313>

Classical ad auctions jointly consider advertisers' bids and ads' click-through rate (CTR). The Generalized Second Price (GSP) auction has long been the benchmark for ad auctions due to its interpretability and ease of implementation. However, the separated CTR assumption in GSP, which posits that the users' clicking depends solely on the ad itself but not the other displayed items, restricts the model's performance. In reality, the page presented to users contains multiple items, and users compare factors such as prices and appearances before making decisions; therefore, other displayed items will influence the target ad's CTR [5, 43]. The effect of external items is called as *externalities* [15] in online advertising. The empirical study on user behaviors [20] also indicates that optimal ad auction design must take externalities into account.

Recently, learning-based auctions such as Deep Neural Auction (DNA) [25] and Score Weighted VCG (SW-VCG) [22] are proposed to better capture externalities and enhance platform revenue. However, these methods are limited by the "prediction-before-allocation" design paradigm, as the prediction process remains agnostic to the context within the allocation results. In essence, these methods only model information within the candidate ad set (*set-level externalities*), failing to incorporate the sequential context that affects the CTR of each ad within the final allocation (*permutation-level externalities*), resulting in a suboptimal allocation. Studies on reranking [29, 32, 39] in recommendation systems similarly consider item correlations among the displayed items. Nonetheless, their focus on enhancing overall user feedback neglects advertiser strategic behaviors, and thus fails to incorporate economic constraints like IC and IR, hindering their applicability to ad auction design.

Designing optimal auction with *permutation-level externalities* for online ad platform faces three critical challenges: (i) The existing "prediction-before-allocation" design paradigm cannot perceive the context in the sequential allocation while predicting ad value. (ii) Employing a VCG-like approach to evaluate all permutations can achieve the optimal allocation, but the high computational complexity makes it infeasible for online deployment. (iii) The intuition behind IC condition is to require that each advertiser's expected value, which is proportional to CTR in the Cost-per-Click setting, is non-decreasing with her bid [26]. Considering that higher slots typically have greater exposure, most existing methods aim to ensure that a higher bid secures the same or a higher slot for the advertiser. However, permutation-level externalities can lead to higher slots yielding lower CTR, leading to the violation of the IC condition. As illustrated in Figure 1, industrial data from *Taobao* advertising platform shows that the permutation-aware CTR of ads is not monotonic with respect to their allocated slot. Consequently, designing an ad auction with an IC constraint under permutation-level externalities is a non-trivial problem.

To address the above challenges, in this work, we propose Contextual Generative Auction (CGA), to optimize platform revenue with the guaranteed economic properties. The framework of CGA adheres to the structure of our theoretically derived optimal Incentive Compatible (IC) auction, which decouples the optimization of allocation and payment. CGA leverages the Generator-Evaluator (G-E) paradigm [13] to capture permutation-level externalities. Specifically, Generator utilizes an autoregressive model to generate an ad sequence as allocation, guided by Evaluator that refines contextual




Figure 1: The Permutation-aware ad CTR as a function of slot on Taobao advertising platform is non-monotonic. Each line represents an ad within a request. When advancing an ad's assigned slot with a fixed set of competing ads, its permutation-aware CTR is not strictly increasing.

interaction within the ad sequence. Additionally, directly calculating the optimal payment rule is infeasible, since the allocation rule of CGA is implemented via a generative model. Inspired by works on multi-item auction design that uses neural networks to effectively recover analytical solutions [11], we restate the IC constraint to require zero expected ex-post regret for ad auction, and learn the optimal payment by pushing gradients through the regret term. Our main contributions are summarized as follows:

- We formulate the multi-slot auction with permutation-level externalities and theoretically derive the optimal IC auction, providing theoretical guarantees for generative auction design.
- To model the permutation-level externalities, we break the "prediction-before-allocation" design paradigm of learning-based auctions by introducing an autoregressive generative model that directly generates allocations. Building on the structure of the optimal mechanism, the proposed CGA leverages the G-E paradigm to optimize allocation and minimizes differentiable ex-post regret to learn the optimal payment rule.
- Experiments on large-scale industrial data and online A/B tests demonstrate that CGA outperforms the existing methods in revenue and CTR, and effectively approximates the optimal auction with nearly maximal revenue and negligible regret.

2 PRELIMINARIES

2.1 Ad Auctions with Permutation-level Externalities

Multi-slot ad auction. The auction for online advertising can be abstracted as a typical multi-slot auction design problem. Formally, when a user request arrives, there are n advertisers² bidding for k ($k \leq n$) ad slots with each advertiser i submits a bid b_i based on her private value v_i for a click to ad_i ³. Advertiser i 's private value is drawn independently from a distribution $f_i(\cdot)$, and $F_i(\cdot)$ denotes the cumulative distribution function (CDF) of the probability density

²Advertisers can either manually adjust bids or employ the platform's auto-bidding agent [3, 8, 41] for automated bidding.

³We consider the single-parameter and Cost-per-Click (CPC) auction setting, i.e., each advertiser submits a bid and gets paid for a click event, which aligns with the practical scenarios of online ad platforms and remains consistent with related work [22, 23, 25].

function (PDF) $f_i(\cdot)$. Let $\mathbf{v} = (v_1, v_2, \dots, v_n)$ denote the value profile for all advertisers, with $\mathbf{v}_{-i}$ representing the value profile excluding the element v_i , similarly for $\mathbf{b}$ and $\mathbf{b}_{-i}$. The auctioneer, *i.e.*, the ad platform, knows the distributions $\mathbf{f} = (f_1, f_2, \dots, f_n)$ (derived from historical data), but does not know the true value profile $\mathbf{v}$.

Given bid profile $\mathbf{b}$, user feature vector $\mathbf{u}$, and the collection of all ad feature vectors $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, the ad auction mechanism is then formalized as $\mathcal{M}(\mathcal{A}(\mathbf{b}; \mathbf{X}, \mathbf{u}), \mathcal{P}(\mathbf{b}; \mathbf{X}, \mathbf{u}))$. The allocation rule $\mathcal{A}$ decides the ad-slot allocation, represented by the allocation matrix $\mathbf{A} = (a_{ij})_{n \times k} \in \{0, 1\}^{n \times k}$, where $a_{ij} = 1$ means that ad i is allocated to slot j and $a_{ij} = 0$ otherwise⁴. Note that an ad is allocated to at most one slot, and each slot must be assigned one ad, hence a feasible allocation matrix $\mathbf{A}$ satisfies $\sum_{j=1}^k a_{ij} \leq 1, \forall i$ and $\sum_{i=1}^n a_{ij} = 1, \forall j$; The payment rule $\mathcal{P}$ decides the payment of each ad, represented by the payment vector $\mathbf{p} = (p_i)_n \in \mathbb{R}^n$.

Permutation-level externalities. The externalities in ad auctions mean that the winning ads' utility is also influenced by other winning ads [15]. For multi-slot ad auctions, the externality effect is reflected in the CTR of ads. Specifically, CTR models can be abstracted as functions mapping from ad features and user features to the probability of a user clicking an ad. We further consider a permutation-aware CTR model Θ , which takes allocation $\mathbf{A}$ and associated features as input, and captures the inter-ad influence within $\mathbf{A}$. Note that model Θ is permutation-aware, meaning that even when two allocations contain the same set of ads, such as $\mathbf{A}_1 = (ad_1, ad_2, ad_3)$ and $\mathbf{A}_2 = (ad_3, ad_2, ad_1)$, the differing permutations cause a different external impact on the CTR of ad_2 . Formally, we denote the CTR of ad_i as $\Theta(\mathbf{x}_i; \mathbf{A}, \mathbf{X}, \mathbf{u})$. For ease of notation, let $\Theta_i(\mathbf{b}; \mathbf{X}, \mathbf{u}) = \Theta(\mathbf{x}_i; \mathcal{A}(\mathbf{b}; \mathbf{X}, \mathbf{u}), \mathbf{X}, \mathbf{u})$. The permutation-level externalities is encoded in the mapping process from allocation $\mathbf{A}$ to CTR via the model Θ .

Problem formulation. Given the mechanism $\mathcal{M}(\mathcal{A}, \mathcal{P})$, the expected utility of an advertiser with valuation v_i is given by:

$$u_i(v_i; \mathbf{b}; \mathbf{X}, \mathbf{u}) = (v_i - p_i(\mathbf{b}; \mathbf{X}, \mathbf{u})) \cdot \Theta_i(\mathbf{b}; \mathbf{X}, \mathbf{u}).$$

Considering that advertisers can employ strategies aimed at maximizing their utilities through misreporting their values, we introduce two essential properties of ad auction for the stability of the advertising platform: *Incentive Compatibility* (IC) and *Individual Rationality* (IR).

DEFINITION 1. An auction is IC, if each advertiser's utility is maximized by reporting truthfully no matter what the other advertisers report. Formally,

$$u_i(v_i; v_i, \mathbf{b}_{-i}; \mathbf{X}, \mathbf{u}) \geq u_i(v_i; b_i, \mathbf{b}_{-i}; \mathbf{X}, \mathbf{u}), \forall i \in [n], \forall v_i, b_i \in \mathbb{R}^+.$$

DEFINITION 2. An auction is IR, if every advertiser's payment does not exceed her submitted bid. Formally,

$$p_i(\mathbf{b}; \mathbf{X}, \mathbf{u}) \leq b_i, \forall i \in [n].$$

The goal is to design an auction $\mathcal{M}$ that maximizes the expected revenue of ad platform:

$$\mathbb{E}_{\mathbf{v} \sim F} \text{Rev}^{\mathcal{M}}(\mathbf{b}, \mathbf{X}, \mathbf{u}) = \sum_{i=1}^n p_i(\mathbf{b}; \mathbf{X}, \mathbf{u}) \Theta_i(\mathbf{b}; \mathbf{X}, \mathbf{u}),$$

⁴For notational simplicity, we slightly abuse matrix notation and rewrite $\mathbf{A} = (a_{A_1}, a_{A_2}, \dots, a_{A_k})$ to denote the allocation of ads to k slots, where $A_i \in [n]$ denotes the index of ad allocated to the i -th slot of allocation $\mathbf{A}$.

and satisfies IC and IR constraints, which can be formulated as:

$$\max_{\mathcal{M}(\mathcal{A}, \mathcal{P})} \mathbb{E}_{\mathbf{v} \sim F} \text{Rev}^{\mathcal{M}}(\mathbf{b}, \mathbf{X}, \mathbf{u}), \text{ s.t. IC and IR constraints.} \quad (1)$$

2.2 Optimal Multi-slot Auction

To solve the above revenue-maximizing auction design problem with IC and IR constraints, an intuitive method is to follow the well-known Myerson auction [26]. Myerson auction can be adapted to the setting with permutation-level externalities⁵ as follows:

DEFINITION 3. (Myerson auction with permutation-level externalities):

- **Allocation:** $\mathcal{A} \in \arg\max_{\mathcal{A}} \sum_{i=A_1}^{A_k} \tilde{\phi}(b_i, F_i) \Theta(\mathbf{x}_i; \mathbf{A}, \mathbf{X}, \mathbf{u})$;
- **Payment:** $p_i = \begin{cases} b_i - \frac{\int_0^{b_i} \Theta_i(t, \mathbf{b}_{-i}; \mathbf{X}, \mathbf{u}) dt}{\Theta_i(\mathbf{b}; \mathbf{X}, \mathbf{u})} & \text{if } \Theta_i(\mathbf{b}; \mathbf{X}, \mathbf{u}) > 0; \\ 0, & \text{otherwise,} \end{cases}$

where $\tilde{\phi}(v, F_i)$ denotes the ironed virtual value function [26], which is monotone in v .

Recall Myerson's Lemma [26], as formulated in Theorem 1.

THEOREM 1. (Myerson's Lemma [26]). For a single-parameter environment, an allocation rule $\mathcal{A}$ is implementable if there exists a payment rule $\mathcal{P}$ such the mechanism $\mathcal{M}(\mathcal{A}, \mathcal{P})$ is IC. The following two claims hold: (1) An allocation rule is implementable if and only if it is monotone. (2) If allocation rule $\mathcal{A}$ is monotone, then there exists a unique payment rule $\mathcal{P}$ such that the mechanism $\mathcal{M}(\mathcal{A}, \mathcal{P})$ is IC. It is given by:

$$p_i(b_i, \mathbf{b}_{-i}) = b_i \mathcal{A}_i(b_i, \mathbf{b}_{-i}) - \int_0^{b_i} \mathcal{A}_i(t, \mathbf{b}_{-i}) dt.$$

Note that introducing permutation-level externalities diverges from traditional Myerson auction by influencing each ad's CTR based on allocation outcomes. As such, an increase in ad_i 's bid does not necessarily increase its CTR due to inter-ad influences within the allocation results. This could potentially violate the monotonic allocation requirement of Theorem 1. We address this by proving that, even with permutation-level externalities, the monotonic allocation condition in Myerson's Lemma holds if the allocation rule maximizes virtual welfare, as demonstrated in Lemma 1. We put the proof in Appendix A.1 due to space limitations.

LEMMA 1. (Monotonic allocation with permutation-level externalities). For every ad_i and $\mathbf{b}_{-i}$, the obtained CTR $\Theta(\mathbf{x}_i; \mathcal{A}(b_i, \mathbf{b}_{-i}), \mathbf{X}, \mathbf{u})$ is non-decreasing in its bid b_i (or we say $\mathcal{A}$ is monotone), if the allocation rule $\mathcal{A}$ maximizes the expected virtual welfare.

THEOREM 2. For a single-parameter environment, maximizing expected revenue over the space of IC auctions is equal to maximizing expected virtual welfare [26].

Theorem 2 [26] equates the goal of maximizing revenue to maximizing expected virtual welfare. Based on Theorem 1 and 2 and Lemma 1, we can deduce the corollary that Myerson auction with

⁵Li et al. [22] directly rewrites Myerson Auction into an externality-aware form, yielding similar conclusions. However, when incorporating permutation-level externalities, unlike traditional single-item auctions, the ad CTR is influenced by the allocation outcome. This impact complicates the transition of the monotone allocation condition, which satisfies IC constraint, from Myerson Auction. Consequently, the preservation of optimality is non-trivial. We discuss this in detail in Lemma 1 and Appendix A.1.

permutation-level externalities in Definition 3 constitutes the optimal solution to Problem (1).

COROLLARY 1. *The ad auction $\mathcal{M}$, characterized by allocation rule $\mathcal{A}$ and payment rule $\mathcal{P}$ in Definition 3, represents the optimal mechanism with permutation-level externalities that maximizes the platform's expected revenue while satisfying IC and IR constraints.*

2.3 Auction Design as a Learning-based Problem

A direct approach to implement the allocation rule $\mathcal{A}$, as outlined in Definition 3, involves enumerating all permutations to select the one maximizing virtual welfare, with a computational complexity of $P(n, k) = \frac{n!}{(n-k)!}$. However, for online advertising, taking *Taobao* as an example, where $n \approx 50$ and $k \approx 5$, the high computational complexity fails to meet online deployment requirements. Therefore, we parameterize the auction mechanism as $\mathcal{M}^w = \langle \mathcal{A}^{w_a}, \mathcal{P}^{w_p} \rangle$ with parameters w_a and w_p , and solves a learning problem to determine these parameters.

To impose IC constraint on learning-based auctions and ensure the differentiability of the optimization process, similar to the original work of learning-based multi-item auction [11], we introduce the concept of *ex-post regret* for each advertiser to quantify the extent of deviation from IC conditions. Specifically, the *ex-post regret* for ad_i is defined as the maximum increase in utility that can be obtained through misreporting b'_i :

$$rgt_i(v_i, X, \mathbf{u}) = \max_{b'_i} \{u_i(v_i; b'_i, \mathbf{b}_{-i}; X, \mathbf{u}) - u_i(v_i; \mathbf{b}; X, \mathbf{u})\}.$$

Hence, the IC constraint is satisfied if and only if the ex-post regret for each ad_i is zero. Given L value samples from the distribution $\mathbf{F}$, we obtain the empirical ex-post regret for ad_i :

$$\widehat{rgt}_i = \frac{1}{L} \sum_{l=1}^L rgt_i(v_i^l, X, \mathbf{u}). \quad (2)$$

Then we reformulate the auction design problem (1) as minimizing the expected negated revenue, subject to the constraint that the empirical ex-post regret for each ad_i is zero and each advertiser's utility is non-negative:

$$\min_w -\mathbb{E}_{\mathbf{v} \sim \mathbf{F}} Rev^{\mathcal{M}^w}, \text{ s.t. } \widehat{rgt}_i = 0, u_i \geq 0, \forall i \in [n]. \quad (3)$$

3 Methodology

In this section, we provide a detailed discussion of the Contextual Generative Auction (CGA). To overcome the limitation of the “prediction-before-allocation” paradigm in capturing permutation-level externalities, CGA employs a G-E architecture, as depicted in Figure 2. The Generator leverages an autoregressive module, perceiving the established preceding context to generate an ad sequence. The Evaluator estimates permutation-aware reward within the whole ad sequence, guiding Generator towards the optimal allocation through policy gradient. For online inference, only Generator is deployed, ensuring a minimal computational latency during ad allocation. Moreover, a dedicated PaymentNet module is introduced to learn the optimal payment rule, trained via optimizing differentiable ex-post regret.

3.1 Generator: Autoregressive Generative Module

According to Theorem 2 and Corollary 1, the objective of Generator is to generate an ad sequence $\mathbf{A}$ of length k from n candidate ads that maximizes the expected virtual welfare. We develop an autoregressive generative module, comprising two key components: a permutation-invariant set-level encoder and a permutation-equivariant decoder.

- *Permutation invariance* [42], an architectural property wherein the outputs remain unaffected by the permutation of inputs, has been shown to enhance the platform revenue of ad auctions [25]. The encoder aims to enhance ad representation by modeling set-level externalities, requiring this property because its input is the candidate ad set.
- *Permutation equivariance* [42] is an architectural property wherein changing the permutation of inputs will correspondingly change the permutation of outputs. In each step, the decoder takes candidate ads as input and outputs the allocation probability for each ad. Therefore, altering the permutation of input ads should result in a corresponding change in the permutation of output logits. This property is widely adopted in automated auction design [10, 18, 31], with Qin *et al.* [30] demonstrating its efficacy in improving the generalization ability of learning-based auctions.

Permutation-invariant set-level encoder. It encodes the candidate ad set and provides the context embedding $\mathbf{c}$ as the initial input of decoder. First, we adopt a self-attention layer to capture the mutual influence among n candidate ads, producing the set-level augmented embedding $\mathbf{h}_i$ for each ad i :

$$[\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n] = \text{self-attention}([\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n]), \quad (4)$$

where $\mathbf{e}_i$ denotes the embedding of ad a_i . Since the encoder processes an unordered set of ads, the above attention layer excludes positional encoding, and we perform sum pooling on $\mathbf{h}_i, i \in [n]$, ensuring the property of *permutation invariance* that is changing the permutation of input ads retains the same output context embedding $\mathbf{c}$:

$$\mathbf{c} = \text{MLP}(\sum_{i=1}^n \mathbf{h}_i; \mathbf{h}_u),$$

where $\mathbf{h}_u$ denotes the user representation.

Permutation-equivariant autoregressive decoder. To address the limitations of the “prediction-before-allocation” paradigm in auction design, we introduce an autoregressive decoder. This model efficiently captures joint distributions over the output ad sequences [4] to generate the ad sequence. Given the context embedding $\mathbf{c}$, the probabilistic generative model learns the conditional probability $p(\mathbf{A}_m | \mathbf{c})$ for each ad sequence $\mathbf{A}_m$. During inference, the model selects a specific allocation from its output space based on this conditional probability. In the multi-slot auction setting, where $\mathbf{A}_m$ consists of k ads: $a_{A_{m_1}}, a_{A_{m_2}}, \dots, a_{A_{m_k}}$, and $a_{A_{m_i}}$ are not independent, the autoregressive decoder models the joint output distribution using the product rule:

$$p(a_{A_{m_1}}, a_{A_{m_2}}, \dots, a_{A_{m_k}} | \mathbf{c}) = p(a_{A_{m_1}} | \mathbf{c}) p(a_{A_{m_2}} | \mathbf{c}, a_{A_{m_1}}) \dots p(a_{A_{m_k}} | \mathbf{c}, a_{A_{m_1}}, a_{A_{m_2}}, \dots, a_{A_{m_{k-1}}}).$$




Figure 2: The architecture of Contextual Generative Auction (CGA). The middle part shows the overall framework of CGA, with dashed outlines and arrows depicting offline training components, and solid lines representing online inference paths. The other two parts provide specific implementations of Generator and Evaluator.

We further employ the GRU cell [9] to model the conditional probability $p(a_i | c, a_{A_{m_1}}, \dots, a_{A_{m_{t-1}}})$ for each candidate ad a_i at slot $t \in [k]$ ($a_{A_{m_0}} = \emptyset$). At the beginning of decoding, the context embedding c initializes the hidden state of GRU, and a special token representing “START” is fed into GRU as the initial input. Iteratively, the t -th GRU cell is formulated as:

$$s_t = \text{GRU}(s_{t-1}, \mathbf{h}_{A_{t-1}}), \quad t = 1, 2, \dots, k,$$

where $\mathbf{h}_{A_{t-1}}$ denotes the encoded representation of the ad chosen in slot $t-1$ ($\mathbf{h}_{A_0} = \text{START}$), and s_{t-1} denotes the preceding contextual information ($s_0 = c$). Consequently, given the state s_t at slot t , we obtain the allocation probability z_i^t of each candidate ad a_i :

$$z_i^t = \text{softmax}([\text{MLP}([s_t; \mathbf{h}_i]) + e^w \tilde{\phi}_i]_{i=1}^n), \quad (5)$$

where w is a learnable parameter such that e^w remains positive, ensuring higher virtual value leads to a greater allocation probability, aligning with the objective of maximizing virtual welfare. Ads allocated prior to slot t are masked, and a sampling strategy based on z^t determines the ad allocated to slot t . This sampling occurs only during training to explore diverse sequence generation strategies. During inference, the ad with the highest value in z^t is selected. This selection process is repeated k times to generate an ad sequence of length k , denoted as $\mathbf{A} = [a_{A_1}, a_{A_2}, \dots, a_{A_k}]$.

Moreover, since the MLPs and GRU cells in Generator operate on each state-ad pair, and the encoder satisfies permutation invariance, the decoder exhibits the property of *permutation equivariance*.

3.2 Evaluator: Permutation-aware Prediction Module

The goal of Evaluator is to predict the permutation-aware CTR $\Theta(x_i; \mathbf{A}, \mathbf{X}, \mathbf{u})$ for each ad a_{A_i} in the ad allocation $\mathbf{A}$. Evaluator takes the allocation embedding $\mathbf{H}_\mathbf{A} = [\mathbf{h}_{A_1}, \mathbf{h}_{A_2}, \dots, \mathbf{h}_{A_k}]$, user

embedding $\mathbf{e}_\mathbf{u}$ and point-wise CTR $\alpha_\mathbf{A} \in [0, 1]^k$ as input, and outputs the permutation-aware CTR. Each embedding $\mathbf{h}_{A_i}$ of ad a_{A_i} is derived from the encoder of Generator, as defined in Equation (4). Since CGA operates at the end of the three-stage advertising system (i.e., matching, prediction and auction), to fully utilize the user interests captured in the preceding stages, Evaluator leverages the point-wise CTR $\alpha_\mathbf{A}$ output from the prediction stage, and constructs a Permutation-level Ad Encoder to estimate a personalized externality-aware calibration vector $\gamma_\mathbf{A} \in (0, 2)^k$. These two vectors are then element-wise multiplied to obtain the permutation-aware CTR: $\Theta_\mathbf{A} = \min(\alpha_\mathbf{A} \odot \gamma_\mathbf{A}, 1)$.

Specifically, we adopt a Bi-LSTM [16] layer and a multi-head self-attention [36] layer to encode the ad sequence embedding $\mathbf{H}_\mathbf{A}$, where Bi-LSTM effectively captures bidirectional sequence information, while self-attention efficiently captures interactions between distantly positioned ads within the sequence. Formally, the sequential representation of self-attention layer is defined as:

$$\mathbf{H}_\mathbf{A}^s = \text{softmax}\left(\frac{\mathbf{Q}_\mathbf{A} \mathbf{K}_\mathbf{A}^T}{\sqrt{d}}\right) \mathbf{V}_\mathbf{A},$$

where d denotes the dimension of embeddings, and $\mathbf{Q}_\mathbf{A}, \mathbf{K}_\mathbf{A}, \mathbf{V}_\mathbf{A}$ represents the query, key and value matrices, respectively, which are transformed linearly from the sum of allocation embedding $\mathbf{H}_\mathbf{A}$ and positional encoding $\mathbf{PE}_\mathbf{A}$ (the sinusoidal version in [36]).

Next, we obtain the forward output state $\mathbf{H}_\mathbf{A}^f$ and backforward output state $\mathbf{H}_\mathbf{A}^b$ of $\mathbf{H}_\mathbf{A}$ through Bi-LSTM layer.

$$\mathbf{H}_\mathbf{A}^f = \text{Bi-LSTM}_{\text{forward}}(\mathbf{H}_\mathbf{A})$$

$$\mathbf{H}_\mathbf{A}^b = \text{Bi-LSTM}_{\text{backward}}(\mathbf{H}_\mathbf{A}).$$

For brevity, we omit Bi-LSTM’s implementation details, including the input gates, forget gates, output gates and cell states.

Subsequently, all sequential representations are concatenated with the user representation $\mathbf{h}_u$, and fed into a Multilayer Perceptron (MLP) to obtain the externality-aware calibration vector:

$$\gamma_A = 2\sigma(r(r([H_A^s; H_A^f; H_A^b; \mathbf{h}_u])),$$

where $\sigma(\cdot)$ and $r(\cdot)$ denote the fully connected layer with sigmoidal and ReLU activation functions, respectively.

Finally, it should be noted that CGA allows any Evaluator model that captures sequential context to improve the G-E framework. In this work, we present an efficient implementation from our practice.

3.3 Payment Computation

According to Corollary 1, the optimal payment is defined in Definition 3, where the integral term can be rewritten as an expectation:

$$\int_0^{b_i} \Theta_i(t, \mathbf{b}_{-i}; \mathbf{X}, \mathbf{u}) dt = b_i E_{t_i \sim U[0, b_i]} [\Theta_i(t_i, \mathbf{b}_{-i}; \mathbf{X}, \mathbf{u})],$$

which can be approximated using Monte Carlo sampling. However, we note that $\Theta_i(t_i, \mathbf{b}_{-i}; \mathbf{X}, \mathbf{u}) = \Theta(\mathbf{x}_i; \mathcal{A}(t_i, \mathbf{b}_{-i}; \mathbf{X}, \mathbf{u}), \mathbf{X}, \mathbf{u})$. Processing each sample requires invoking Generator $\mathcal{A}$ and Evaluator Θ , resulting in high computational costs during inference. Reducing the number of samples can mitigate this but increases payment variance, directly affecting platform revenue and advertiser utility.

Motivated by the application of neural networks to effectively approximate the optimal mechanism in multi-item auctions in [11], we introduce *PaymentNet* to learn the optimal payment rule. Specifically, the inputs include the allocation embedding $\mathbf{H}_A \in \mathbb{R}^{k \times d}$, the self-exclusion bidding profile $\mathbf{B}^- = [\mathbf{b}_{-A_1}, \mathbf{b}_{-A_2}, \dots, \mathbf{b}_{-A_k}] \in \mathbb{R}^{k \times (k-1)}$, and the expected value profile $\mathbf{Z} \cdot \Theta \in \mathbb{R}^{k \times 1}$, where $\mathbf{Z} = [z_{A_1}^1, z_{A_2}^2, \dots, z_{A_k}^k]$ denotes the allocation probability defined in Equation (5) and $\Theta = [\Theta_{A_1}, \Theta_{A_2}, \dots, \Theta_{A_k}]$ denotes the permutation-aware CTR estimated by Evaluator. Moreover, to satisfy IR constraint, as defined in Definition 2, PaymentNet employs a sigmoidal activation function to compute the payment rate $\tilde{\mathbf{p}} \in [0, 1]^k$, and subsequently outputs the payment $\mathbf{p} = \tilde{\mathbf{p}} \odot \mathbf{b}$. Formally, the payment rate is defined as:

$$\tilde{\mathbf{p}} = \sigma(r(r([\mathbf{H}_A; \mathbf{B}^-; \mathbf{Z} \cdot \Theta])),$$

where $\sigma(\cdot)$ and $r(\cdot)$ denote the fully connected layer with sigmoidal and ReLU activation functions, respectively.

4 Optimization and Training

According to Corollary 1, the optimal allocation rule only requires maximizing virtual welfare and is independent of the payment rule. Therefore, we decouple the optimization of G-E framework and PaymentNet, which maintains the optimality of CGA.

4.1 Optimization of G-E Framework

In the G-E framework, as Evaluator captures permutation-level externalities and guides Generator to obtain the optimal allocation, we first train Evaluator Θ to convergence using list-wise ad click data. Each sample $l \in \mathcal{D}$ is an ad sequence of length k exposed to a user, with the label $y_i^l \in \{0, 1\}$ indicating whether the user clicks

on the i -th ad. The binary cross-entropy loss is defined as:

$$\mathcal{L}_E = -\frac{1}{|\mathcal{D}|} \sum_{l \in \mathcal{D}} \sum_{i=1}^k (y_i^l \log \theta_i^l + (1 - y_i^l) \log(1 - \theta_i^l)), \quad (6)$$

where θ_i^l denotes the permutation-aware CTR of the i -th ad in the ad sequence l as predicted by Evaluator.

After training Evaluator Θ to convergence, we freeze its parameters and train Generator using policy gradient with rewards from Evaluator. At each slot t , the contextual information $\mathbf{s}_t$ serves as the state, and the allocation probability z_i^t output by Generator serves as the action. Since $\mathbf{s}_t$ only encodes the preceding context, to capture the bi-directional contextual information, we use the well-trained Evaluator to estimate the permutation-aware reward of each candidate ad within the whole ad sequence $\mathbf{A}$. Motivated by the effective use of the policy gradient method in the reranking task [13], we decompose this permutation-aware reward into two parts.

Self-Reward. Aligning with the objective of optimal allocation, we use Evaluator Θ to estimate the expected virtual welfare of each selected ad a_{A_i} , termed as self-reward:

$$r_{A_i}^{self} = \tilde{\phi}_i \cdot \Theta_i(\mathbf{b}; \mathbf{A}, \mathbf{u}). \quad (7)$$

External-Reward. Each selected ad not only contributes its reward but also affects the CTR of other ads due to the permutation-level externalities. Similar to the marginal contribution applied in classical VCG mechanism [35], we model this external effect as external reward:

$$r_{A_i}^{external} = \sum_{j \in \mathbf{A}_{-i}} \tilde{\phi}_j \Theta_j(\mathbf{b}; \mathbf{A}, \mathbf{u}) - \sum_{j \in \mathbf{A}_{-i}} \tilde{\phi}_j \Theta_j(\mathbf{b}_{-i}; \mathbf{A}_{-i}, \mathbf{u}), \quad (8)$$

where $\mathbf{A}_{-i}$ denotes the ad sequence excluding a_{A_i} .

Combining the above two rewards, we obtain the permutation-aware reward of selecting a_{A_i} , defined as:

$$\begin{aligned} r_{A_i} &= r_{A_i}^{self} + r_{A_i}^{external} \\ &= \tilde{\phi}_i \Theta_i(\mathbf{b}; \mathbf{A}, \mathbf{u}) + \sum_{j \in \mathbf{A}_{-i}} \tilde{\phi}_j \Theta_j(\mathbf{b}; \mathbf{A}, \mathbf{u}) - \sum_{j \in \mathbf{A}_{-i}} \tilde{\phi}_j \Theta_j(\mathbf{b}_{-i}; \mathbf{A}_{-i}, \mathbf{u}) \\ &= \sum_{j \in \mathbf{A}} \tilde{\phi}_j \Theta_j(\mathbf{b}; \mathbf{A}, \mathbf{u}) - \sum_{j \in \mathbf{A}_{-i}} \tilde{\phi}_j \Theta_j(\mathbf{b}_{-i}; \mathbf{A}_{-i}, \mathbf{u}). \end{aligned}$$

Finally, the loss function of Generator is defined as:

$$\mathcal{L}_G = -\frac{1}{|\mathcal{D}|} \sum_{s \in \mathcal{D}} \sum_{i \in k} r_{A_i^s} \log z_{A_i^s}, \quad (9)$$

where s denotes a sample from dataset $\mathcal{D}$, representing the candidate ad set for a request, A_i^s denotes the i -th ad of the ad allocation output by Generator based on the input ad set s , and $z_{A_i^s}$ denotes the allocation probability in Equation (5).

4.2 Optimization of PaymentNet

After training Generator $\mathcal{A}$ and Evaluator Θ to convergence, we freeze their parameters and then apply the augmented Lagrangian method to solve the constrained optimization problem (3) to optimize PaymentNet. The Lagrangian function augmented with a

quadratic penalty term for violating the IC constraint is defined as:

$$\mathcal{L}_P = -\frac{1}{|D|} \sum_{s \in \mathcal{D}} \left(\sum_{i \in k} \mathcal{P}_i(A^s) \Theta_i(A^s) - \sum_{i \in k} \lambda_i \widehat{rgt}_i^s - \frac{\rho}{2} \sum_{i \in k} (\widehat{rgt}_i^s)^2 \right), \quad (10)$$

where A^s denotes the allocation output by the freezed Generator $\mathcal{A}$ with the input of ad set s , λ_i denotes Lagrange multipliers, and $\rho > 0$ denotes the hyperparameter for the IC penalty term.

The optimization process based on the Lagrangian function (10) includes the iteration of (i) updating PaymentNet: $w_p^{new} = \argmin_{w_p} \mathcal{L}_P(w_p^{old}; \lambda^{old})$, and (ii) updating the Lagrange multipliers: $\lambda^{new} = \lambda^{old} + \rho \widehat{rgt}(w_p^{new})$. Note that problem (3) is non-convex. Thus, the above Lagrangian method is not guaranteed to converge to the global optimum. However, the experimental results indicate that the optimized CGA can approximate the optimal revenue with minimal regret, as discussed in Section 5.2.

5 Experiments

In this section, we conduct offline experiments and online A/B tests to evaluate the effectiveness of CGA, and mainly focus on the following questions:

- **RQ1:** How does CGA perform in terms of key metrics like platform revenue and CTR, compared with existing ad auctions?
- **RQ2:** What are the effects of modeling permutation-level externalities compared to set-level and W/O externalities?
- **RQ3:** How do various designs of CGA affect its performance?
- **RQ4:** How does our proposed CGA perform in the real-world ad auction scenarios with efficient deployment?

5.1 Experiment Setup

5.1.1 Dataset. In offline experiments, we evaluate the performance of CGA using real logs collected from a leading e-commerce platform *Taobao* during January 2024. The training set comprises 500,000 randomly selected auctions from January 20 to January 23, involving 1,100,875 unique advertisers while the test set comprises 100,000 randomly selected auctions from January 25, involving 452,671 unique advertisers. For each auction sample in the dataset, approximately 30 advertisers are selected by ad system as candidates for the auction stage, where they submit their bids to compete for slots. We set the number of slots $k = 3$ in offline experiments.

5.1.2 Baseline Methods. We compare CGA with representative auction mechanisms. The baseline methods are categorized into three groups based on the granularity of externality modeling.

W/O externalities:

- **GSP:** GSP ranks ads based on the product of bids and predicted CTR, without modeling externalities.

Set-level externalities:

- **DNA [25]:** Building on GSP, DNA predicts each ad's rankscore by modeling set-level externalities of candidate ads and ranks accordingly for allocation, without considering the mutual influence among ads in the ad list.

- **SW-VCG [22]:** SW-VCG formalizes multi-slot auction design as a maximum weighted bipartite matching problem between ads and slots, which captures only set-level externalities of candidate ads to estimate edge weights.

Permutation-level externalities:

- **VCG Auction [35] with Permutation-level Externalities (VCG):** VCG selects the ad permutation that maximizes social welfare for allocation and combines this with a payment rule to satisfy IC. We evaluate all permutations of the candidate ads with CGA's Evaluator, enabling VCG to capture permutation-level externalities.

- **EdgeNet [33]:** EdgeNet utilizes Transformer to replace DNA's set encoder and employs a greedy strategy to sequentially allocate ads, modeling partial permutation-level externalities.

- **Optimal:** To evaluate CGA's approximation to the theoretical upper bound, according to Corollary 1, we construct this baseline by traversing all permutations to maximize virtual welfare and using Monte Carlo sampling to calculate the payment.

5.1.3 Performance Metrics. We consider the following metrics to measure platform revenue, user experience, and ex-post regret of advertisers, respectively.

- Revenue Per Mille: $RPM = \frac{\sum click \times payment}{\sum impression} \times 1000$.

- Click-Through Rate: $CTR = \frac{\sum click}{\sum impression}$.

- IC metric: $\Psi = \frac{1}{|D|} \sum_{s \in \mathcal{D}} \sum_{i \in k} \frac{\widehat{rgt}_i^s}{u_i(v_i^s, b^s; X^s, u^s)}$, where $\widehat{rgt}_i^s$ is defined in Equation (2). This is a common metric for IC testing of ad auction mechanisms [7, 23, 25, 33], which measures the relative increase in utility that advertisers can obtain by manipulating their bids. Consistent with the evaluation process in [38], we conduct counterfactual perturbation on each advertiser's bid by replacing b_i with $\alpha \times b_i$, where $\alpha \in \{0.2 \times j \mid j = 1, 2, \dots, 10\}$.

5.1.4 Implementation. In CGA, we set the embedding size of features as 8. To capture richer information from different representation subspaces, a multi-head attention mechanism with 4 attention heads is used in all attention layers. The externality-aware calibration vector γ_A and the payment rate $\tilde{p}$ are derived using a MLP with hidden layers of sizes 128 and 32. CGA is trained with the Adam optimizer at a learning rate of 1e-3 and a batch size of 512. All learning-based baseline methods have hyperparameters tuned through grid search, exploring learning rates of {1e-5, 1e-4, 1e-3, 1e-2}, batch sizes of {256, 512, 1024, 2048}, and hidden sizes of {8, 16, 32, 64}.

5.2 Offline Comparison (RQ1 & RQ2)

CGA and baseline-Optimal involve the ironed virtual value function, which requires knowledge of each advertiser's value distribution $f_i(\cdot)$. Estimating $f_i(\cdot)$ from historical data is well-covered in existing research [6, 21, 27, 28]. This paper does not delve into this estimation as CGA focuses on employing generative models to capture permutation-level externalities, independent of distribution estimation techniques. To avoid biases from distribution estimation methods in comparative experiments, following [22], we preset the conditional distribution of each advertiser's value based on their feature vectors as *uniform* and *exponential distributions*, respectively, and regenerate the bids accordingly. For SW-VCG, since its predicted ad score Scr_i aims to fit the virtual value, we directly set Scr_i to $\tilde{\phi}_i$ for comparison.

The experimental results under the above two distributions are presented in Table 1. Our key observations are:

- (1) As externality modeling advances from W/O externalities to set-level and finally to permutation-level, auction performance

Table 1: Performance comparison for key metrics. Lift percentage indicates the relative change of baseline methods over CGA.

Value Distribution	Externalities	Model	RPM	CTR	Ψ
Uniform	W/O Externalities	GSP	1141.40 (-9.7%)	0.04359 (-7.1%)	14.7%
	Set-level	DNA	1195.75 (-5.4%)	0.04537 (-3.3%)	8.4%
		SW-VCG	1190.70 (-5.8%)	0.04528 (-3.5%)	5.6%
	Permutation-level	EdgeNet	1205.87 (-4.6%)	0.04561 (-2.8%)	4.9%
		VCG	1085.78 (-14.1%)	0.04791 (+2.1%)	0.0
		Optimal	1318.36 (+4.3%)	0.04772 (+1.7%)	0.0
		CGA	1264.01	0.04692	2.1%
Exponential	W/O Externalities	GSP	1179.22 (-12.2%)	0.04050 (-9.8%)	18.3%
	Set-level	DNA	1257.11 (-6.4%)	0.04283 (-4.6%)	10.1%
		SW-VCG	1247.71 (-7.1%)	0.04266 (-5.0%)	7.4%
	Permutation-level	EdgeNet	1266.52 (-5.7%)	0.04315 (-3.9%)	6.7%
		VCG	1321.85 (-15.8%)	0.04602 (+2.5%)	0.0
		Optimal	1415.60 (+5.4%)	0.04589 (+2.2%)	0.0
		CGA	1343.07	0.04490	3.7%

improves across three metrics (higher RPM and CTR with lower regret)⁶. This highlights the importance of modeling fine-grained externalities in auction design (RQ2). To further investigate the impact of externalities, we conduct comparative experiments on the CTR prediction task. The results are shown in Table 2, where "W/O Externalities" indicates the use of point-wise pCTR and "Set-level" refers to using the Set-level Ad Encoder, as described in Section 3.2, to replace Evaluator's permutation-level ad encoder. The results demonstrate that Evaluator with permutation-level externalities improves the accuracy of predictive values, explaining the improvement in auction metrics.

(2) Compared to baseline-Optimal based on enumeration, CGA achieves approximately 95% revenue maximization and negligible ex-post regret. This suggests that CGA enables efficient ad allocation and closely approximates the optimal auction mechanism.

Table 2: Performance comparison of externalities modeling.

	Logloss	AUC
W/O Externalities	0.2311	0.7366
Set-level Externalities	0.2286	0.7391
Permutation-level Evaluator Θ	0.2237	0.7476

5.3 Ablation Study (RQ3)

To verify the effectiveness of CGA's various design considerations, we construct the following variants:

- CGA- Θ removes the Evaluator Θ and uses point-wise CTR instead to evaluate each ad allocation's virtual welfare.
- CGA (end2end) directly trains Generator and PaymentNet using $\mathcal{L}_P$ from Equation (10). In contrast, CGA first trains Generator to convergence, freezes its parameters, and then trains PaymentNet.
- CGA- r^{self} removes the self reward r^{self} , defined in Equation (7), in the loss function $\mathcal{L}_G$ of Generator.

⁶VCG optimizes social welfare at the expense of revenue, independent of externalities.

- CGA- $r^{external}$ removes the external reward $r^{external}$, defined in Equation (8), in the loss function $\mathcal{L}_G$ of Generator.
- CGA- $\tilde{\phi}$ replaces virtual value $\tilde{\phi}_i$ with b_i in $\mathcal{L}_G$.

Table 3: Ablation study of CGA.

Model	RPM	CTR	Ψ
CGA	1264.01	0.04692	2.1%
CGA- Θ	1214.71 (-3.9%)	0.04579 (-2.4%)	3.7%
CGA (end2end)	1181.85 (-6.5%)	0.04504 (-4.0%)	10.5%
CGA- r^{self}	1122.44 (-11.2%)	0.04302 (-8.3%)	18.9%
CGA- $r^{external}$	1210.92 (-4.2%)	0.04575 (-2.5%)	3.4%
CGA- $\tilde{\phi}$	1231.15 (-2.6%)	0.04603 (-1.9%)	2.8%

The results are shown in Table 3, from which we observe:

(1) CGA- Θ performs worse than CGA, indicating that the Generator alone cannot fully capture permutation-level externalities, as the autoregressive model perceives only preceding context. The G-E framework helps distill complete sequential knowledge into Generator via policy gradient.

(2) CGA (end2end) performs worse. One likely reason is that CGA uses PaymentNet to learn the optimal payment, so the revenue gradient must pass through PaymentNet before reaching Generator while conducting end-to-end training, complicating convergence to the optimal allocation. Furthermore, Corollary 1 ensures that decoupled optimization maintains optimality, thus allowing Generator to focus solely on maximizing virtual welfare.

(3) CGA- r^{self} performs the worst due to the absence of each selected ad's own value increment. CGA outperforms CGA- $r^{external}$ because it uses external rewards to model the impact of each selected ad on the final allocation list.

(4) Estimating value distribution $f_i(\cdot)$ requires empirical assumptions about bidding strategies, since Ad platforms can only observe bids, causing bias in the estimated $\hat{f}_i(\cdot)$. To assess the impact of removing value distribution on CGA's performance, we construct the

variant $\text{CGA-}\tilde{\phi}$. As shown in Table 3, $\text{CGA-}\tilde{\phi}$ shows only a minor decline in revenue (2.6%) and CTR (1.9%) compared to CGA. This small drop may be due to CGA’s modeling of externalities, which captures inter-ad correlations and partially reflects the missing value distribution information through other advertisers’ bids.

5.4 Online A/B Test (RQ4)

To verify CGA’s effectiveness in the real world, we compare CGA with the fully deployed DNA in the Taobao advertising system through online A/B tests. Table 4 presents the experimental results of the 7-day online A/B testing, utilizing 2% of total production traffic. CGA achieves a 3.2% improvement in RPM with only a 3 ms average increase (1.6% relatively) in online response time (RT) per request, suggesting that CGA can efficiently explore the allocation space via generative models and boost platform revenue. Moreover, the heightened return on investment (ROI) for advertisers indicates that CGA’s revenue enhancement results not from inflated payments but from optimized ad allocation by capturing permutation-level externalities, as evidenced by significant CTR and gross merchandise volume (GMV) improvements.

Table 4: Experimental results from Online A/B tests.

Relative change in metrics	RPM	CTR	GMV	ROI	RT
CGA over baseline-DNA	+3.2%	+1.4%	+6.4%	+3.5%	+1.6%

6 Related Work

GSP [12] and its variants, like uGSP [2], are widely used in online advertising due to their interpretability and high revenue guarantee, but they do not consider the impact of other ads on user clicks, neglecting externalities [14, 17], leading to suboptimal performance.

Advances in computing have led researchers to explore learning-based auctions [44]. DeepGSP [45] and DNA [25] extend GSP using online feedback for end-to-end learning. However, DNA’s rankscore prediction faces the evaluation-before-ranking problem [40], restricting its scope to set-level externalities and yielding suboptimal allocations. SW-VCG [22] separates optimal auction design into designing monotone score functions and solving ad-slot maximum bipartite matching. However, SW-VCG still overlooks the mutual influence of ads in the final sequence, as the edge weights in the bipartite graph cannot be predetermined considering permutation-level externalities. Following VCG, NMA [23] proposed an enumeration-based framework to select the optimal allocation. While globally optimal, its high computational complexity makes it impractical for real-time online inference [34]. EdgeNet [33] employs a PointerNet-based structure for greedy ad allocation but ignores the impact of succeeding ads, remaining insufficient for achieving optimal results.

Studies of reranking in RS parallels externality modeling by leveraging contextual information to optimize item sequences. Reranking methods can be divided into one-stage and two-stage approaches [32]. One-stage methods [1, 29] estimate refined scores for each item within the initial list and rerank them greedily. These methods encounter the evaluation-before-ranking problem similar to DNA: reranking alters the permutation, introducing different mutual influences. Two-stage methods [13, 32, 34] typically employ a G-E

framework. The Generator produces multiple feasible sequences, while the Evaluator selects the optimal sequence based on the estimated list value. This approach enables comprehensive exploration of the permutation space [32], which provides valuable insights for the implementation of CGA. While the G-E framework effectively captures permutation-level externalities, it lacks the capacity to express key economic constraints such as IC and fails to directly optimize platform revenue. Our theoretical results decouple the allocation and payment in optimal auctions with permutation-level externalities. This allows the allocation to employ a general G-E framework focused solely on maximizing expected virtual welfare, while the optimal payment rule is learned through differentiable ex-post regret. Consequently, CGA can be viewed as a unified framework bridging reranking in RS and ad auction theory.

7 CONCLUSION

In this work, we have proposed the Contextual Generative Auction (CGA), to incorporate permutation-level externalities in online multi-slot ad auctions. Our primary theoretical results demonstrate that the classic Myerson Auction maintains its optimality when extended to permutation-level externalities. This insight drives the design of the CGA framework, which decouples the optimization of allocation and payment. Specifically, we develop an autoregressive generative model with G-E learning paradigm to optimize allocation, and learn the optimal payment by quantifying IC constraint into expected ex-post regret. Extensive offline and online experiments verify the effectiveness of CGA. Furthermore, the autoregressive generation process of CGA is not limited to specific generative model and can accommodate various advanced solutions [24, 37]. Future research will extend this contextual generative mechanism to integrate heterogeneous items from different channels.

Acknowledgments

This work was supported in part by National Key R&D Program of China (No. 2023YFB4502400), in part by China NSF grant No. 62322206, 62132018, 62025204, U2268204, 62272307, 62372296, in part by Alibaba Group through Alibaba Innovative Research Program. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

References

- [1] Qingyao Ai, Keping Bi, Jiafeng Guo, and W Bruce Croft. 2018. Learning a deep listwise context model for ranking refinement. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 135–144.
- [2] Yoram Bachrach, Sofia Ceppi, Ian A Kash, Peter Key, and David Kurokawa. 2014. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 75–92.
- [3] Santiago R Balseiro, Yuan Deng, Jieming Mao, Vahab S Mirrokni, and Song Zuo. 2021. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 132–133.
- [4] Lucas Beyer, Bo Wan, Gagan Madan, Filip Pavetic, Andreas Steiner, Alexander Kolesnikov, André Susano Pinto, Emanuele Bugliarello, Xiao Wang, Qihang Yu, et al. 2023. A study of autoregressive decoders for multi-tasking in computer vision. *arXiv preprint arXiv:2303.17376* (2023).
- [5] Chi Chen, Hui Chen, Kangzhi Zhao, Junsheng Zhou, Li He, Hongbo Deng, Jian Xu, Bo Zheng, Yong Zhang, and Chunxiao Xing. 2022. Extr: click-through rate prediction with externalities in e-commerce sponsored search. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2732–2740.

- [6] Yeshwanth Cherapanamjeri, Constantinos Daskalakis, Andrew Ilyas, and Manolis Zampetakis. 2022. Estimation of standard auction models. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 602–603.
- [7] Yuan Deng, Sébastien Lahaie, Vahab Mirrokni, and Song Zuo. 2020. A data-driven metric of incentive compatibility. In *Proceedings of The Web Conference 2020*. 1796–1806.
- [8] Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. 2021. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference 2021*. 3965–3973.
- [9] Rahul Dey and Fathi M Salem. 2017. Gate-variants of gated recurrent unit (GRU) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*. 1597–1600.
- [10] Zhijian Duan, Haoran Sun, Yurong Chen, and Xiaotie Deng. 2023. A scalable neural network for DSIC affine maximizer auction design. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*. 56169–56185.
- [11] Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. 2019. Optimal auctions through deep learning. In *International Conference on Machine Learning*. PMLR, 1706–1715.
- [12] 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.
- [13] Yufei Feng, Binbin Hu, Yu Gong, Fei Sun, Qingwen Liu, and Wenwu Ou. 2021. GRN: Generative Rerank Network for Context-wise Recommendation. *arXiv preprint arXiv:2104.00860* (2021).
- [14] Nicola Gatti, Alessandro Lazaric, and Francesco Trovò. 2012. A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 605–622.
- [15] Arpita Ghosh and Mohammad Mahdian. 2008. Externalities in online advertising. In *Proceedings of the 17th international conference on World Wide Web*. 161–168.
- [16] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [17] Patrick Hummel and R Preston McAfee. 2014. Position auctions with externalities. In *Web and Internet Economics: 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings 10*. Springer, 417–422.
- [18] Dmitry Ivanov, Iskander Safulin, Igor Filippov, and Ksenia Balabaeva. 2022. Optimal-er auctions through attention. *Advances in Neural Information Processing Systems* 35 (2022), 34734–34747.
- [19] Bernard J Jansen and Tracy Mullen. 2008. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business* 6, 2 (2008), 114–131.
- [20] Przemysław Jeziorski and Ilya Segal. 2015. What makes them click: Empirical analysis of consumer demand for search advertising. *American Economic Journal: Microeconomics* 7, 3 (2015), 24–53.
- [21] Albert Xin Jiang and Kevin Leyton-Brown. 2005. Estimating bidders' valuation distributions in online auctions. In *Proceedings of IJCAI-05 workshop on game theoretic and decision theoretic agents*. 98–107.
- [22] Ningyuan Li, Yunxuan Ma, Yang Zhao, Zhijian Duan, Yurong Chen, Zhilin Zhang, Jian Xu, Bo Zheng, and Xiaotie Deng. 2023. Learning-Based Ad Auction Design with Externalities: The Framework and A Matching-Based Approach. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1291–1302.
- [23] 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).
- [24] Shuchang Liu, Qingpeng Cai, Zhankui He, Bowen Sun, Julian McAuley, Dong Zheng, Peng Jiang, and Kun Gai. 2023. Generative flow network for listwise recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1524–1534.
- [25] Xiangyu Liu, Chuan Yu, Zhilin Zhang, Zhenzhe Zheng, Yu Rong, Hongtao Lv, Da Huo, Yiqing Wang, Dagui Chen, Jian Xu, Fan Wu, Guihai Chen, and Xiaoqiang Zhu. 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] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.
- [27] Gali Noti and Vasilis Syrgkanis. 2021. Bid prediction in repeated auctions with learning. In *Proceedings of the Web Conference 2021*. 3953–3964.
- [28] Michael Ostrovsky and Michael Schwarz. 2023. Reserve prices in internet advertising auctions: A field experiment. *Journal of Political Economy* 131, 12 (2023), 3352–3376.
- [29] Changhua Pei, Yi Zhang, Yongfeng Zhang, Fei Sun, Xiao Lin, Hanxiao Sun, Jian Wu, Peng Jiang, Junfeng Ge, Wenwu Ou, et al. 2019. Personalized re-ranking for recommendation. In *Proceedings of the 13th ACM conference on recommender systems*. 3–11.
- [30] Tian Qin, Fengxiang He, Dingfeng Shi, Wenbing Huang, and Dacheng Tao. 2022. Benefits of permutation-equivariance in auction mechanisms. *Advances in Neural Information Processing Systems* 35 (2022), 18131–18142.
- [31] 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.
- [32] Yuxin Ren, Qiya Yang, Yichun Wu, Wei Xu, Yalong Wang, and Zhiqiang Zhang. 2024. Non-autoregressive Generative Models for Reranking Recommendation. *arXiv preprint arXiv:2402.06871* (2024).
- [33] Guangyuan Shen, Shengjie Sun, Dehong Gao, Duanxiao Song, Libin Yang, Zhen Wang, Yongping Shi, and Wei Ning. 2023. EdgeNet: Encoder-decoder generative Network for Auction Design in E-commerce Online Advertising. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4274–4278.
- [34] Xiaowen Shi, Fan Yang, Ze Wang, Xiaoxu Wu, Muzhi Guan, Guogang Liao, Wang Yongkang, Xingxing Wang, and Dong Wang. 2023. PIER: Permutation-Level Interest-Based End-to-End Re-ranking Framework in E-commerce. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4823–4831.
- [35] Hal R Varian and Christopher Harris. 2014. The VCG auction in theory and practice. *American Economic Review* 104, 5 (2014), 442–445.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017), 6000–6010.
- [37] Wenjie Wang, Yiyang Xu, Fuli Feng, Xinyu Lin, Xiangnan He, and Tat-Seng Chua. 2023. Diffusion recommender model. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 832–841.
- [38] Yiqing Wang, Xiangyu Liu, Zhenzhe Zheng, Zhilin Zhang, Miao Xu, Chuan Yu, and Fan Wu. 2022. On designing a two-stage auction for online advertising. In *Proceedings of the ACM Web Conference 2022*. 90–99.
- [39] Yunjia Xi, Weiwen Liu, Xinyi Dai, Ruiming Tang, Weinan Zhang, Qing Liu, Xi-qiang He, and Yong Yu. 2021. Context-aware reranking with utility maximization for recommendation. *arXiv preprint arXiv:2110.09059* (2021).
- [40] Yunjia Xi, Weiwen Liu, Xinyi Dai, Ruiming Tang, Weinan Zhang, Qing Liu, Xi-qiang He, and Yong Yu. 2021. Context-aware reranking with utility maximization for recommendation. *arXiv preprint arXiv:2110.09059* (2021).
- [41] Yidan Xing, Zhilin Zhang, Zhenzhe Zheng, Chuan Yu, Jian Xu, Fan Wu, and Guihai Chen. 2023. Truthful auctions for automated bidding in online advertising. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*. 2915–2922.
- [42] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. 2017. Deep sets. *Advances in neural information processing systems* 30 (2017).
- [43] Ruizhe Zhang, Xiaohui Xie, Jiaxin Mao, Yiqun Liu, Min Zhang, and Shaoping Ma. 2021. Constructing a comparison-based click model for web search. In *Proceedings of the Web Conference 2021*. 270–283.
- [44] Zhanhao Zhang. 2021. A survey of online auction mechanism design using deep learning approaches. *arXiv preprint arXiv:2110.06880* (2021).
- [45] 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.

A APPENDIX

A.1 Proof of Lemma 1

Denote A^* as the allocation that maximizes the expected virtual welfare (select randomly if there exist multiple optimal allocations):

$$A^* = \mathcal{A}(b_t, \mathbf{b}_{-t}; X, \mathbf{u}) = \operatorname{argmax}_A \sum_{i=1}^n \tilde{\phi}(b_i, F_i) \cdot \Theta(\mathbf{x}_i; A, X, \mathbf{u}),$$

which can be obtained through enumeration. *W.l.o.g.*, suppose $ad_t \in A^*$, and the other ads in A^* constitute the set A_{-t}^* , then the virtual welfare of A^* can be written as:

$$\tilde{\Phi}(A^*; b_t) = \tilde{\phi}(b_t, F_t) \Theta(\mathbf{x}_t; A^*, X, \mathbf{u}) + \sum_{j \in A_{-t}^*} \tilde{\phi}(b_j, F_j) \Theta(\mathbf{x}_j; A^*, X, \mathbf{u}).$$

When the bid of ad_t rises to $b'_t > b_t$, suppose the resulting allocation maximizing virtual welfare is A' , i.e.,

$$A' = \mathcal{A}(b'_t, \mathbf{b}_{-t}; X, \mathbf{u}).$$

We can directly obtain that $ad_t \in A'$; otherwise, suppose $ad_t \notin A'$, then we have $\tilde{\Phi}(A'; b'_t) = \tilde{\Phi}(A'; b_t)$. Since A^* maximizes virtual

welfare when ad_t bids b_t and $\phi(b, F)$ is monotone non-decreasing in b [26], it follows that $\tilde{\Phi}(A'; b_t) \leq \tilde{\Phi}(A^*, b_t) \leq \tilde{\Phi}(A^*; b'_t)$. If the first ' $\leq$ ' holds as equality, meaning when ad_t bids b'_t , both allocations A' and A^* have equal and maximized virtual welfare, we choose A^* as the allocation result *w.l.o.g.* Thus, when ad_t increases its bid, the allocation remains unchanged, and so does the CTR of ad_t , validating the proposition. Otherwise, $\tilde{\Phi}(A'; b_t) < \tilde{\Phi}(A^*, b_t)$, hence, $\tilde{\Phi}(A'; b'_t) < \tilde{\Phi}(A^*; b'_t)$, contradicting the definition of A' . Therefore, we have $ad_t \in A'$.

Similarly, the virtual welfare of A' can be written as:

$$\tilde{\Phi}(A'; b'_t) = \tilde{\phi}(b'_t, F_t) \Theta(x_t; A', X, u) + \sum_{j \in A'_{-t}} \tilde{\phi}(b_j, F_j) \Theta(x_j; A', X, u).$$

Based on the maximized virtual welfare property of A^* and A' , we can deduce the following two inequalities:

$$\begin{aligned} \tilde{\phi}(b_t, F_t) \Theta(x_t; A^*, X, u) + \sum_{j \in A^*_{-t}} \tilde{\phi}(b_j, F_j) \Theta(x_j; A^*, X, u) &\geq \\ \tilde{\phi}(b_t, F_t) \Theta(x_t; A', X, u) + \sum_{j \in A'_{-t}} \tilde{\phi}(b_j, F_j) \Theta(x_j; A', X, u), \end{aligned} \quad (11)$$

$$\begin{aligned} \tilde{\phi}(b'_t, F_t) \Theta(x_t; A^*, X, u) + \sum_{j \in A^*_{-t}} \tilde{\phi}(b_j, F_j) \Theta(x_j; A^*, X, u) &\leq \\ \tilde{\phi}(b'_t, F_t) \Theta(x_t; A', X, u) + \sum_{j \in A'_{-t}} \tilde{\phi}(b_j, F_j) \Theta(x_j; A', X, u). \end{aligned} \quad (12)$$

By subtracting Inequality (12) from Inequality (11), we obtain:

$$\begin{aligned} (\tilde{\phi}(b_t, F_t) - \tilde{\phi}(b'_t, F_t)) \Theta(x_t; A^*, X, u) &\geq \\ (\tilde{\phi}(b_t, F_t) - \tilde{\phi}(b'_t, F_t)) \Theta(x_t; A', X, u). \end{aligned} \quad (13)$$

Since $\tilde{\phi}(t, F_t)$ is monotone non-decreasing in t and $b_t < b'_t$, we have:

$$\tilde{\phi}(b_t, F_t) \leq \tilde{\phi}(b'_t, F_t). \quad (14)$$

- Case 1: $\tilde{\phi}(b_t, F_t) - \tilde{\phi}(b'_t, F_t) < 0$. Then we have:

$$\Theta(x_t; A^*, X, u) \leq \Theta(x_t; A', X, u), \text{ i.e.,}$$

$$\Theta(x_t; \mathcal{A}(b_t, b_{-t}; X, u), X, u) \leq \Theta(x_t; \mathcal{A}(b'_t, b_{-t}; X), X, u),$$

demonstrating that the allocation rule $\mathcal{A}$ is monotone.

- Case 2: $\tilde{\phi}(b_t, F_t) - \tilde{\phi}(b'_t, F_t) = 0$. By swapping the LHS and RHS of Inequality (12) and then adding it to Inequality (11), we obtain:

$$\begin{aligned} \tilde{\phi}(b_t, F_t) \Theta(x_t; A^*, X, u) + \tilde{\phi}(b'_t, F_t) \Theta(x_t; A', X, u) &\geq \\ \tilde{\phi}(b_t, F_t) \Theta(x_t; A', X, u) + \tilde{\phi}(b'_t, F_t) \Theta(x_t; A^*, X, u). \end{aligned} \quad (15)$$

Substituting $\tilde{\phi}(b_t, F_t) = \tilde{\phi}(b'_t, F_t)$ into Equation (15), we have, the ' $\geq$ ' in Inequality (15) holds as equality. Therefore, both Inequality (11) and Inequality (12) hold as equalities, which means that both A^* and A' are the allocations that maximize virtual welfare after ad_t increases her bid to b'_t . Therefore, *w.l.o.g.*, we choose A^* as the final allocation result, indicating that the allocation outcome remains unchanged after the increase in ad_t 's bid. Consequently, the CTR of ad_t also remains unchanged. Thus, the allocation rule $\mathcal{A}$ is monotone.