# Two-Stage Auctions with Bid Refinement for Online Advertising

Yidan Xing
School of Computer Science,
Shanghai Jiao Tong University
Shanghai, China
katexing@sjtu.edu.cn

Rui Guo
Taobao & Tmall Group of Alibaba
Beijing, China
xianqiao.gr@alibaba-inc.com

Yixin Tao
Shanghai University of Finance and Economics
Shanghai, China
taoyixin@mail.shufe.edu.cn

Dagui Chen
Taobao & Tmall Group of Alibaba
Beijing, China
dagui.cdg@alibaba-inc.com

Zhenzhe Zheng*
School of Computer Science,
Shanghai Jiao Tong University
Shanghai, China
zhengzhenzhe@sjtu.edu.cn

Jian Xu
Taobao & Tmall Group of Alibaba
Beijing, China
xiyu.xj@alibaba-inc.com

Fan Wu
School of Computer Science,
Shanghai Jiao Tong University
Shanghai, China
fwu@cs.sjtu.edu.cn

## Abstract

To balance prediction accuracy and system latency, large-scale online ad auctions employ two-stage architectures. These systems first retrieve a candidate ads subset using coarse quality metrics before finalizing auction outcomes with refined metrics. However, existing implementations typically elicit user-specific bids only once, overlooking the impact of real-time quality metrics on advertiser valuations and thereby limiting allocation efficiency. Motivated by recent industry practice, we investigate the design of two-stage auctions that allow advertisers to submit and update their bids, with second-stage bids serving as refinements of the initial ones. We derive the incentive-compatible (IC) conditions and analyze the revenue properties of this two-stage auction. Notably, an additional entry fee is required to prevent inflated initial bids, which compromises the standard ex-post individual rationality (IR) property. To address this, we propose a dynamic two-stage auction that adopts realization-dependent entry fees with discounts. By leveraging historical bidding information, our mechanism guarantees approximate ex-ante IC across both stages and restores ex-post IR.

## CCS Concepts

• **Information systems → Computational advertising**; • **Theory of computation → Algorithmic mechanism design**.

## Keywords

Mechanism Design; Online Advertising; Two-Stage Auction

*Zhenzhe Zheng is the corresponding author.


*KDD '26, Jeju Island, Republic of Korea*

ACM ISBN 979-8-4007-2259-2/2026/08
https://doi.org/10.1145/3770855.3817645

**ACM Reference Format:**
Yidan Xing, Rui Guo, Yixin Tao, Dagui Chen, Zhenzhe Zheng, Jian Xu, and Fan Wu. 2026. Two-Stage Auctions with Bid Refinement for Online Advertising. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26), August 09–13, 2026, Jeju Island, Republic of Korea.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3770855.3817645

## 1 Introduction

With the advancement of real-time bidding techniques [31], the frequency and scale of ad auctions have increased dramatically. Upon each user request, advertisers submit user-specific bids to the platform, and the platform then determines the ads to display via real-time auctions that consider both bid values and ad quality metrics—such as click-through rate (CTR) and conversion rate (CVR). In addition to ad content, these metrics incorporate granular user features and historical behaviors to assess the ad's relevance to each unique request. To balance prediction accuracy and system latency across a large number of ad candidates, modern online ad systems typically adopt a two-stage architecture [18, 36]. In this design, the platform first computes coarse predictions of ad quality for all candidates and, together with the submitted bids, uses them to select a subset of ads during a *pre-auction stage.* Only this subset proceeds to a second, *formal auction stage*, where the platform could then afford to compute the refined ad quality metrics with computationally intensive methods for finalizing allocation and payment. While the refinement of quality metrics is inherent to this architecture, the corresponding protocol for handling bids throughout this two-stage process remains a non-standardized design choice.

In industrial settings, a common practice is to retrieve the candidate ad subset using average historical bids, and elicit user-specific real-time bids only in the second stage. An alternative approach, proposed in the literature [36, 40], collects bids only during the pre-auction stage and uses the same bid values throughout both stages. However, advertisers' valuations and corresponding bids depend critically on real-time predictions of ad quality. Bids based on historical or outdated metrics often fail to reflect the advertiser's true valuation at the moment. For example, advertisers optimizing for conversions base their bids primarily on user-specific CVR predictions, which improve with real-time modeling and become more accurate when refined CVR predictions are available. Consequently, to improve allocation efficiency, platforms should ideally elicit and use real-time, up-to-date bids in both stages of the auction.

Driven by this rationale, Taobao, one of the major online ad platforms, has upgraded its system architecture to support this dual real-time bidding approach at scale. While Taobao's current implementation retains the auction rules of existing approaches, the system incorporates updated bids at both stages. Empirical analysis across billions of auctions reveals that the refinement of ad quality metrics leads to an average twofold change in valuations, underscoring the significance of bid refinement. However, this pioneering approach is susceptible to strategic manipulations: because final auction outcomes depend solely on the refined second-stage bids, advertisers can inflate their pre-auction bids to ensure their entry to the formal auction without incurring any cost. As a result, Taobao currently restricts this feature to trusted internal advertisers, limiting the scope of bid refinement in general ad auctions.

Motivated by these industry practices, we aim to formally investigate the design of two-stage ad auctions that allow advertisers to submit and refine their bids from a game-theoretical perspective. Specifically, we focus on two-stage incentive-compatible (IC) auctions, where advertisers are incentivized to truthfully report their valuations at both stages. Unlike in traditional multi-dimensional auctions, an advertiser's initial valuation does not contribute to her utility directly; instead, it implicitly influences subsequent allocation and payment. This asymmetry introduces an incentive for advertisers to inflate their initial bids to increase their chances of being selected for the second stage, especially under improperly designed mechanisms.

To preserve the well-ordering of bidding, we provide formal definitions and analysis for two-stage IC conditions. Our result enables the platform to truthfully implement any monotone allocation rule under a newly proposed two-stage monotonicity concept. A key finding is that advertisers granted access to the second stage must pay an entry fee alongside standard single-stage payments, irrespective of their refined bids. To guide revenue optimization, we derive a closed-form expression for the expected revenue of any two-stage auction, which can be viewed as a transformation of the classical virtual value function [29]. However, due to the indispensability of entry fees, a broad class of common two-stage IC auctions fails to satisfy ex-post individual rationality (IR) conditions for general non-trivial allocation rules, *i.e.*, some advertisers may be charged without receiving an allocation. Furthermore, these entry fees result in higher expected payments for advertisers compared to incumbent auction mechanisms.

To overcome this limitation, we propose a dynamic two-stage auction with discounts in entry fees that simultaneously achieves ex-post IR and approximate ex-ante IC. This auction incorporates realization-dependent entry fees tailored to the advertiser's realized capacity to pay, and allows the platform to provide flexible discounts on this basis. Inspired by non-monetary mechanism design, we track advertisers' historical bidding behavior and compare their cumulative discount usage against prior expectations. Advertisers exhibiting abnormally high discount consumption are subject to exclusion, ensuring that the utility gain from strategic manipulation vanishes over time. For the widely adopted rank-score allocation rules, our auction can either charge a full entry fee by setting an equivalent bid-dependent reserve price, or completely eliminate the entry fees to align with standard single-stage payments. Crucially, as the discount computation and monitoring process is simple based on entry fees, our dynamic two-stage auction remains practical for implementation. Our results offer valuable insights into designing two-stage ad auctions that effectively integrate bid refinement while preserving economic robustness in online advertising.

The main contributions can be summarized as follows:

- We study the auction design problem for a new two-stage model with bid refinement in online advertising. We provide industrial evidence to support this model and motivate the necessity for a rigorous game-theoretical analysis.
- We derive the IC conditions and analyze the revenue performances of this model. We find that general two-stage IC auctions are unable to satisfy ex-post IR due to additional entry fees.
- We propose a dynamic two-stage auction with flexible discounts to achieve ex-post IR and approximate ex-ante IC. This mechanism can convert any monotone two-stage allocation rule to a bid refinement scheme and remains simple in auction format.

## 2 Two-Stage Auctions: Industry Practice

In this section, we provide the background on two-stage ad auctions and present the recent deployment by Taobao to demonstrate the feasibility and effects of bid refinement.

### 2.1 Two-Stage Auctions in Online Ad System

The two-stage architecture has been widely adopted as a fundamental infrastructure to address real-time user requests in both industrial recommender system [7, 10, 20] and online ad system [12, 36, 40]. In two-stage ad auctions, the formal auction stage operates similarly to standard (single-stage) ad auctions but typically includes only a few *hundred* ad candidates. This constraint arises from the computational cost of producing refined predictions for ad quality metrics [28, 41]. Due to these resource limits, the pre-auction stage must filter *tens of thousands* of relevant ad candidates down to a manageable subset. The typical goal is to identify a candidate subset that is likely to yield high revenue in the formal auction. This filtering is usually based on a *pre-ranking score*, computed using coarse (low-cost) prediction of quality metrics and advertiser bids. The specific method used to compute pre-ranking scores can vary across platforms [12, 36], but generally, ads with the highest pre-ranking scores are admitted to the formal auction.

A key modification explored in this work—and recently adopted in practice by Taobao—is to *collect advertiser bids at both stages*, rather than eliciting a single set of real-time bids. When feasible, this allows advertisers' bids to dynamically reflect their current valuations, enhancing the platform's ability to make more accurate and efficient allocation decisions.

*2.1.1 Deployment Overview.* Previously, Taobao adopted a *time-weighted average of historical bids* to simulate bid values in the pre-auction stage (see Figure 1a). This was combined with coarse quality metrics computed *in parallel* to determine the pre-ranking

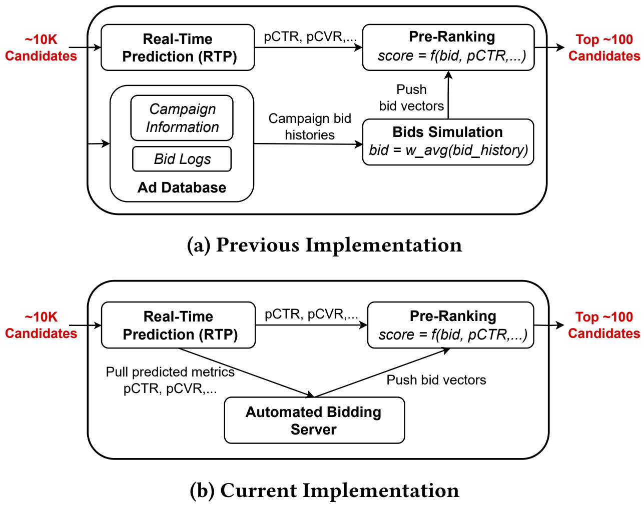


**Figure 1: Online Deployment of Pre-Auction Stage**

score. Real-time, user-specific bids were only collected after an advertiser's ad passed the pre-auction filter and entered the formal auction stage, where refined metrics were available.

To enable real-time bidding in both stages, Taobao redesigned its system architecture so that, in the pre-auction stage, coarse quality metrics are computed first, followed *sequentially* by the advertiser's bidding decision, mirroring the flow used in the formal auction stage (see Figure 1b). Such implementation for the pre-auction stage is more complex than that of the formal auction stage due to the significantly larger scale of ad candidates, making distributed computation more demanding. Despite the increased latency, the improvement of utilizing real-time, user-specific bids for subset selection outweighs its engineering complexity.

*2.1.2 Bid Discrepancies Between Stages.* To quantify the discrepancies between bids submitted at the pre-auction and formal auction stages, we define the *absolute logarithm of pre-bid ratio (ALPR)* for each ad request as

$$\mathrm{ALPR} = \left(\sum_{i\in C}\left|\log_2 \frac{\hat{v}_i}{\hat{w}_i}\right|\right) / |C|,$$

where $C$ denotes the set of advertisers who enter the formal auction stage for the current request. Here, $\hat{v}_i$ and $\hat{w}_i$ represent advertiser $i$'s initial (pre-auction) and refined (formal auction) bids, respectively. We analyzed ALPR across all ad requests from date 09-29 to 10-06, covering billions of auction instances. The average ALPR during this period was *1.04*, indicating that, on average, an advertiser's refined bid was approximately either twice or half the value of their initial bid. These bid discrepancies largely stem from differences between coarse and refined quality metric predictions used in the two stages. This significant variation underscores the importance of bid refinement in the formal auction stage. Even though initial bids can be collected in the pre-auction stage at scale, relying solely on them without subsequent refinement could severely undermine the efficiency of ad allocation.

Taobao has *fully deployed* this new two-stage bidding architecture in its major display advertising services, supporting over 300,000 advertisers and processing more than 10 billion ad requests per day[1]. Representative deployment scenarios include "Recommended for You" ad slots on the homepage and post-purchase pages of Taobao's mobile app. The resulting P99 latency, *i.e.*, the time by which 99% of ad requests are completed, is approximately 355 milliseconds, indicating that most auctions are resolved within acceptable real-time constraints.

## 2.2 Motivation of Auction Design

The feasibility of the above deployment relies on a crucial condition: all participants are internal advertisers treated as *trusted*. That is, they either use automated bidding tools provided by Taobao or adopt a fixed, pre-determined bid for all relevant auctions. Due to the trustworthy guarantee of involved advertisers, Taobao is able to treat initial bids solely as reference signals to help select a promising candidate subset.

When the platform possesses full knowledge of how participants evaluate an impression across two stages, it is also feasible to adopt *automatic bid refinement*, thereby saving the efforts of explicitly eliciting refined bids. In real-time bidding, advertisers typically adopt some bidding formulas parametrized by ad quality metrics and control variables reflecting the advertiser's anticipation of real-time market conditions [1, 19, 37]; these formulas effectively serve as their valuation function in personalized contexts. If the platform had access to these functions, it could automatically refine each bidder's initial bid upon the update of ad quality metrics, rendering the second-stage bid elicitation unnecessary.

However, for general advertisers, the platform typically lacks access to these specific valuation functions. Advertisers may employ some complex, proprietary strategies [6, 39] to maintain a strategic competitive advantage. Moreover, it is difficult to reach a consensus on the credibility of ad metric predictions and automatic bid refinement when fully conducted by the platform. Thus, eliciting bids at both stages is more suitable for general scenarios to leverage refined valuations. Yet, this flexibility grants strategic advertisers significant opportunities to manipulate bids and extract higher utility, which would severely distort the pre-ranking scores and degrade auction performance. Given the scalability and practical viability of collecting real-time bids at both stages, it becomes essential to develop a formal game-theoretic understanding of two-stage auctions. In particular, properly resolving the involved incentive issues is key to ensuring that the auction remains robust and effective when extended to a broader set of advertisers.

## 3 Notations and Definitions

Based on the new features of two-stage ad auctions discussed above, we formulate our formal model as follows. Suppose there are $n$ advertisers[2] competing for display opportunities provided by an ad platform. The process of each auction can be described as follows.

1) Pre-auction stage: Each advertiser $i \in [n]$ learns her initial valuation $v_i$ on the current display opportunity. Due to varying relevance between advertisers and users, we assume $v_i$ is drawn independently from a personalized distribution $f_i(\cdot)$. While $f_i$ is known by the platform, $v_i$ is privately revealed to advertiser $i$. The platform asks advertisers to submit initial bids $\hat{v}_i$, and determines

[1] Statistics in this section are counted during the same period as the ALPR metric.
[2] We use advertisers with bidders interchangeably in this work.

the bidder subset to enter the formal auction stage with subset selection rule $y_S$. We use $y_S(\hat{\mathbf{v}})$ to denote the probability that a bidder subset $S \subseteq [n]$ is selected by the platform to enter the formal auction stage under the initial bid profile $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_n)$.

2) Formal auction stage: Each advertiser $i$ entering the formal auction stage learns her refined valuation $w_i$ on the display opportunity, independently drawn from a distribution $g_i(\cdot; v_i)$. As $w_i$ is a refined evaluation based on $v_i$, we assume the distribution of $w_i$ depends on her initial valuation $v_i$. We adopt the concept of *first-order stochastic dominance* [9] to describe the relation between distributions $g_i(\cdot; v_i)$ of different initial values $v_i$ as

$$\forall i, v_i' \geq v_i, w_i : \; G_i(w_i; v_i') \leq G_i(w_i; v_i),$$

where $G_i(\cdot; v_i)$ is the CDF of $g_i(\cdot; v_i)$. That is, when an advertiser's initial valuation increases, the probability of her refined valuation exceeding any fixed $w_i$ is non-decreasing. Distributions $g_i$ are public information, while $w_i$ are privately known by advertiser $i$. The platform asks advertisers to submit their refined bids $\hat{w}_i$ to determine the auction outcome based on the allocation rule $x$ and payment rule $p$ of the formal auction stage. We use $x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ to denote the allocation to bidder $i$ when the selected bidder subset is $S$ and their refined bid profile is $\hat{\mathbf{w}} = (\hat{w}_1, \dots, \hat{w}_n)$, and denote the corresponding payment charged to bidder $i$ as $p_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$, with $\hat{w}_i = 0$ and $x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = 0$ for $i \notin S$. We assume there is an upper bound $\overline{u}$ for $v_i, \hat{v}_i, w_i$ and $\hat{w}_i$, *i.e.*, $v_i, \hat{v}_i, w_i, \hat{w}_i \leq \overline{u}$.

The realized utility of bidder $i$ with refined valuation $w_i$ when the bids profile is $(\hat{\mathbf{v}}, \hat{\mathbf{w}})$ and the entering bidder subset is $S$ could then be defined as

$$u_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; w_i, S) = w_i \cdot x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) - p_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S).$$

We also define the expected utility of bidder $i$ in pre-auction stage when the initial bid profile is $\hat{\mathbf{v}}$ and her true initial valuation is $v_i$ as

$$U_i(\hat{\mathbf{v}}; v_i) = \sum_{S \subseteq [n]} y_S(\hat{\mathbf{v}}) \cdot \mathbb{E}_{\mathbf{w} \sim \mathbf{g}(\cdot; v_i, \hat{\mathbf{v}}_{-i})} [u_i(\hat{\mathbf{v}}, \mathbf{w}; w_i, S)],$$

assuming all the bidders report true refined valuations in the formal auction stage with initial valuations $(v_i, \hat{\mathbf{v}}_{-i})$, where we use $\mathbf{w} \sim \mathbf{g}(\cdot; v_i, \hat{\mathbf{v}}_{-i})$ to represent $w_i \sim g_i(\cdot; v_i)$ and $w_j \sim g_j(\cdot; \hat{v}_j)$ for $j \neq i$.

In this work, we consider the design of two-stage incentive compatible (IC) direct-revelation auctions. Since reporting the initial bid and refined bid formulates a sequential decision problem for advertisers, the IC conditions for the two stages are defined in a backward manner. The allocation rule of a two-stage auction is defined by $(y_S, x)$, which constitutes a complete auction together with the payment rule $p$.

Definition 3.1. *A two-stage auction mechanism $(y_S, x, p)$ satisfies two-stage IC if the following conditions hold:*

*(1) (Ex-post IC for formal auction stage)*

$$\forall \hat{\mathbf{v}}, \hat{\mathbf{w}}, i, w_i, S : u_i(\hat{\mathbf{v}}, w_i, \hat{\mathbf{w}}_{-i}; w_i, S) \geq u_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; w_i, S). \tag{1}$$

*(2) (Interim IC for pre-auction stage)*

$$\forall \hat{\mathbf{v}}, v_i, i : U_i(v_i, \hat{\mathbf{v}}_{-i}; v_i) \geq U_i(\hat{\mathbf{v}}; v_i). \tag{2}$$

Given that an advertiser has entered the formal auction stage with the pre-auction stage context $\hat{\mathbf{v}}$ and $S$, the remaining auction process becomes similar to the classical single-stage auctions. Following this idea, conditions (1) for the formal auction stage are defined to elicit true refined valuations under any potential pre-auction stage context. By conditions (1), we could assume advertisers to always report true refined valuations in the formal auction, which simplifies the computation of expected utility for advertisers in the pre-auction stage. To see conditions (2) satisfy classical Bayesian-Nash incentive compatible (BNIC) conditions for initial valuations, we note that $U_i(\hat{\mathbf{v}}; v_i)$ becomes the true expected utility for advertiser $i$ when other advertisers bid truthfully.

We also desire the auction mechanism to satisfy individual rationality (IR) conditions.

Definition 3.2. *A two-stage auction $(y_S, x, p)$ satisfies two-stage interim IR if $\forall \mathbf{v}, i : U_i(\mathbf{v}; v_i) \geq 0$; and satisfies ex-post IR if $\forall \mathbf{v}, \mathbf{w}, i, S :$ $u_i(\mathbf{v}, \mathbf{w}; w_i, S) \geq 0$.*

As ex-post IR guarantees non-negative realized utility for any potential realization of bidder subset and refined bids, it provides a stronger guarantee than the two-stage interim IR on expected utility. For any two-stage auction considered in the work, we assume $p_i(\mathbf{v}, \mathbf{w}; S) \geq 0$, with $x_i(0, \mathbf{v}_{-i}, 0, \mathbf{w}_{-i}; S) = p_i(0, \mathbf{v}_{-i}, 0, \mathbf{w}_{-i}; S) = 0$, $\forall i, S, \mathbf{v}, \mathbf{w}$, to guarantee a safe outside option. For convenience in notations, we define the *two-stage combined* allocation rule $z$ and payment rule $q$ for two-stage auction $(y_S, x, p)$ as

$$\begin{aligned} z_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}) &= \sum_{S \subseteq [n]} y_S(\hat{\mathbf{v}}) \cdot x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S), \\ q_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}) &= \sum_{S \subseteq [n]} y_S(\hat{\mathbf{v}}) \cdot p_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S), \end{aligned} \tag{3}$$

which represents expected allocation and payment for a realization of $(\hat{\mathbf{v}}, \hat{\mathbf{w}})$ over the randomness of subset selection. We may also use $\mathbb{E}_{S \sim y_S(\hat{\mathbf{v}})}$ to represent such randomness, *e.g.*, $z_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}) = \mathbb{E}_{S \sim y_S(\hat{\mathbf{v}})} [x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)]$.

# 4 Characterization and Properties

In this section, we present the characterization results for two-stage IC auctions and analyze their associated revenue and IR properties. Detailed proofs of our results are provided in Appendix A.

## 4.1 Conditions for Two-Stage IC

After a bidder enters the formal auction stage, the bidder subset $S$ and initial bids $\hat{\mathbf{v}}$ have been fixed. Therefore, the ex-post IC conditions (1) for refined bids $\hat{\mathbf{w}}$ are analogous to the classical IC conditions in single-stage auctions [29]. By this observation, we have $p_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \int_0^{\hat{w}_i} \tilde{w}_i \cdot \frac{\partial x_i(\hat{\mathbf{v}}, \tilde{w}_i, \hat{\mathbf{w}}_{-i}; S)}{\partial \tilde{w}_i} \, \mathrm{d}\tilde{w}_i + p_i(\hat{\mathbf{v}}, 0, \hat{\mathbf{w}}_{-i}; S)$, and $x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ non-decreasing in $\hat{w}_i$, $\forall i, S, \hat{\mathbf{v}}, \hat{\mathbf{w}}$, to incentivize truthful report of $w_i$. In classical auctions, the "starting point" payment term when the valuation of a bidder is 0, analogous to $p_i(\hat{\mathbf{v}}, 0, \hat{\mathbf{w}}_{-i}; S)$ here, is usually assigned to be 0 to ensure the IR property. However, when bidders need to first compete in a pre-auction stage in our two-stage context, the design of term $p_i(\hat{\mathbf{v}}, 0, \hat{\mathbf{w}}_{-i}; S)$ becomes the key to incentivize true initial bids, as simply setting $p_i(\hat{\mathbf{v}}, 0, \hat{\mathbf{w}}_{-i}; S) = 0$ represents no opportunity cost to participate in the formal auction and leads to overbidding of initial valuations.

To design an appropriate $p_i(\hat{\mathbf{v}}, 0, \hat{\mathbf{w}}_{-i}; S)$, the underlying insight is to view the outcome of bidder subset selection as an allocation of potential opportunities to participate in formal auctions. Since the value of those opportunities is essentially the expected utility a bidder may obtain in the formal auction, we could adapt the spirit of Myerson's result and apply it iteratively on the expected utility derived from the single-stage characterization. This provides a reduction from the design of two-stage IC auctions to a two-stage allocation rule.

Theorem 4.1. *For any two-stage allocation rule* $(y_S, x)$ *satisfying*

- $x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ *non-decreasing in* $\hat{w}_i$*; and*
- $z_i(\hat{\mathbf{v}}, \hat{\mathbf{w}})$ *non-decreasing in* $\hat{v}_i$

*for any* $i \in [n], S \subseteq [n]$ *and potential bid profiles* $\hat{\mathbf{v}}, \hat{\mathbf{w}}$*, define the corresponding payment rule* $p$ *as* $p_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = p_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) + \tau_i(\hat{\mathbf{v}})$*, with standard single-stage payment*

$$p_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \int_0^{\hat{w}_i} \tilde{w}_i \cdot \frac{\partial x_i(\hat{\mathbf{v}}, \tilde{w}_i, \hat{\mathbf{w}}_{-i}; S)}{\partial \tilde{w}_i} \, \mathrm{d}\tilde{w}_i$$

*and an additional entry fee*

$$\tau_i(\hat{\mathbf{v}}) = \mathbb{E}_{\mathbf{w}_{-i} \sim \mathbf{g}_{-i}(\cdot; \hat{\mathbf{v}}_{-i})} \left[ \int_0^{\hat{v}_i} \int_0^{\overline{u}} (1 - G_i(w_i; \tilde{v}_i)) \cdot \frac{\partial z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, w_i, \mathbf{w}_{-i})}{\partial \tilde{v}_i} \, \mathrm{d}w_i \, \mathrm{d}\tilde{v}_i \right], \tag{4}$$

*then the two-stage auction* $(y_S, x, p)$ *is two-stage IC and interim IR.*

We term the allocation rules $(y_S, x)$ satisfying the non-decreasing conditions in Theorem 4.1 to be *two-stage monotone.* That is, fixing the realized valuations of other bidders, reporting a higher initial bid $\hat{v}_i$ brings the advertiser a larger two-stage combined allocation $z_i$, and reporting a higher refined bid $\hat{w}_i$ brings the advertiser a larger realized allocation $x_i$. As we will see in Section 4.2, these conditions are natural to hold for allocation rules in ad auctions.

In Theorem 4.1, the payment charged to advertisers are split into two parts: (1) the standard single-stage payment, *i.e.*, $p_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \int_0^{\hat{w}_i} \tilde{w}_i \cdot \frac{\partial x_{i,S}(\hat{\mathbf{v}}, \tilde{w}_i, \hat{\mathbf{w}}_{-i})}{\partial \tilde{w}_i} \, \mathrm{d}\tilde{w}_i$; and (2) an additional *entry fee* $\tau_i(\hat{\mathbf{v}})$ corresponding to the opportunity cost of participating in the formal auction. In Equation (4), term $\int_0^{\overline{u}} (1 - G_i(w_i; \tilde{v}_i)) z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, w_i, \hat{\mathbf{w}}_{-i}) \, \mathrm{d}w_i$ represents the expected utility of advertiser $i$ in the formal auction when her initial valuation is $\tilde{v}_i$, fixing the bids of other advertisers.

Due to the inherent uncertainty of the future in two-stage auctions, there may exist a multitude of payment rules to constitute a two-stage IC auction for the same allocation rule, which is distinct from the single-stage characterization result. Specifically, using any entry fee rule $\tau_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}_{-i}; S)$ conditional on $S$ and $\hat{\mathbf{w}}_{-i}$ in substitute of $\tau_i(\hat{\mathbf{v}})$, with $\mathbb{E}_{\hat{\mathbf{w}}_{-i} \sim \mathbf{g}(\cdot; \hat{\mathbf{v}}_{-i}), S \sim y_S}[\tau_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}_{-i}; S)] = \tau_i(\hat{\mathbf{v}})$, also provides a two-stage IC and interim IR mechanism. We present the payment format as $\tau_i(\hat{\mathbf{v}})$ in Theorem 4.1 due to its simplicity in operation: the platform only needs to charge a fixed entry fee for each bidder depending on their initial bids, which could be paid up before the formal auction begins. For ease of presentation, Theorem 4.1 is written using ordinary derivatives. When the allocation rule is monotone but discontinuous, the derivative notation is interpreted in the Lebesgue–Stieltjes sense induced by the monotone allocation rule. For continuously differentiable allocation rules, this convention coincides with the displayed formula.

## 4.2 Example: Allocation with Rank Scores

In online advertising, a group of allocation rules based on *rank scores* is widely adopted due to its simplicity in implementation and flexibility in design [23, 26, 33], covering both subset selection [12, 13, 36] and auction allocation [26, 33].

*4.2.1 Rank Scores in Online Advertising.* The working principle of the allocation rule with rank scores is formally described as follows. The ad platform determines the (personalized) rank score functions $h_{i,0}, h_{i,1} : \mathbb{R}^+ \to \mathbb{R}$ for each advertiser $i$ before the auction starts, with $h_{i,0}$ for the pre-auction stage, and $h_{i,1}$ for the formal auction stage. Note that $h_i$ can depend on other contextual features or prior information irrelevant of submitted bids.

Given bids of advertiser $i$, the platform compute the corresponding rank scores for her, denoted as $h_{i,0}(\hat{v}_i)$ for rank score in the pre-auction stage and $h_{i,1}(\hat{w}_i)$ for the formal auction stage. The platform greedily chooses $k$ advertisers ($k < n$) with the highest initial rank scores $h_{i,0}(\hat{v}_i)$ to enter the formal auction, and finally allocates to the advertiser with the highest refined rank score $h_{i,1}(\hat{w}_i)$. The allocation rule can be formally defined as:

$$y_S(\hat{\mathbf{v}}) = \mathbb{I}\{S = \operatorname*{argmax}_{S' \subset [n], |S'| = k} \textstyle\sum_{i \in S'} h_{i,0}(\hat{v}_i)\},$$
$$x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \mathbb{I}\{i = \operatorname*{argmax}_{j \in S} h_{j,1}(\hat{w}_j)\},$$

where ties are broken in an arbitrary fixed order. As the rank scores are non-decreasing in bids for both stages, fixing the bids of other advertisers, reporting a higher bid would not lead the final allocation to decrease, thus satisfying the two-stage monotone property.

*4.2.2 Computation of Entry Fees.* As a consequence of the non-decreasing rank score functions, there exist some threshold bids for each advertiser to enter the second stage or win the auction, leading the entry fee in Equation (4) to be largely simplified. To facilitate our discussion, we use $h_{-i,0}^{(k)}(\hat{\mathbf{v}}_{-i})$ to denote the $k^{th}$ highest initial rank score besides that of bidder $i$, and use $S^{(k)}(\hat{\mathbf{v}})$ to denote the subset of $k$ advertisers entering the formal auction. Formally, we may define the *threshold initial bid* and *threshold refined bid* respectively as

$$r_{i,0}(\hat{\mathbf{v}}_{-i}) = \inf_{v_i}\{h_{i,0}(v_i) \geq h_{-i,0}^{(k)}(\hat{\mathbf{v}}_{-i})\}$$
$$r_{i,1}(\hat{\mathbf{w}}_{-i}; S) = \inf_{w_i}\{h_{i,1}(w_i) \geq h_{j,1}(\hat{w}_j), \forall j \in S \backslash \{i\}\}.$$

That is, an advertiser $i$ enters the formal auction only if $\hat{v}_i \geq r_{i,0}(\hat{\mathbf{v}}_{-i})$, and then wins the auction only if $\hat{w}_i \geq r_{i,1}(\hat{\mathbf{w}}_{-i}; S^{(k)}(\hat{\mathbf{v}}))$.

Proposition 4.2. *In the allocation with rank scores setting, charging the following entry fee satisfies two-stage IC:*

$$\tau_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}_{-i}; S) = \begin{cases} \mathop{\mathbb{E}}_{\tilde{w}_i \sim G_i(\cdot | r_{i,0}(\hat{\mathbf{v}}_{-i}))} \left[ \left(\tilde{w}_i - r_{i,1}(\hat{\mathbf{w}}_{-i}; S)\right)^+ \right], & \textit{if } i \in S, \\ 0, & \textit{otherwise}, \end{cases} \tag{5}$$

*where* $(x)^+ := \max(x, 0)$.

Observing $\left(\tilde{w}_i - r_{i,1}(\hat{\mathbf{w}}_{-i}; S)\right)^+$ is the utility of an advertiser with valuation $\tilde{w}_i$ in standard single-stage auctions, the entry fee can be simply computed by simulating the expected utility of $i$ under different $\tilde{w}_i$, assuming her initial valuation is the threshold initial bid $r_{i,0}(\hat{\mathbf{v}}_{-i})$ with fixed $S$ and $\hat{\mathbf{w}}_{-i}$.

## 4.3 Revenue and IR Properties

Next, we consider the revenue and IR properties to inspect the overall impacts brought by additional entry fees in two-stage auctions.

Theorem 4.3. *Suppose $w_i = v_i + \epsilon_i$, where $\epsilon_i$ is an additive noise independent of $v_i$, and $g_i(\cdot; v_i)$ is the induced conditional distribution supported on $v_i + \text{supp}(\epsilon_i) \subseteq \mathbb{R}_+$. Consider a two-stage IC auction with allocation rule $(y_S, x)$ and $y_S(0, \mathbf{v}_{-i}) = 0, \forall S \ni i$. Then its expected revenue is*

$$\mathbb{E}_{\mathbf{v}\sim\mathbf{f},\mathbf{w}\sim\mathbf{g}(\cdot;\mathbf{v})}\left[\sum_{i=1}^{n} z_i(\mathbf{v},\mathbf{w}) \cdot \left(w_i - \frac{1-F_i(v_i)}{f_i(v_i)}\right)\right], \tag{6}$$

*where $z_i(\mathbf{v}, \mathbf{w})$ is the combined allocation rule of $(y_S, x)$.*

Recall that the expected revenue of standard single-stage IC auction only eliciting initial valuations $\mathbf{v}$ (with allocation rule $a$) is $\mathbb{E}_{\mathbf{v}\sim\mathbf{f}}\left[\sum_{i=1}^{n} a_i(\mathbf{v}) \cdot \left(v_i - \frac{1-F_i(v_i)}{f_i(v_i)}\right)\right]$, the virtual value function of two-stage IC auctions replaces the initial value $v_i$ by the refined valuation $w_i$, while the inverse hazard rate term remains unchanged. In other words, not only could the social welfare of the auction be improved by allocating based on the latest valuations, but the platform may also exploit the difference between $v_i$ and $w_i$ to enhance revenue. As our results provide a truthful reduction for eliciting bids, we can derive approximation results on social welfare and revenue based on existing monotone subset selection algorithms [13] with moderate regularity assumptions on virtual value function. Importantly, as our two-stage auction generalizes the single-stage auction, any outcomes achieved by single-stage auctions with initial valuations are trivially attainable here.

*Failure of Ex-post IR.* While entry fees are necessary to elicit true valuations in a two-stage auction, a flip side of this design is the violation of ex-post IR properties: advertisers are liable for positive entry fees regardless of whether they receive the final allocation. Unfortunately, such situations hold for general non-trivial two-stage IC auctions. To formalize this idea, for any allocation rule $(y_S, x)$, define the expected combined allocation of bidder $i$ under initial bid $\hat{v}_i$, fixing other initial bids $\hat{\mathbf{v}}_{-i}$, as $Z_i(\hat{v}_i, \hat{\mathbf{v}}_{-i}; \mathbf{v}) := \mathbb{E}_{\mathbf{w}\sim\mathbf{g}(\cdot;\mathbf{v})}\left[z_i(\hat{v}_i, \hat{\mathbf{v}}_{-i}, \mathbf{w})\right]$. We call a two-stage allocation rule *trivial* if, for any $i, \mathbf{v}, \hat{\mathbf{v}}_{-i}$, the function $Z_i(\hat{v}_i, \hat{\mathbf{v}}_{-i}; \mathbf{v})$ is constant in $\hat{v}_i$. Otherwise, the allocation rule is *non-trivial.* For Proposition 4.4, we assume the refined-value distributions are regular and fully supported: for any $i, v_i \in [0, \overline{u}]$, $G_i(\cdot; v_i)$ is atomless and admits a density $g_i(w_i; v_i) > 0$ on $(0, \overline{u})$.

Proposition 4.4. *Under the regular full-support assumption above, for any non-trivial two-stage monotone allocation rule, there does not exist an ex-post IR auction to implement it and satisfy two-stage IC.*

## 5 Ex-Post IR Two-Stage Auction with Discounts

Our characterization results outline a promising paradigm of two-stage ad auctions: the platform could implement any desirable two-stage monotone allocation rule for general advertisers. Despite those desirable properties, the accompanying features of two-stage IC auctions, including the charge of additional entry fee and the violation of ex-post IR conditions (Proposition 4.4), diverge from the current payment rules familiar to advertisers, which may pose additional challenges for its promotion. With the above considerations, this section investigates a design for two-stage auctions that balances the necessary auction properties for online advertising.

### 5.1 From Static to Dynamic Auctions

Compared with standard single-stage auctions, the distinguishing feature of two-stage IC auctions is their reliance on entry fees to incentivize truthful reporting of initial valuations. However, entry fees are an inherent feature of static two-stage IC auctions. To resolve this tension between IC and IR, we turn to the design of dynamic auctions. Distinct from classical dynamic auctions, our objective is to leverage historical bidding information to detect and discourage strategic misreporting, thereby reducing the reliance on entry fees while still ensuring approximate two-stage IC.

*5.1.1 More Notations for Dynamic Setting.* To formalize the concepts of dynamic auctions, we focus on the setting of finite-horizon auctions with length $T$. In each period $t \in [T]$, the ad platform provides a new display opportunity to sell to the advertisers, where the values of the opportunities are *i.i.d.* across time for each bidder $i$, following $f_i$ and $g_i$. To adapt the previous notations to a dynamic setting, we use the superscript $t$ to denote the corresponding attribute at time period $t$. For example, $v_i^t$ and $w_i^t$ denote the initial and refined valuation for advertiser $i$ at time $t$ respectively, with $v_i^t \sim f_i$ and $w_i^t \sim g_i(\cdot; v_i^t)$, and $S^t$ denotes the bidder subset at time $t$. With a little abuse of notation, we also use $(x_i^t, p_i^t)$ and $(\mathbf{x}^t, \mathbf{p}^t)$ to denote the realization of allocation and payment for advertiser $i$ and all the advertisers at time period $t$, respectively. In the dynamic auction setting, both the platform and advertisers may rely on the historical information to determine the auction outcomes or make bidding decisions adaptively. We use $\mathcal{H}^{t,0} = \{(\hat{\mathbf{v}}^{t'}, S^{t'}, \hat{\mathbf{w}}^{t'}, \mathbf{x}^{t'}, \mathbf{p}^{t'})\}_{t'<t}$ and $\mathcal{H}^{t,1} = \{\mathcal{H}^{t,0}, \hat{\mathbf{v}}^t, S^t\}$ to denote the public history at the pre-auction stage and formal auction stage of time period $t$, respectively. We also use $\mathcal{T}_i^{t,0} = \{(v_i^{t'}, w_i^{t'})\}_{t'<t}$ and $\mathcal{T}_i^{t,1} = \{\mathcal{T}^{t,0}, v_i^t\}$ to denote the private history types of advertiser $i$ at the pre-auction stage and formal auction stage of time period $t$. We use $\Delta\mathcal{H}$ and $\Delta\mathcal{T}_i$ to denote the space of public history and private history of advertiser $i$.

*Utilities of Advertisers.* A bidding strategy of advertiser $i$ could be represented as $\mathcal{B}_i : \Delta\mathcal{H}\times\Delta\mathcal{T}_i\times[0, \overline{u}] \to [0, \overline{u}]$. That is, the advertiser relies on the public history, her private history, as well as her current valuation to make a bidding decision. We use $\mathcal{I}_i$ to denote the truthful strategy that always report true (initial or refined) valuation in the current stage. Assuming all other advertisers report truthfully, the expected utility of adopting strategy $\mathcal{B}_i$ for advertiser $i$ under auction mechanism $\mathcal{M}$ is defined as

$$\mathcal{U}_i(\mathcal{B}_\mathrm{i}; \mathcal{M}) = \mathbb{E}\left[\sum_{t=1}^{T} w_i^t \cdot x_i^t - p_i^t\right],$$

where the expectation is taken over the realization of advertisers' initial valuations $\mathbf{v}^t$ and refined valuations $\mathbf{w}^t$, the randomness of advertiser $i$'s strategy $\mathcal{B}_i$ in reporting initial bids $\hat{v}_i^t$ and refined bids $\hat{w}_i^t$, and the randomness of mechanism $\mathcal{M}$ in selecting $S^t$ and implementing $(\mathbf{x}^t, \mathbf{p}^t)$. Note that $(x_i^t, p_i^t)$ are random variables depending on the aforementioned randomness.

*IC Properties of Dynamic Two-Stage Auction.* Due to intricate dynamics in repeated auctions, an exact non-trivial dynamic IC auction typically has complex allocation and payment rules [4], leading them inappropriate for ad auctions. Therefore, we turn to consider the *$\epsilon$-approximate IC* on time-averaged utility in an *ex-ante* sense, defined as $\frac{\mathcal{U}_i(\mathcal{B}_i;\mathcal{M})}{T} - \frac{\mathcal{U}_i(\mathcal{I}_i;\mathcal{M})}{T} \le \epsilon$, $\forall i, \mathcal{B}_i$, *i.e.*, any dishonest strategy could not achieve an expected utility increase greater than $\epsilon$ over truthful reporting in each time step. As we will see later, our auction could achieve $\epsilon$-approximate ex-ante IC, with $\epsilon$ converging to 0 with the increase of time period length $T$.

## 5.2 Realization-Dependent Entry Fee

To prevent chain misreporting in $v_i$ and $w_i$, the platform has to charge an entry fee accounting for the overall uncertainty of potential $w_i \sim G_i(\cdot|v_i)$ and $S \sim y_S$ when only $v_i$ is revealed, unavoidably breaking the ex-post IR condition. With the utilization of historical information across time, we will be able to relax the requirement on the entry fee to accounting only for a specific realization of $\hat{w}_i$ and $S$. To distinguish with the entry fee $\tau_i(\hat{\mathbf{v}})$ and $\tau_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}_{-i}; S)$ in static auctions, we use $\hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ to denote such *realization-dependent* entry fee rules in dynamic environment. Before presenting our design to realize this relaxation (Section 5.3), we first demonstrate the condition of $\hat{\tau}_i$ to be feasible for our design, and provide two detailed ex-post IR entry fee rules as candidates for implementation.

Formally speaking, we can design and implement any realization-dependent, non-negative entry fee rules $\hat{\tau}$ to be approximate two-stage IC, as long as it satisfies the *aligning expectation* condition

$$\mathbb{E}_{\hat{\mathbf{w}} \sim \mathbf{g}(\cdot;\hat{\mathbf{v}}), S \sim y_S(\hat{\mathbf{v}})} \left[ \hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) \right] = \tau_i(\hat{\mathbf{v}}). \tag{7}$$

This aligning expectation condition arises as a natural relaxation of the static entry fee $\tau_i(\hat{\mathbf{v}})$. Since $\tau_i(\hat{\mathbf{v}})$ guarantees interim IR, once we relax it to be realization-dependent, finding such an ex-post IR entry fee rule for every realization should be flexible in principle.

A simple yet effective way is to charge entry fee *proportional* to the realized second-stage utility, *i.e.*, $\hat{w}_i \cdot x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) - p_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$, with a constant proportion determined by the static entry fee $\tau_i(\hat{\mathbf{v}})$ and her expected utility $U_i(\hat{\mathbf{v}}; \hat{v}_i)$ in the static auction. Specifically, for the allocation with rank score setting, a convenient approach to charge entry fee is to leverage the classical concept of *quantile transformation* in mechanism design [5, 17]. Given the c.d.f. of refined value $G_i(\cdot; v_i)$, its quantile function $\phi_i(\cdot; v_i) : [0,1] \to [0, \overline{u}]$ is defined as $\phi_i(q; v_i) = H^{-1}(q)$, with $H(\cdot) = G_i(\cdot; v_i)$. That is, $\phi_i$ maps a quantile in $[0,1]$ to the corresponding realization of $w_i$ in the relative location of distribution $G_i(\cdot; v_i)$. As the entry fee in Equation (13) charge the expected utility obtained by threshold initial value $r_{i,0}(\hat{\mathbf{v}})$, we can simulate the virtual $\tilde{w}_i$ induced by $r_{i,0}(\hat{\mathbf{v}})$ utilizing the quantile of realized $\hat{w}_i$ induced by $\hat{v}_i$, *i.e.*, charging

$$\hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \Big( \phi_i\big(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\big) - r_{i,1}(\hat{\mathbf{w}}_{-i}; S) \Big)^+. \tag{8}$$

On the other hand, the standard single-stage payment here is $p_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = r_{i,1}(\hat{\mathbf{w}}_{-i}; S) \cdot x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$. Recall $x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \mathbb{I}\{\hat{w}_i \geq r_{i,1}(\hat{\mathbf{w}}_{-i}; S)\}$, this gives the overall payment $\hat{p}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ as

$$x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) \cdot \max\left\{ r_{i,1}(\hat{\mathbf{w}}_{-i}; S), \phi_i\big(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\big) \right\}.$$

Since $\phi_i\big(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\big)$ and $\hat{w}_i$ has the same quantile in distribution $G_i(\cdot; r_{i,0}(\hat{\mathbf{v}}_{-i}))$ and $G_i(\cdot; \hat{v}_i)$ respectively, we also have $\phi_i\big(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\big) \leq \hat{w}_i$. As $r_{i,1}(\hat{\mathbf{w}}_{-i}; S)$ represents the standard single-stage winning payment, this format aligns with the established *reserve price* scheme in ad auctions [30, 35], despite the dependence of this new "reserve price" $\phi_i\big(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\big)$ on $\hat{w}_i$, providing a convenient approach for implementation.

**Algorithm 1:** Two-Stage Auction with Discounts

**Input:** The bids profile $(\hat{\mathbf{v}}^t, \hat{\mathbf{w}}^t)_{t \in [T]}$
**Parameter:** A two-stage monotone allocation rule $(y_S, x)$, a non-negative, expectation-aligning and ex-post IR entry fee rule $\hat{\tau}$, discount proportion $k_i$ and tolerance parameter $\epsilon_i$
**Output:** Auction outcomes $(\mathbf{x}^t, \mathbf{p}^t)_{t \in [T]}$

Compute $\overline{\tau_i} = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\tau_i(\mathbf{v})], \forall i$;
Initialize $L_i = k_i T \cdot \overline{\tau_i} + \epsilon_i$ and $d_i = 0$ for each bidder $i$;
Initialize the active bidder set $A^0 = [n]$;
**for** $t = 1, 2, \ldots, T$ **do**
  Set $A^t = A^{t-1}$;
  Get the initial bids profile $\hat{\mathbf{v}}^t$ from bidders;
  For each bidder $i \notin A^{t-1}$, sample $\hat{v}_i^t \sim f_i$ to replace her original $\hat{v}_i^t$;
  Determine the entering bidder subset by $S^t \sim y_S(\hat{\mathbf{v}}^t)$;
  Get the refined bids profile $\hat{\mathbf{w}}^t$ from bidders;
  For each bidder $i \notin A^{t-1}$ and $i \in S^t$, sample $\hat{w}_i^t \sim g_i(\cdot; \hat{v}_i^t)$ to replace her original $\hat{w}_i^t$;
  **for** *each bidder* $i \in A^{t-1}$ **do**
    Update $d_i = d_i + \tau_i(\hat{\mathbf{v}}^t) - (1 - k_i) \cdot \hat{\tau}_i(\hat{\mathbf{v}}^t, \hat{\mathbf{w}}^t; S^t)$;
    Compute $p_{i,1}^t = \int_0^{\hat{w}_i^t} \tilde{w}_i \cdot \frac{\partial x_i(\hat{\mathbf{v}}^t, \tilde{w}_i, \hat{\mathbf{w}}_{-i}^t; S^t)}{\partial \tilde{w}_i} d\tilde{w}_i$;
    Assign $p_i^t = (1 - k_i) \cdot \hat{\tau}_i(\hat{\mathbf{v}}^t, \hat{\mathbf{w}}^t; S^t) + p_{i,1}^t$;
    **if** $d_i > L_i$ **then**
      Update $A^t = A^t \backslash \{i\}$;
  Assign $\mathbf{x}^t = \big(x_i(\hat{\mathbf{v}}^t, \hat{\mathbf{w}}^t; S^t)\big)_{i \in [n]}$, and replace $x_i^t$ by $x_i(0, \hat{\mathbf{v}}_{-i}^t, 0, \hat{\mathbf{w}}_{-i}^t; S^t)$ for $i \notin A^{t-1}$;
  Assign $p_i^t = 0$ for $i \notin A^{t-1}$;
  Implement $(\mathbf{x}^t, \mathbf{p}^t)$ for period $t$.

Theorem 5.1. *Given any two-stage monotone allocation rule $(y_S, x)$, the following realization-dependent entry fee rules satisfy the aligning expectation and ex-post IR condition.*

- *(Proportional Entry Fee) For general setting, let $\eta = \frac{\tau_i(\hat{\mathbf{v}})}{U_i(\hat{\mathbf{v}}; \hat{v}_i) + \tau_i(\hat{\mathbf{v}})}$, charge entry fee $\hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \eta \cdot \big(\hat{w}_i \cdot x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) - p_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)\big)$, where $U_i(\hat{\mathbf{v}}; \hat{v}_i)$, $\tau_i(\hat{\mathbf{v}})$ are defined by $(y_S, x)$ as in Theorem 4.1.*
- *(Quantile-Based Entry Fee) In the allocation with rank scores setting, charge the entry fee $\hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ according to Equation (8)*[3] *for $i \in S$, and charge 0 for $i \notin S$.*

## 5.3 Auction Design with Discounts in Entry Fees

*5.3.1 Discounts in Entry Fees.* When promoting two-stage auctions, even if ex-post IR is guaranteed under a realization-dependent entry fee, advertisers may take issue with incurring an additional expense of entry fees, despite its necessity to elicit true values and boost social welfare in static two-stage auctions. This motivates us to provide discounts on entry fees based on the platform's choices. Since an ad platform may prefer to strike a balance between advertiser satisfaction and platform revenue, we introduce

[3]The formula assumes atomless conditional distributions. When atoms exist, the same argument can be recovered by using standard randomized quantile within each atom.

a parameter $k_i$ to regulate the proportion of discounts offered to advertiser $i$. With $k_i = 0$, no discount is provided, and the auction charges the full (realization-dependent) entry fee defined by the platform; conversely, when $k_i = 1$, a full discount is provided to advertiser $i$, such that she does not pay any entry fee. The theoretical guarantee of our design, *i.e.*, ex-post IR and approximate two-stage IC, holds for arbitrary $k_i \in [0, 1]$, including (1) charging the full realization-dependent entry fee ($k_i = 0$) by, for example, setting quantile-based reserve price in allocation with rank score (Theorem 5.1); and (2) eliminating entry fees ($k_i = 1$) and preserving the standard single-stage payment format.

*5.3.2 Auction Procedure and Analysis.* The design of our auction is inspired by the non-monetary artificial currency mechanism [14]. The formal process of our auction is presented in Algorithm 1. The platform still implements the two-stage allocation rule in each time period (Line 8 and 17). Besides the standard single-stage payment $p_{i,1}$ (Line 13), we also charge the realization-dependent entry fee scaled by the discount parameter, *i.e.*, $(1 - k_i) \cdot \hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S^t)$ (Line 14). As we deviate from the static entry fee $\tau_i(\hat{\mathbf{v}})$, the original IC guarantee (Theorem 4.1) for each time period no longer holds. Since the entry fee now depends on reported $\hat{w}_i$, bidders can deliberately report a lower refined bid to reduce the entry fee in $\hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S^t)$. Even if we adopt $\hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S^t) = \tau_i(\hat{\mathbf{v}})$, the provided discount $k_i$ still leaves room for bidders to report higher initial bids $\hat{v}_i$.

To avoid advertisers exploiting these beneficial designs, we endow each advertiser with a discount upper limit $L_i$, such that an advertiser would be excluded from the auctions when her consumed discounts, denoted by $d_i$, exceed $L_i$ (Line 15-16). We compute $d_i$ as the difference between the static two-stage IC entry fee should have been charged, *i.e.*, $\tau_i(\hat{\mathbf{v}})$, and the actual amount of charged entry fees, *i.e.*, $(1-k_i) \cdot \hat{\tau}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S^t)$ (Line 12). Correspondingly, we set the upper limit $L_i$ according to the expectation of consumed discounts under prior distribution (Line 1-2), where we define $\overline{\tau_i} = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}}[\tau_i(\mathbf{v})]$ according to Equation (4), and $\epsilon_i$ is a parameter controlling the tolerance of out-of-distribution realization. As a technical detail, we replace the bids of excluded advertisers with sampled ones to preserve properties of the static auction in analysis (Line 7 and 10).

The intuition behind setting a discount upper limit $L_i$ is as follows. For a truthful bidder $i$, her expected allocation under Algorithm 1 remains the same as in the static two-stage auction, provided that $L_i$ is not exhausted. As the expected total discount consumed by a truthful strategy $\mathcal{I}_i$ over $T$ rounds is $T \cdot k_i \overline{\tau_i}$ (by aligning expectation condition of $\hat{\tau}_i$), the probability that this consumption exceeds $L_i$ is upper bounded by Hoeffding inequality, thus limiting the risk of mistakenly exclude truthful bidders. Conversely, the additional discounts a dishonest advertiser can exploit are bounded by a tolerance parameter $\epsilon_i$. This setup introduces a tradeoff: selecting $\epsilon_i$ involves balancing the upper bound of dishonest utility increment against the risk of false penalties for truthful bidders. A larger $\epsilon_i$ reduces false exclusion risks but weakens approximate IC.

Theorem 5.2. *For any auction mechanism $\mathcal{M}_d$ defined by Algorithm 1, supposing all the other advertisers report truthfully, the time-averaged utility increment of any potential strategy $\mathcal{B}_i$ compared with the truthful strategy $\mathcal{I}_i$ is upper bounded by $2\overline{u}\sqrt{\frac{\ln T}{T}} + \frac{2\overline{u}}{T}$ under choice of $\epsilon_i = 2\overline{u}\sqrt{T \ln T}$. Moreover, the probability that a truthful strategy would lead advertiser $i$ to be mistakenly excluded from auctions,* i.e.*, $i \notin A^T$, is upper bounded by $T^{-1}$.*

Theorem 5.2 implies our auction satisfies an $\epsilon$-approximate ex-ante IC condition with $\epsilon \to 0$ as $T \to \infty$. To implement Algorithm 1, the platform only needs to track each bidder's discount consumption $d_i$ and exclude them if $d_i$ exceeds $L_i$. Since the computation of discount consumption is a simple arithmetic based on static and realization-dependent entry fees (Line 12), its implementation is straightforward. In the case of full discounting ($k_i = 1$), initial bids serve solely as a reference for the platform to select the entering bidder subset, yet our design of $L_i$ ensures incentive alignment and prevents overbidding in the pre-auction stage.

## 6 Related Work

As the auction design for the formal auctions has been extensively studied, previous work on two-stage ad auctions focuses on designing the bidder subset selection rule with two major information acquisition settings. In the first setting, the information augmented in the formal auction is the refined estimation of CTR, while advertisers retain the same private valuations across two stages. Both deep learning methods [36, 40] and theory-based method [12, 13] have been proposed to select a bidder subset with higher expected social welfare, while guaranteeing the IC conditions. An online learning variant has also been studied [24]. In the second setting, advertisers learn their valuations only after entering the formal auction stage, and the platform needs to select the bidder subset based on their prior distributions [3, 15]. In contrast, we consider a scenario where advertisers have different valuations across two stages, which better outlines the change in advertiser valuations.

While economics literature has studied similar settings where bidders could learn additional signals to refine their valuations in a two-stage auction [11, 38], our focus and characterization results are distinct. The line of work on indicative bidding [22, 32, 38] examines this setting due to the high cost incurred by learning the value of an item in high-value asset sales. They assume the second stage of the auction to be fixed as standard auctions (*e.g.*, second-price auction), and mainly focus on the equilibrium properties for different auction formats in the first stage. Another related line of work investigates how the auctioneer could maximize profits by disclosing additional signals to bidders and enabling them to update valuations [11, 25, 27]. Despite the structural similarities between their expected revenue results and our Theorem 4.3, they concentrate on the specific revenue-maximizing auction, while our results in Section 4 provide a general reduction to construct two-stage IC auctions with a different IC concept for online ad auctions.

Our objective of reducing entry fees induced by IC conditions aligns with the broader principles of non-monetary mechanism design [2, 8, 14, 16, 21]. The phenomenon of impossibility to achieve ex-post IR is also discovered in a cost-per-action auction setting [34]. In particular, we adapt the idea of the artificial currency mechanism [14] to provide discounts in entry fees of two-stage auctions.

## 7 Conclusion

In this work, we consider an auction model in which bidder valuations change across two stages. Motivated by industry practice in online advertising, we study the design of two-stage incentive-compatible auctions that allow bidders to submit and refine their bids. We characterize feasible payment rules to truthfully implement any two-stage monotone allocation rule, and analyze its revenue and IR properties. As the payment in two-stage IC auctions features an additional entry fee, general two-stage IC auctions fail to satisfy ex-post IR. Towards this limitation, we propose a dynamic two-stage auction with flexible discounts in entry fees. Our mechanism allows realization-dependent entry fee, and can achieve approximate ex-ante two-stage IC and ex-post IR with the payment format of single-stage auctions when full discount is applied.

This work presents the formal game-theoretical analysis of two-stage auctions with bid refinement in online advertising. Our findings highlight the importance of entry fees in preventing strategic misreporting, and demonstrate the feasibility of using discount upper bounds as an alternative to regulate advertiser behavior in bid refinement. Moving forward, we aim to leverage these theoretical insights to bridge the gap towards the practical deployment of two-stage auctions in large-scale ad systems.

## Acknowledgement

This work was supported 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] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. 2019. Autobidding with constraints. In *Proceedings of the 15th International Conference on Web and Internet Economics*. Springer, 17–30.
[2] Santiago R Balseiro, Huseyin Gurkan, and Peng Sun. 2019. Multiagent mechanism design without money. *Operations Research* 67, 5 (2019), 1417–1436.
[3] Xiaohui Bei, Nick Gravin, Pinyan Lu, and Zhihao Gavin Tang. 2023. Bidder subset selection problem in auction design. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 3788–3801.
[4] Dirk Bergemann and Juuso Välimäki. 2019. Dynamic mechanism design: An introduction. *Journal of Economic Literature* 57, 2 (2019), 235–274.
[5] Jeremy Bulow and Paul Klemperer. 1996. Auctions versus Negotiations. *American Economic Review* 86, 1 (1996), 180–194.
[6] Leng Cai, Junxuan He, Yikai Li, Junjie Liang, Yuanping Lin, Ziming Quan, Yawen Zeng, and Jin Xu. 2025. RTBAgent: A LLM-based agent system for real-time bidding. In *Companion Proceedings of the ACM on Web Conference 2025*. 104–113.
[7] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for Youtube Recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.
[8] Yan Dai, Moïse Blanchard, and Patrick Jaillet. 2025. Non-monetary mechanism design without distributional information: Using scarce audits wisely. *arXiv preprint arXiv:2502.08412* (2025).
[9] Darinka Dentcheva and Andrzej Ruszczynski. 2003. Optimization with stochastic dominance constraints. *SIAM Journal on Optimization* 14, 2 (2003), 548–566.
[10] Chantat Eksombatchai, Pranav Jindal, Jerry Zitao Liu, Yuchen Liu, Rahul Sharma, Charles Sugnet, Mark Ulrich, and Jure Leskovec. 2018. Pixie: A system for recommending 3+ billion items to 200+ million users in real-time. In *Proceedings of the 2018 World Wide Web Conference*. 1775–1784.
[11] Péter Eső and Balazs Szentes. 2007. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies* 74, 3 (2007), 705–731.
[12] Zhikang Fan, Lan Hu, Ruirui Wang, Zhongrui Ma, Yue Wang, Ye Qi, and Weiran Shen. 2025. Two-stage auction design in online advertising. In *Proceedings of the ACM Web Conference 2025*.
[13] Gagan Goel, Renato Paes Leme, Jon Schneider, David Thompson, and Hanrui Zhang. 2023. Eligibility mechanisms: Auctions meet information retrieval. In *Proceedings of the ACM Web Conference 2023*. 3541–3549.
[14] Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. 2021. From monetary to nonmonetary mechanism design via artificial currencies. *Mathematics of Operations Research* 46, 3 (2021), 835–855.
[15] Nikolai Gravin, Yixuan Even Xu, and Renfei Zhou. 2024. Bidder selection problem in position auctions: A fast and simple algorithm via poisson approximation. In *Proceedings of the ACM Web Conference 2024*. 89–98.
[16] Mingyu Guo, Vincent Conitzer, and Daniel M Reeves. 2009. Competitive repeated allocation without payments. In *International Workshop on Internet and Network Economics*. Springer, 244–255.
[17] Jason D Hartline. 2013. *Mechanism design and approximation*. Chapter 3.
[18] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, and Joaquin Quiñonero Candela. 2014. Practical lessons from predicting clicks on ads at facebook. In *Proceedings of the 8th International Workshop on Data Mining for Online Advertising*. 1–9.
[19] Yue He, Xiujun Chen, Di Wu, Junwei Pan, Qing Tan, Chuan Yu, Jian Xu, and Xiaoqiang Zhu. 2021. A unified solution to constrained bidding in online display advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM, 2993–3001.
[20] Jiri Hron, Karl Krauth, Michael Jordan, and Niki Kilbertus. 2021. On component interactions in two-stage recommender systems. *Proceedings of the 35th International Conference on Neural Information Processing Systems* 34 (2021), 2744–2757.
[21] Matthew O Jackson and Hugo F Sonnenschein. 2007. Overcoming incentive constraints by linking decisions 1. *Econometrica* 75, 1 (2007), 241–257.
[22] John Kagel, Svetlana Pevnitskaya, and Lixin Ye. 2008. Indicative bidding: An experimental analysis. *Games and Economic Behavior* 62, 2 (2008), 697–721.
[23] 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.
[24] Haoming Li, Yumou Liu, Zhenzhe Zheng, Zhilin Zhang, Jian Xu, and Fan Wu. 2024. Truthful bandit mechanisms for repeated two-stage ad auctions. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1565–1575.
[25] Hao Li and Xianwen Shi. 2017. Discriminatory Information Disclosure. *American Economic Review* 107, 11 (November 2017), 3363–85.
[26] 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.
[27] Jingfeng Lu, Lixin Ye, and Xin Feng. 2021. Orchestrating information acquisition. *American Economic Journal: Microeconomics* 13, 4 (2021), 420–465.
[28] Xiao Ma, Liqin Zhao, Guan Huang, Zhi Wang, Zelin Hu, Xiaoqiang Zhu, and Kun Gai. 2018. Entire space multi-task model: An effective approach for estimating post-click conversion rate. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 1137–1140.
[29] Roger B Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
[30] Michael Ostrovsky and Michael Schwarz. 2023. Reserve prices in internet advertising auctions: A field experiment. *Journal of Political Economy* 131, 12 (2023), 3352–3376.
[31] Weitong Ou, Bo Chen, Xinyi Dai, Weinan Zhang, Weiwen Liu, Ruiming Tang, and Yong Yu. 2023. A survey on bid optimization in real-time bidding display advertising. *ACM Transactions on Knowledge Discovery from Data* 18, 3 (2023), 1–31.
[32] Daniel Quint and Kenneth Hendricks. 2018. A theory of indicative bidding. *American Economic Journal: Microeconomics* 10, 2 (2018), 118–151.
[33] Ben Roberts, Dinan Gunawardena, Ian A Kash, and Peter Key. 2016. Ranking and tradeoffs in sponsored search auctions. *ACM Transactions on Economics and Computation (TEAC)* 4, 3 (2016), 1–21.
[34] Weiran Shen, Zihe Wang, and Song Zuo. 2018. Ex-post IR dynamic auctions with cost-per-action payments. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 505–511.
[35] David RM Thompson and Kevin Leyton-Brown. 2013. Revenue optimization in the generalized second-price auction. In *Proceedings of the 14th ACM Conference on Electronic Commerce*. 837–852.
[36] 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.
[37] Xun Yang, Yasong Li, Hao Wang, Di Wu, Qing Tan, Jian Xu, and Kun Gai. 2019. Bid optimization by multivariable control in display Advertising. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1966–1974.
[38] Lixin Ye. 2007. Indicative bidding and a theory of two-stage auctions. *Games and Economic Behavior* 58, 1 (2007), 181–207.
[39] Weinan Zhang, Shuai Yuan, and Jun Wang. 2014. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1077–1086.
[40] Zhishan Zhao, Jingyue Gao, Yu Zhang, Shuguang Han, Siyuan Lou, Xiang-Rong Sheng, Zhe Wang, Han Zhu, Yuning Jiang, Jian Xu, et al. 2023. COPR: Consistency-oriented pre-Ranking for online advertising. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4974–4980.

[41] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1059–1068.

# A Missing Proofs

## A.1 Proof of Theorem 4.1

*Part 1. Characterization of Payment Rule.* Fixing $\hat{\mathbf{v}}$ and $S$, the ex-post IC condition for the formal auction stage is equivalent to the classical IC conditions for single-item auctions. By Myerson's seminal result [29], following conditions are necessary and sufficient to guarantee the IC conditions for the formal auction stage: $\forall \hat{\mathbf{v}}, S$, (1) $x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ is non-decreasing in $\hat{w}_i$; and (2) the payment satisfies $p_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) = \int_0^{\hat{w}_i} \tilde{w}_i \cdot \frac{\partial x_i(\hat{\mathbf{v}}, \tilde{w}_i, \hat{\mathbf{w}}_{-i}; S)}{\partial \tilde{w}_i} \,\mathrm{d}\tilde{w}_i + p_i(\hat{\mathbf{v}}, 0, \hat{\mathbf{w}}_{-i}; S)$, with $u_i(\hat{\mathbf{v}}, \mathbf{w}; w_i, S) = \int_0^{w_i} x_i(\hat{\mathbf{v}}, \tilde{w}_i, \mathbf{w}_{-i}; S)\,\mathrm{d}\tilde{w}_i - p_i(\hat{\mathbf{v}}, 0, \mathbf{w}_{-i}; S)$.

Then, the expected utility $U_i(\hat{\mathbf{v}}; v_i)$ equals

$$\mathbb{E}_{\mathbf{w}\sim\mathbf{g}(\cdot; v_i, \hat{\mathbf{v}}_{-i})}\left[\int_0^{w_i} z_i(\hat{\mathbf{v}}, \tilde{w}_i, \mathbf{w}_{-i})\,\mathrm{d}\tilde{w}_i - q_i(\hat{\mathbf{v}}, 0, \mathbf{w}_{-i})\right].$$

To satisfy the interim IC conditions in the pre-auction stage for advertiser $i$, considering the true initial valuation to be $v_i$ and the strategy to misreport $\hat{v}_i$, we require

$$\begin{aligned}&\mathbb{E}_{\mathbf{w}\sim\mathbf{g}(\cdot; v_i, \hat{\mathbf{v}}_{-i})}\left[\int_0^{w_i} z_i(v_i, \hat{\mathbf{v}}_{-i}, \tilde{w}_i, \mathbf{w}_{-i}) - z_i(\hat{\mathbf{v}}, \tilde{w}_i, \mathbf{w}_{-i})\,\mathrm{d}\tilde{w}_i\right] \\ &\geq \mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\left[q_i(v_i, \hat{\mathbf{v}}_{-i}, 0, \mathbf{w}_{-i}) - q_i(\hat{\mathbf{v}}, 0, \mathbf{w}_{-i})\right].\end{aligned} \tag{9}$$

On the other hand, considering the true initial valuation to be $\hat{v}_i$ and the strategy to misreport $v_i$, we require

$$\begin{aligned}&\mathbb{E}_{\mathbf{w}\sim\mathbf{g}(\cdot;\hat{\mathbf{v}})}\left[\int_0^{w_i} z_i(v_i, \hat{\mathbf{v}}_{-i}, \tilde{w}_i, \mathbf{w}_{-i}) - z_i(\hat{\mathbf{v}}, \tilde{w}_i, \mathbf{w}_{-i})\,\mathrm{d}\tilde{w}_i\right] \\ &\leq \mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\left[q_i(v_i, \hat{\mathbf{v}}_{-i}, 0, \mathbf{w}_{-i}) - q_i(\hat{\mathbf{v}}, 0, \mathbf{w}_{-i})\right].\end{aligned} \tag{10}$$

Combining them and taking $\hat{v}_i \to v_i^+$, we have

$$\begin{aligned}&\mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\left[\frac{\partial q_i(v_i, \hat{\mathbf{v}}_{-i}, 0, \mathbf{w}_{-i})}{\partial v_i}\right] \\ &= \mathbb{E}_{\mathbf{w}\sim\mathbf{g}(\cdot; v_i, \hat{\mathbf{v}}_{-i})}\left[\int_0^{w_i} \frac{\partial z_i(v_i, \hat{\mathbf{v}}_{-i}, \tilde{w}_i, \mathbf{w}_{-i})}{\partial v_i}\,\mathrm{d}\tilde{w}_i\right].\end{aligned}$$

Considering that $q_i(0, \hat{\mathbf{v}}_{-i}, 0, \mathbf{w}_{-i}) = 0$ and integrating over $v_i$, we then have

$$\begin{aligned}&\mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\left[q_i(v_i, \hat{\mathbf{v}}_{-i}, 0, \mathbf{w}_{-i})\right] \\ &= \int_0^{v_i} \mathbb{E}_{\mathbf{w}\sim\mathbf{g}(\cdot;\tilde{v}_i, \hat{\mathbf{v}}_{-i})}\left[\int_0^{w_i} \frac{\partial z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, \tilde{w}_i, \mathbf{w}_{-i})}{\partial \tilde{v}_i}\,\mathrm{d}\tilde{w}_i\right]\mathrm{d}\tilde{v}_i \\ &= \mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\left[\int_0^{v_i}\int_0^{\overline{u}} g_i(w_i;\tilde{v}_i) \int_0^{w_i} \frac{\partial z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, \tilde{w}_i, \mathbf{w}_{-i})}{\partial \tilde{v}_i}\,\mathrm{d}\tilde{w}_i\,\mathrm{d}w_i\,\mathrm{d}\tilde{v}_i\right],\end{aligned} \tag{11}$$

which could be further simplified to the expression in the statement of Theorem 3. Here, we utilize the transformation

$$\int_0^{\overline{u}} g_i(w_i; v_i)\int_0^{w_i} h(\tilde{w}_i)\,\mathrm{d}\tilde{w}_i\,\mathrm{d}w_i = \int_0^{\overline{u}} (1 - G(w_i; v_i)) h(w_i)\,\mathrm{d}w_i$$

for c.d.f. $G(\cdot; v_i)$. Since an advertiser evaluates the auction outcomes from an interim perspective in the pre-auction stage, any design of $q_i(v_i, \hat{\mathbf{v}}_{-i}, 0, \mathbf{w}_{-i})$ satisfying Equation (11) are equivalent for two-stage IC conditions. Our payment rule in Theorem 4.1 assigns the payment $p_i(v_i, \hat{\mathbf{v}}_{-i}, 0, \mathbf{w}_{-i}; S)$ to be the same for any $S$ and $\mathbf{w}_{-i}$.

*Part 2. Investigate Feasible Allocation Rule.* Substituting the payment rule back to Inequality (9), we require

$$\begin{aligned}&\mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\Big[\int_0^{\overline{u}} (1 - G(w_i; v_i))\cdot \\ &\qquad (z_i(v_i, \hat{\mathbf{v}}_{-i}, w_i, \mathbf{w}_{-i}) - z_i(\hat{\mathbf{v}}, w_i, \mathbf{w}_{-i}))\,\mathrm{d}w_i\Big] \\ &\geq \mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\left[\int_{\hat{v}_i}^{v_i}\int_0^{\overline{u}} (1 - G_i(w_i;\tilde{v}_i))\cdot\frac{\partial z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, w_i, \mathbf{w}_{-i})}{\partial \tilde{v}_i}\,\mathrm{d}w_i\,\mathrm{d}\tilde{v}_i\right].\end{aligned} \tag{12}$$

We would now demonstrate the allocation rule with $\frac{\partial z_i(\mathbf{v},\mathbf{w})}{\partial v_i} \geq 0, \forall \mathbf{v}, \mathbf{w}$, could satisfy the interim IC conditions in the pre-auction stage. The statement holds if we could show

$$\begin{aligned}&\int_0^{\overline{u}} (1 - G(w_i; v_i))\cdot(z_i(v_i, \hat{\mathbf{v}}_{-i}, w_i, \mathbf{w}_{-i}) - z_i(\hat{\mathbf{v}}, w_i, \mathbf{w}_{-i}))\,\mathrm{d}w_i \\ &\geq \int_{\hat{v}_i}^{v_i}\int_0^{\overline{u}} (1 - G_i(w_i;\tilde{v}_i))\cdot\frac{\partial z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, w_i, \mathbf{w}_{-i})}{\partial \tilde{v}_i}\,\mathrm{d}w_i\,\mathrm{d}\tilde{v}_i\end{aligned}$$

for any $v_i$, $\hat{\mathbf{v}}$ and $\mathbf{w}_{-i}$. By first-order stochastic dominance of $G_i$, if $\hat{v}_i < v_i$, then $G_i(w_i;\hat{v}_i) \geq G_i(w_i;\tilde{v}_i) \geq G_i(w_i; v_i)$ for $\tilde{v}_i \in [\hat{v}_i, v_i]$ given fixed $w_i$. Therefore, applying $1 - G_i(w_i;\tilde{v}_i) \leq 1 - G_i(w_i; v_i)$ proves the statement. The case for $\hat{v}_i > v_i$ holds similarly.

*Part 3. Verification of Interim IR.* With expectation of terms $q_i(\hat{\mathbf{v}}, 0, \mathbf{w}_{-i})$ specified in Equation (11), $U_i(v_i, \hat{\mathbf{v}}_{-i}; v_i)$ could be represented as

$$\begin{aligned}&\mathbb{E}_{\mathbf{w}_{-i}\sim\mathbf{g}_{-i}(\cdot;\hat{\mathbf{v}}_{-i})}\Big[\int_0^{\overline{u}} (1 - G(w_i; v_i))\cdot z_i(v_i, \hat{\mathbf{v}}_{-i}, \mathbf{w})\,\mathrm{d}w_i \\ &\qquad - \int_0^{v_i}\int_0^{\overline{u}} (1 - G_i(w_i;\tilde{v}_i))\cdot\frac{\partial z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, \mathbf{w})}{\partial \tilde{v}_i}\,\mathrm{d}w_i\,\mathrm{d}\tilde{v}_i\Big],\end{aligned}$$

which is non-negative by the monotonicity of $G(w_i;\cdot)$.

## A.2 Proof of Proposition 4.2

Proof. We apply the Lebesgue–Stieltjes version of Theorem 4.1.

**Case 1:** $i \notin S^{(k)}(\hat{\mathbf{v}})$, *i.e.*, advertiser $i$ has $h_{i,0}(\hat{v}_i) < h^{(k)}_{-i,0}(\hat{\mathbf{v}}_{-i})$ and does not enter the formal auction stage. By the non-decreasing property of $h_{i,0}$, the advertiser has no chance to enter the formal auction with any $\tilde{v}_i \leq \hat{v}_i$; therefore, $z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, w_i, \mathbf{w}_{-i})$ always keep 0 for $\tilde{v}_i \leq \hat{v}_i$, leading Equation (4) and total payment to be 0.

**Case 2:** $i \in S^{(k)}(\hat{\mathbf{v}})$, *i.e.*, advertiser $i$ has $h_{i,0}(\hat{v}_i) \geq h^{(k)}_{-i,0}(\hat{\mathbf{v}}_{-i})$ and enters the formal auction stage. As the rank scores in two stages are independent, we can omit the influence of $\hat{v}_i$ on $x_i$ given $S^{(k)}$. As $z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, w_i, \mathbf{w}_{-i}) = \sum_{S\subseteq[n]} y_S(\tilde{v}_i, \hat{\mathbf{v}}_{-i})\cdot x_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, \mathbf{w}; S)$, which may jump from 0 to $\mathbb{I}\{w_i \geq r_{i,1}(\mathbf{w}_{-i}; S)\}$ only when $\tilde{v}_i = r_{i,0}(\hat{\mathbf{v}}_{-i})$ for $S = S^{(k)}(\hat{\mathbf{v}})$, $\tau_i(\hat{\mathbf{v}})$ in Equation (4) is equal to the expectation of

$$\begin{aligned}&\int_{r_{i,1}(\hat{\mathbf{w}}_{-i}; S^{(k)}(\hat{\mathbf{v}}))}^{\overline{u}} \left(1 - G_i\big(w_i; r_{i,0}(\hat{\mathbf{v}}_{-i})\big)\right)\mathrm{d}w_i \\ &= \mathop{\mathbb{E}}_{\tilde{w}_i\sim g_i(\cdot\mid r_{i,0}(\hat{\mathbf{v}}_{-i}))}\left[\left(\tilde{w}_i - r_{i,1}\big(\mathbf{w}_{-i}; S^{(k)}(\hat{\mathbf{v}})\big)\right)^+\right],\end{aligned} \tag{13}$$

over the randomness of $\mathbf{w}_{-i} \sim \mathbf{g}(\cdot;\hat{\mathbf{v}}_{-i})$. As $y_S$ is deterministic, our entry fee rule is equivalent to $\tau_i(\hat{\mathbf{v}})$ in expectation. □

## A.3 Proof of Theorem 4.3

Proof. The first part of revenue on the entry fee from $i$ is

$$\begin{aligned}
&\int_{\mathbf{v}} \mathbf{f}(\mathbf{v}) \int_{\mathbf{w}} \mathbf{g}(\mathbf{w};\mathbf{v}) \cdot q_i(\mathbf{v}, 0, \mathbf{w}_{-i})\, \mathrm{d}\mathbf{w}\, \mathrm{d}\mathbf{v} \\
=& \int_{\mathbf{v}} \mathbf{f}(\mathbf{v}) \int_{\mathbf{w}_{-i}} \mathbf{g}_{-i}(\mathbf{w}_{-i}; \mathbf{v}_{-i}) \int_0^{v_i} \int_0^{+\infty} (1 - G_i(w_i; \tilde{v}_i))\cdot \\
&\qquad \frac{\partial z_i(\tilde{v}_i, \mathbf{v}_{-i}, w_i, \mathbf{w}_{-i})}{\partial \tilde{v}_i} \mathrm{d}w_i\, \mathrm{d}\tilde{v}_i\, \mathrm{d}\mathbf{w}_{-i}\, \mathrm{d}\mathbf{v} \qquad (14)\\
=& \int_{\mathbf{v}_{-i}} \mathbf{f}_{-i}(\mathbf{v}_{-i}) \int_{\mathbf{w}_{-i}} \mathbf{g}_{-i}(\mathbf{w}_{-i}; \mathbf{v}_{-i}) \int_0^{+\infty} \int_0^{+\infty} (1 - F_i(v_i))\cdot \\
&(1 - G_i(w_i; v_i)) \cdot \frac{\partial z_i(v_i, \mathbf{v}_{-i}, w_i, \mathbf{w}_{-i})}{\partial v_i} \mathrm{d}w_i\, \mathrm{d}v_i\, \mathrm{d}\mathbf{w}_{-i}\, \mathrm{d}\mathbf{v}_{-i}.
\end{aligned}$$

Denote the c.d.f. of variable $\epsilon_i$ as $H_i$. By the additive noise assumption, we have $G_i(w_i; v_i) = H_i(w_i - v_i)$, which indicates $\frac{\partial G_i(w_i;v_i)}{\partial v_i} = -g_i(w_i; v_i)$. Applying integration by parts, we have

$$\begin{aligned}
&\int_0^{+\infty} (1 - F_i(v_i)) \cdot (1 - G_i(w_i; v_i)) \cdot \frac{\partial z_i(v_i, \mathbf{v}_{-i}, w_i, \mathbf{w}_{-i})}{\partial v_i} \mathrm{d}v_i \\
=&(1 - F_i(v_i)) \cdot (1 - G_i(w_i; v_i)) \cdot z_i(v_i, \mathbf{v}_{-i}, w_i, \mathbf{w}_{-i})|_{v_i=0}^{+\infty} \\
&- \int_0^{+\infty} [g_i(w_i; v_i) \cdot (1 - F_i(v_i)) - f_i(v_i) \cdot (1 - G_i(w_i; v_i))] \\
&\qquad \cdot z_i(v_i, \mathbf{v}_{-i}, w_i, \mathbf{w}_{-i})\, \mathrm{d}v_i
\end{aligned}$$

By our assumption on $y_S$ and $z_i(0, \mathbf{v}_{-i}, w_i, \mathbf{w}_{-i}) = 0$, Equation (14) can thus be calculated as

$$\int_{\mathbf{v}} \mathbf{f}(\mathbf{v}) \int_{\mathbf{w}} \mathbf{g}(\mathbf{w};\mathbf{v}) \cdot \left( \frac{1 - G_i(w_i; v_i)}{g_i(w_i; v_i)} - \frac{1 - F_i(v_i)}{f_i(v_i)} \right) \cdot z_i(\mathbf{v}, \mathbf{w})\, \mathrm{d}\mathbf{w}\, \mathrm{d}\mathbf{v}.$$

The second part of revenue is

$$\begin{aligned}
&\int_{\mathbf{v}} \mathbf{f}(\mathbf{v}) \int_{\mathbf{w}} \mathbf{g}(\mathbf{w};\mathbf{v}) \int_0^{w_i} \tilde{w}_i \cdot \frac{\partial z_i(\mathbf{v}, \tilde{w}_i, \mathbf{w}_{-i})}{\partial \tilde{w}_i} \mathrm{d}\tilde{w}_i\, \mathrm{d}\mathbf{w}\, \mathrm{d}\mathbf{v} \\
=& \int_{\mathbf{v}} \mathbf{f}(\mathbf{v}) \int_{\mathbf{w}} \mathbf{g}(\mathbf{w};\mathbf{v}) \cdot (w_i \cdot z_i(\mathbf{v}, \mathbf{w}) \\
&\qquad - \int_0^{w_i} z_i(\mathbf{v}, \tilde{w}_i, \mathbf{w}_{-i})\, \mathrm{d}\tilde{w}_i)\, \mathrm{d}\mathbf{w}\, \mathrm{d}\mathbf{v} \\
=& \int_{\mathbf{v}} \mathbf{f}(\mathbf{v}) \int_{\mathbf{w}} \mathbf{g}(\mathbf{w};\mathbf{v}) \cdot z_i(\mathbf{v}, \mathbf{w}) \left( w_i - \frac{1 - G_i(w_i; v_i)}{g_i(w_i; v_i)} \right) \mathrm{d}\mathbf{w}\, \mathrm{d}\mathbf{v}.
\end{aligned}$$

Combining the two parts, the total expected revenue from advertiser $i$ is thus $\int_{\mathbf{v}} \mathbf{f}(\mathbf{v}) \int_{\mathbf{w}} \mathbf{g}(\mathbf{w};\mathbf{v}) \cdot z_i(\mathbf{v}, \mathbf{w}) \left( w_i - \frac{1-F_i(v_i)}{f_i(v_i)} \right) \mathrm{d}\mathbf{w}\, \mathrm{d}\mathbf{v}$. □

## A.4 Proof of Proposition 4.4

Proof. Suppose a two-stage auction satisfies two-stage IC and ex-post IR simultaneously. By ex-post IR and non-negativity, taking $w_i = 0$, we must have $q_i(\mathbf{v}, 0, \mathbf{w}_{-i}) = 0, \forall \mathbf{v}, \mathbf{w}_{-i}$. Suppose we have two initial bids $a < b$ such that $Z_i(b, \hat{\mathbf{v}}_{-i}; \mathbf{v}) > Z_i(a, \hat{\mathbf{v}}_{-i}; \mathbf{v})$. This indicates $\int_a^b \int_0^{\overline{u}} g_i(w_i; \tilde{v}_i) \int_0^{w_i} \frac{\partial z_i(\tilde{v}_i, \hat{\mathbf{v}}_{-i}, \tilde{w}_i, \mathbf{w}_{-i})}{\partial \tilde{v}_i} \mathrm{d}\tilde{w}_i\, \mathrm{d}w_i\, \mathrm{d}\tilde{v}_i$, in expectation over $\mathbf{w}_{-i} \sim g_{-i}(\cdot; \hat{\mathbf{v}}_{-i})$, is strictly positive under our regular full-supported assumption, contradicting with Equation (11). Therefore, we must have $Z_i(\hat{v}_i, \hat{\mathbf{v}}_{-i}; \mathbf{v})$ constant in $\hat{v}_i$. □

## A.5 Proof of Theorem 5.1

Proof. (Proportional Entry Fee) As our choice of $\eta$ is based on the rule of the static two-stage auction, by its interim IR property, we have $U_i(\hat{\mathbf{v}}; \hat{v}_i) \geq 0$, indicating $\eta \leq 1$. The realized utility after charging $\hat{\tau}_i$ is $(1-\eta)\cdot\left(\hat{w}_i \cdot x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) - \int_0^{\hat{w}_i} \tilde{w}_i \cdot \frac{\partial x_i(\hat{\mathbf{v}}, \tilde{w}_i, \hat{\mathbf{w}}_{-i}; S)}{\partial \tilde{w}_i} \mathrm{d}\tilde{w}_i\right)$, which is always non-negative. To verify the expectation of this design aligns with the original entry fee, note that by definition, $\mathbb{E}_{\hat{\mathbf{w}} \sim \mathbf{g}(\cdot; \hat{\mathbf{v}}), S \sim y_S(\hat{\mathbf{v}})} \left[\hat{w}_i \cdot x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S) - p_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)\right] = U_i(\hat{\mathbf{v}}; \hat{v}_i) + \tau_i(\hat{\mathbf{v}})$, thus the two terms exactly align under the choice of $\eta$.

(Quantile-Based Entry Fee) We have shown in the main text that the overall payment $\hat{p}_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)$ under this entry fee becomes $x_i(\hat{\mathbf{v}}, \hat{\mathbf{w}}; S)\cdot\max\left\{r_{i,1}(\hat{\mathbf{v}}, \hat{\mathbf{w}}_{-i}; S), \phi_i\left(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\right)\right\}$. As we have $x_{i,S}(\hat{\mathbf{v}}, \hat{\mathbf{w}}) > 0$ indicating $\hat{v}_i \geq r_{i,0}(\hat{\mathbf{v}}_{-i})$ and $\hat{w}_i \geq r_{i,1}(\hat{\mathbf{w}}_{-i}; S)$, the first-order stochastic dominance property of $G$ further indicates $\hat{w}_i \geq \phi_i(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i}))$, thus guarantees the ex-post IR. By the equivalent form of entry fee in (13), we only need to verify

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\tilde{w}_i \sim G_i(\cdot | r_{i,0}(\hat{\mathbf{v}}_{-i}))} \left[\left(\tilde{w}_i - r_{i,1}(\hat{\mathbf{w}}_{-i}; S)\right)^+\right] \\
=& \mathop{\mathbb{E}}_{\hat{w}_i \sim G_i(\cdot | \hat{v}_i)} \left[\left(\phi_i\left(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\right) - r_{i,1}(\hat{\mathbf{w}}_{-i}; S)\right)^+\right],
\end{aligned}$$

which holds by applying a substitution in integration with $\tilde{w}_i = \phi_i\left(G_i(\hat{w}_i; \hat{v}_i); r_{i,0}(\hat{\mathbf{v}}_{-i})\right)$. □

## A.6 Proof of Theorem 5.2

Proof. Given the allocation rules $(y_S, x)$ and payment rules $(\hat{\tau}, p_{i,1})$, we use $\mathcal{M}_d$ to denote our mechanism in Algorithm 1, and use $\mathcal{M}$ to denote the corresponding static mechanism with standard entry fee rule $\tau$ in Theorem 4.1. We aim to bound the utility difference $\mathcal{U}_i(\mathcal{B}_i; \mathcal{M}_d) - \mathcal{U}_i(\mathcal{I}_i; \mathcal{M}_d)$ for any potential strategy $\mathcal{B}_i$. We use $\mathcal{P}_{i,0}(\mathcal{B}_i; \mathcal{M})$ to denote the expectation of the sum of entry fees terms $\tau_i$ or $\hat{\tau}_i$ (depending on the specific mechanism) under the same set of randomness. We may use $u(\mathcal{B}_i; \mathcal{M})$ or $p_{i,0}(\mathcal{B}_i; \mathcal{M})$ to denote a realization of a bidder's accumulative utility or entry fee over the entire $T$ periods. We may use $\tau$ or $\hat{\tau}$ only when indicating entry fee at a single time step.

In $\mathcal{M}_d$, a bidder $i$ is excluded from the auction, *i.e.*, treated as reporting $\hat{v}_i = \hat{w}_i = 0$, once her $L_i$ is broken. Therefore, for any strategy $\mathcal{B}_i$, we could equivalently construct a new distorted strategy $\tilde{\mathcal{B}}_i$ which always reports $(0, 0)$ after her $L_i$ is broken. By this definition, we have $\mathcal{U}_i(\mathcal{B}_i; \mathcal{M}_d) = \mathcal{U}_i(\tilde{\mathcal{B}}_i; \mathcal{M}_d)$ and $\mathcal{P}_{i,0}(\mathcal{B}_i; \mathcal{M}_d) = \mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}_d)$. Since $\tilde{\mathcal{B}}_i$ is designed to always not excluded in $\mathcal{M}_d$, the total value of allocation aligns for $\tilde{\mathcal{B}}$ in $\mathcal{M}_d$ and $\mathcal{M}$, *i.e.*, $\mathcal{U}_i(\tilde{\mathcal{B}}_i; \mathcal{M}_d) + (1 - k_i)\mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}_d) = \mathcal{U}_i(\tilde{\mathcal{B}}_i; \mathcal{M}) + \mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M})$, where we have canceled out the static second-stage payment terms $\mathcal{P}_{i,1}$ as they are defined as the same in two mechanisms. Specifically, for the distorted truthful strategy $\tilde{\mathcal{I}}_i$, by aligning expectation property of $\tilde{p}_0$, we have $\mathbb{E}\left[\hat{\tau}_i(\hat{\mathbf{v}}^t, \hat{\mathbf{w}}^t; S^t) \mid \mathcal{H}^{t,0}, \hat{\mathbf{v}}^t\right] = \tau_i(\hat{\mathbf{v}}^t)$. Let $\alpha_i^t = \mathbb{I}\{i \in A^{t-1}\}$, which is determined by the history before period $t$. Therefore, $\mathcal{P}_{i,0}(\tilde{\mathcal{I}}_i; \mathcal{M}_d) = \sum_{t=1}^T \mathbb{E}[\alpha_i^t \hat{\tau}_i(\hat{\mathbf{v}}^t, \hat{\mathbf{w}}^t; S^t)] = \sum_{t=1}^T \mathbb{E}[\alpha_i^t \tau_i(\hat{\mathbf{v}}^t)] = \mathcal{P}_{i,0}(\tilde{\mathcal{I}}_i; \mathcal{M})$. This indicates $\mathcal{U}_i(\mathcal{I}_i; \mathcal{M}_d) = \mathcal{U}_i(\tilde{\mathcal{I}}_i; \mathcal{M}_d) = \mathcal{U}_i(\tilde{\mathcal{I}}_i; \mathcal{M}) + k_i \cdot \mathcal{P}_{i,0}(\tilde{\mathcal{I}}_i; \mathcal{M})$. Therefore, we have

$$\begin{aligned}
&\mathcal{U}_i(\mathcal{B}_i; \mathcal{M}_d) - \mathcal{U}_i(\mathcal{I}_i; \mathcal{M}_d) \\
=& \mathcal{U}_i(\tilde{\mathcal{B}}_i; \mathcal{M}) + \mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}) - (1 - k_i)\mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}_d) - \mathcal{U}_i(\mathcal{I}_i; \mathcal{M}_d) \\
\leq& (\mathcal{U}_i(\mathcal{I}_i; \mathcal{M}) - \mathcal{U}_i(\mathcal{I}_i; \mathcal{M}_d)) + \mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}) - (1 - k_i)\mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}_d) \\
=& \left(\mathcal{U}_i(\mathcal{I}_i; \mathcal{M}) - \mathcal{U}_i(\tilde{\mathcal{I}}_i; \mathcal{M})\right) + \mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}) - (1 - k_i)\mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i; \mathcal{M}_d) \\
&- k_i \cdot \mathcal{P}_{i,0}(\tilde{\mathcal{I}}_i; \mathcal{M}).
\end{aligned}$$

where the inequality is by the IC property of static mechanism $\mathcal{M}$.

Since bidder $i$ is served in period $t$ only if $d_i^{t-1} \le L_I$, and the per-period increment $\tau_i(\hat{\mathbf{v}}) - (1-k_i)\cdot\hat{\tau}_i(\hat{\mathbf{v}},\hat{\mathbf{w}};S^t)$ is upper bounded by $\overline{u}$, we always have $\mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i;\mathcal{M}) - (1-k_i)\mathcal{P}_{i,0}(\tilde{\mathcal{B}}_i;\mathcal{M}_d) \le L_i + \overline{u}$. Recall $L_i = \epsilon_i + k_i\mathcal{P}_{i,0}(\mathcal{I}_i;\mathcal{M})$, this indicates

$$\begin{aligned}&\mathcal{U}_i(\mathcal{B}_i;\mathcal{M}_d) - \mathcal{U}_i(\mathcal{I}_i;\mathcal{M}_d)\\ \le &\left(\mathcal{U}_i(\mathcal{I}_i;\mathcal{M}) - \mathcal{U}_i(\tilde{\mathcal{I}}_i;\mathcal{M})\right) + L_i + \overline{u} - k_i\cdot\mathcal{P}_{i,0}(\tilde{\mathcal{I}}_i;\mathcal{M})\\ = &\left(\mathcal{U}_i(\mathcal{I}_i;\mathcal{M}) + k_i\cdot\mathcal{P}_{i,0}(\mathcal{I}_i;\mathcal{M})\right) - \left(\mathcal{U}_i(\tilde{\mathcal{I}}_i;\mathcal{M}) + k_i\cdot\mathcal{P}_{i,0}(\tilde{\mathcal{I}}_i;\mathcal{M})\right)\\ &+ \epsilon_i + \overline{u}.\end{aligned}$$

For events that $i \in A$ throughout $t = 1,\dots,T$ in $\mathcal{M}_d$, we have realizations $u_i(\mathcal{I}_i;\mathcal{M}) + k_i\cdot p_{i,0}(\mathcal{I}_i;\mathcal{M}) = u_i(\tilde{\mathcal{I}}_i;\mathcal{M}) + k_i\cdot p_{i,0}(\tilde{\mathcal{I}}_i;\mathcal{M})$. On the other hand, for any potential realization, we always have $\left(u_i(\mathcal{I}_i;\mathcal{M}) + k_i\cdot p_{i,0}(\mathcal{I}_i;\mathcal{M})\right) - \left(u_i(\tilde{\mathcal{I}}_i;\mathcal{M}) + k_i\cdot p_{i,0}(\tilde{\mathcal{I}}_i;\mathcal{M})\right) \le T\overline{u}$, as $\tilde{\mathcal{I}}_i$ gets outside payoff 0 after exclusion and the realization of first term in each period is upper bounded by $\overline{u}$. Denote $d_i^t$ as the value of $d_i$ achieved by truthful strategy $\mathcal{I}_i$ in Algorithm 1 at time $t$, *i.e.*, $d_i^t(\mathcal{I}_i;\mathcal{M}) = \sum_{t'=1}^{t} -(1-k_i)\cdot\hat{\tau}_i^{t'} + \tau_i^{t'}$, we have $\mathcal{U}_i(\mathcal{B}_i;\mathcal{M}_d) - \mathcal{U}_i(\mathcal{I}_i;\mathcal{M}_d) \le \Pr\left[\exists t\in[T], d_i^t(\mathcal{I}_i;\mathcal{M}) > L_i\right]\cdot T\overline{u} + \epsilon_i + \overline{u} \le \sum_{t=1}^{T}\Pr[d_i^t(\mathcal{I}_i;\mathcal{M}) > L_i]\cdot T\overline{u} + \epsilon_i + \overline{u}$, where the second inequality is by union bound. By Hoeffding inequality, for $X := \sum_{i\in[n]} X_i$ where $X_i, i\in[n]$ are independently distributed in $[a_i,b_i]$, we have $\Pr\left[X > \mathrm{E}[X] + \epsilon\right] \le \exp\left(-\frac{2\epsilon^2}{\sum_{i=1}^{n}(b_i-a_i)^2}\right)$. By design of aligning-expectation entry fee, we have $\mathbb{E}[d_i^t(\mathcal{I}_i;\mathcal{M})] = t\cdot k_i\overline{\tau_i}$, with $\tau_i^t - (1-k_i)\hat{\tau}_i^t \in [-\overline{u},\overline{u}]$ for any potential realization. Since $L_i = T\cdot k_i\overline{\tau_i} + \epsilon_i$, for any $t \le T$, $\Pr[d_i^t(\mathcal{I}_i;\mathcal{M}) > L_i] \le \Pr\left[d_i^t(\mathcal{I}_i;\mathcal{M}) - \mathbb{E}[d_i^t(\mathcal{I}_i;\mathcal{M})] > \epsilon_i\right] \le \exp\left(-\frac{\epsilon_i^2}{2t\overline{u}^2}\right) \le \exp\left(-\frac{\epsilon_i^2}{2T\overline{u}^2}\right)$. Applying $\Pr[d_i^t(\mathcal{I}_i;\mathcal{M}) > L_i] \le \exp\left(-\frac{\epsilon_i^2}{2T\overline{u}^2}\right)$ and taking $\epsilon_i = 2\overline{u}\sqrt{T\ln T}$, we have $\mathcal{U}_i(\mathcal{B}_i;\mathcal{M}_d) - \mathcal{U}_i(\mathcal{I}_i;\mathcal{M}_d) \le 2\overline{u} + 2\overline{u}\sqrt{T\ln T}$, which indicates our utility increment statement. This also indicates the probability of wrong exclusion, *i.e.*, the term $\Pr\left[\exists t, d_i^t(\mathcal{I}_i;\mathcal{M}) > L_i\right]$, is bounded by $T\exp\left(-\frac{\epsilon_i^2}{2T\overline{u}^2}\right) = T^{-1}$. □

# B A Synthetic Illustration of Bid Refinement

In this section, we present simple synthetic experiments to illustrate the two-stage auction environment and the role of bid refinement. The experiments focus on providing intuition for the modeling motivation and the potential value of bid refinement, while the mechanism itself is analyzed theoretically in the main text.

*Experiment Setting.* Each auction contains $n = 30$ bidders and $k = 5$ entrants. Motivated by the standard bid × CTR abstraction, we decompose bidder $i$'s per-impression value into a value-per-click component $\nu_i$ and the quality component $c_i$. Thus, $c_i$ plays the role of the coarse CTR estimated in the pre-auction stage, and the initial valuation is $v_i = \nu_i c_i$. In the formal stage, the refined information of $\nu_i$ and $c_i$ becomes available. We model this by $\omega_i = \max\{0, \nu_i + \epsilon_i^\nu\}$, $c_i' = \mathrm{clip}(c_i + \epsilon_i^c, 0, 1)$ with second-stage refinement variables $\epsilon_i^\nu, \epsilon_i^c$ and refined valuation $w_i = \omega_i c_i'$. The value-per-click components are privately revealed, while the quality components are publicly known. As the primary auction design objective, we evaluate social welfare, defined as the realized refined value of the winner; therefore, we instantiate each auction with the natural greedy selection rule for its available information.

*Baseline Auctions.* We consider three baselines: two-stage auctions that only elicit the first- or second-stage bid, and the single-stage auction. In detail, $B_1$ selects entrants by $\nu_i c_i$ and allocates among entrants by $\nu_i c_i'$. $B_2$ selects entrants by $\mathbb{E}[\nu_i]c_i$ and allocates by $w_i$. $B_3$ is a single-stage auction directly allocates based on $v_i$. We also consider anoracle OPT that allocates based on refined $w_i$ over all bidders, which drops the second-stage entrant constraint. We implement each auction to be incentive compatible, and thus assume bidders to report truthfully.

*Distribution Configuration.* We use uniform distribution in simulations. To inspect the performances under various informativeness of the first-stage attributes, we consider 9 noise levels for two components. Let $V_1, V_2, V_3$ denote the value-component settings

$$\begin{aligned}V_1 &: \nu_i \sim U[0.5, 5.5], \quad \epsilon_i^\nu \sim U[-0.5, 0.5],\\ V_2 &: \nu_i \sim U[1.5, 4.5], \quad \epsilon_i^\nu \sim U[-1.5, 1.5],\\ V_3 &: \nu_i \sim U[2.5, 3.5], \quad \epsilon_i^\nu \sim U[-2.5, 2.5],\end{aligned}$$

and let $Q_1, Q_2, Q_3$ denote the quality-component settings

$$\begin{aligned}Q_1 &: c_i \sim U[0.1, 0.9], \quad \epsilon_i^c \sim U[-0.1, 0.1],\\ Q_2 &: c_i \sim U[0.25, 0.75], \quad \epsilon_i^c \sim U[-0.25, 0.25],\\ Q_3 &: c_i \sim U[0.4, 0.6], \quad \epsilon_i^c \sim U[-0.4, 0.4].\end{aligned}$$

Configuration $C_{\ell m}$ combines $V_\ell$ and $Q_m$.

*Experiment Results.* Our results are averaged over $10^5$ runs. As shown in Table 1, our bid refinement scheme consistently achieves better social welfare due to its utilization of up-to-date, accurate information. Since $B_1$ refines only the quality (CTR) component, it performs better when the value component is stable, but degrades as the noise in $\nu_i$ increases. Conversely, $B_2$ benefits from refined formal-stage values and becomes stronger when value refinement is important. $B_3$ uses neither refined bids nor refined CTRs, and therefore deteriorates most clearly as the discrepancy between the two stages increases. The remaining gap between bid refinement to OPT reflects the intrinsic limitation of two-stage architecture: bidders excluded in the pre-auction stage cannot be recovered later. This supports the role of our bid refinement scheme as a truthful way to exploit refined valuations within the operational constraints of two-stage ad auctions.

**Table 1: Average social welfare in static synthetic auctions.**

| | $B_1$ | $B_2$ | $B_3$ | Ours | OPT |
|---|---|---|---|---|---|
| $C_{11}$ | 4.052 | 3.927 | 3.990 | **4.093** | 4.094 |
| $C_{12}$ | 3.850 | 3.510 | 3.423 | **3.882** | 3.902 |
| $C_{13}$ | 3.950 | 3.218 | 2.902 | **3.974** | 4.096 |
| $C_{21}$ | 3.433 | 3.690 | 3.371 | **3.838** | 3.861 |
| $C_{22}$ | 3.271 | 3.310 | 2.883 | **3.576** | 3.683 |
| $C_{23}$ | 3.325 | 3.060 | 2.422 | **3.551** | 3.867 |
| $C_{31}$ | 2.874 | 3.928 | 2.803 | **3.950** | 4.094 |
| $C_{32}$ | 2.751 | 3.515 | 2.380 | **3.563** | 3.906 |
| $C_{33}$ | 2.728 | 3.215 | 1.957 | **3.337** | 4.093 |