

Ex-post IR Dynamic Auctions with Cost-per-action Payments

Weiran Shen¹, Zihe Wang², Song Zuo¹

¹ Institute for Interdisciplinary Information Sciences, Tsinghua University

² Institute for Theoretical Computer Science, Shanghai University of Finance
emersonswr@gmail.com, wang.zihe@mail.shufe.edu.cn, songzuo.z@gmail.com

Abstract

Motivated by online ad auctions, we consider a repeated auction between one seller and many buyers, where each buyer only has an estimation of her value in each period until she actually receives the item in that period. The seller is allowed to conduct a dynamic auction but must guarantee ex-post individual rationality.

In this paper, we use a structure that we call *credit accounts* to enable a general reduction from any *incentive compatible* and *ex-ante individual rational* dynamic auction to an *approximate incentive compatible* and *ex-post individually rational* dynamic auction with credit accounts. Our reduction obtains stronger individual rationality guarantees at the cost of weaker incentive compatibility. Surprisingly, our reduction works without any common knowledge assumption. Finally, as a complement to our reduction, we prove that there is no non-trivial auction that is exactly incentive compatible and ex-post individually rational under this setting.

1 Introduction

Internet advertising has been playing a very important role in the advertising industry. Most online advertising platforms, such as search engines and social media, have gone through the evolution from the *cost-per-mille impressions* (CPM) model to the *cost-per-click* (CPC) model, where the former is aligned with traditional advertising while the latter focuses more on performance. In the CPC model, when a user requests a certain web page, the platform collects bids from the advertisers and based on the bids, determines whose ad to display on the page. The corresponding advertiser is charged when her advertisement is clicked by the user. Such an advertising model is called the CPC model because the advertiser only needs to pay when her advertisement is clicked. This CPC model has been the de facto model for most major online advertising platforms, and is proven to be profitable [Edelman *et al.*, 2007]. However, despite its success, this model is criticized to have the click fraud problem, i.e., the competitors of an advertiser, or even the platform itself, may deliberately create false clicks to increase the advertiser’s cost or to extract more revenue. Furthermore, the advertisers have to pay

for clicks that do not lead to final purchase of their products. Although one may argue that in expectation the advertisers are indeed profitable, it may still be a serious problem for small companies that cannot ignore such risks.

A relatively new model that has gained much research attention recently is the *cost-per-action* (CPA) advertising model. In contrast to the CPC model, the CPA model is even more performance-oriented and focuses directly on user actions. In the CPA model, the advertisers are only charged when the users make certain actions, such as purchases or transactions. It seems that the CPA model and the CPC model are almost the same except for the payment. However, this advertising model clears the uncertainty faced by the advertiser and can potentially decrease the vulnerability to click fraud. Besides these advantages, the CPA model also gives more incentives to the platforms to deliver high-quality impressions to the users. In 2007, the CPA model was described as the “Holy Grail” of targeted advertising by Google [Spencer, 2007]. Currently, many online advertising platforms, including Google, eBay, Amazon, Facebook, Baidu and WeChat have already started to test the CPA model.

Another essential difference between these two models is that the platform cannot directly observe the users’ actions on the advertisers’ websites whereas the users’ clicks are observable by both the platform and the advertiser. Such an undesirable property may cause the advertisers to hide the users’ actions to avoid payments. This also poses challenges in putting the CPA model in practice to replace the CPC model that is currently dominant in the advertising industry.

This paper is directly motivated by the above challenge. In this paper, we aim to tackle the incentive problem and present a new reduction framework that we call the *credit account mechanism*. Our mechanism solves the incentive issue by adding a credit account for each advertiser on top of the original mechanism that follows the “allocate-report-pay” scheme. In our mechanism, the advertisers are given a certain amount of “credit quota” and an advertiser cannot win the auction if her credit runs out of her “quota”. When an impression is allocated, the winner observes her value and reports to the platform. Our mechanism then charges her by some modified payment based on the original mechanism. The payment is chosen such that it never exceeds the bidder’s reported value so that ex-post individual rationality is guaranteed. In the meanwhile, the winner’s credit balance is updated

according to the difference between her actual payment and her expected payment in the original mechanism.

Intuitively, the credit account tracks the difference between one’s actual payments and expected payments. An honest advertiser only has a negligible chance of consuming all her credit, while an advertiser shading her reports will quickly run out of credit and will be blocked in future auctions.

Our contributions Firstly, we formalize the credit accounts framework.¹ As for incentive compatibility (IC) and individual rationality (IR) properties, our framework can reduce any IC and ex-ante IR mechanism to a credit account mechanism that gives the same allocation rule as the original mechanism with high probability (Theorem 4.1) and guarantees approximate IC (see Definition 3.3) and ex-post IR (see Definition 2.4). We emphasize that the notion of approximate IC serves the motivating example of online advertising, since the benefit of deviation on average is vanishingly small as the number of periods grows (Theorem 4.2).

Such a reduction then naturally induces a trade-off between the strength of truthfulness and the probability of desired implementation (see Section 4). In particular, it also applies to second price auctions (Proposition 4.3).

Finally, as a complement to the constructed credit account mechanisms, we show that the exact IC and the ex-post IR cannot be achieved simultaneously, unless the mechanism is trivial (Theorem 5.1). In this sense, credit account mechanisms have achieved the strongest properties we can hope for.

1.1 Related Works

Ever since Myerson’s seminal paper on designing revenue optimal auctions [Myerson, 1981], there have been intensive researches on analyzing and designing one-shot auction mechanisms. For example, Edelman *et al.* [2007] and Varian [2007] study the performance of the generalized second price auction (GSP). Hartline and Roughgarden [2009], Shen and Tang [2017] and Bachrach *et al.* [2014] provide mechanisms that can tradeoff among different objectives. There is also a rich literature on multi-item auctions [Cai *et al.*, 2012; Daskalakis *et al.*, 2013; Wang and Tang, 2014; Yao, 2014; Tang and Wang, 2017; Yao, 2017], and on repeated auctions motivated by online advertising [Amin *et al.*, 2013; Amin *et al.*, 2014; Kanoria and Nazerzadeh, 2014; Balseiro *et al.*, 2017; Epasto *et al.*, 2018; Tang and Zeng, 2018].

A closely related line of work is dynamic mechanism design (see Bergemann and Välimäki [2017] for a comprehensive survey). Our mechanism is also related to works on auctions with unknown types, for example the “common value” literature [Klemperer, 1998; Bergemann and Morris, 2013].

There is also a series of works that focus on designing mechanisms with the CPA advertising model. Nazerzadeh *et al.* [2013] study the setting where the advertisers’ value may evolve over time. They present a mechanism that satisfies asymptotic IR and asymptotic IC. However, their mechanism does not exactly fall into the CPA advertising model, since

¹In fact, we are not the first to use the idea of credit accounts in mechanism design. Similar reputation based structures are used for other settings [Liu *et al.*, 2014; Hajaj *et al.*, 2015].

the winner still needs to pay even if the user does not click on her ad. Hu *et al.* [2015] compare the CPC model and the CPA model. Their results show that the CPA model is better in incentivizing the platform to improve the purchase rate, but suffers from the adverse selection problem. Agarwal *et al.* [2009] consider a similar setting where the advertisers report both the predefined actions and the action probabilities. They show that at equilibrium, the advertisers may report skewed bids. However, their results only hold in one-shot games.

Our mechanism also borrows some highlevel ideas from the “bank account” mechanism, where the seller maintains a “bank account” for each buyer during the dynamic auction [Mirrokni *et al.*, 2016a; Mirrokni *et al.*, 2016b; Mirrokni *et al.*, 2018a; Mirrokni *et al.*, 2018b]. Although with similar names, the “credit account” in this paper is fundamentally different: (i) the bank account mechanisms are designed under the common knowledge assumption to ensure dynamic IC, while the credit account mechanism guarantees approximate dynamic IC *without* any common knowledge assumption; (ii) the “balance” in bank accounts can be thought of as money, where the buyers might pay their bank account balance, while the “credit” in the credit accounts is more like a reliability measure of the buyers based on their past behaviors.

2 Preliminaries and Setup

Setup We study the problem of cost-per-action mechanism design in the environment with one seller and multiple buyers. In particular, we will focus on the setting with repeated sales, that is, in each period $t \in [T]$, the seller has a new item to sell to the buyers. If not sold, the item is then destroyed. Throughout this paper, we consider the finite horizon case without discount factor, hence $T < \infty$.²

Similar to the standard setting, in each period, what the seller does is to allocate the item to one (or none) of the buyers and charge the buyers some amount of money as payments. Formally, we use $x_t^i \in [0, 1]$ to denote the probability of buyer $i \in [n]$ receiving the item of period t and $p_t^i \in \mathbb{R}$ to denote the payment from buyer i to the seller in period t .

The main difference between our setup and the standard setting is how the buyers value the received items:

- In the standard setting, there is a private value v_t^i of buyer i receiving the t -th item, only known to buyer i .³
- In our setup, each buyer is uncertain about her exact value for the item until she actually receives it.

Formally, an *estimation* of buyer i for the t -th item is a private distribution F_t^i of her possible *exact values* and her *exact value* $v_t^i \in [0, \bar{v}]$ is realized only after receiving the item.⁴ In particular, $v_t^i \sim F_t^i$ and we assume that each v_t (i.e., the vector of v_t^1, \dots, v_t^n) is sampled from the joint distribution \mathcal{F}_t

²Note that the extension to infinite horizon case with discount factor is straightforward.

³In Bayesian settings, the other agents (including the seller and the buyers other than i) might only have a common knowledge of the distribution of v_t^i .

⁴ \bar{v} is just a finite upper bound on the buyer’s value and does not restrict the bid of the buyers (see Mechanism 2.1).

but independently over time. Hence F_t^i is a marginal distribution of \mathcal{F}_t for buyer i . We also allow each F_t^i to be different across both the buyers and periods.

Although such a setup is quite different from the canonical auction setting, it does help us to capture the nature of on-line advertising. In online advertising, there are two types of ads, brand based ads and performance based ads. For the first type of ads, the advertisers acquire value whenever the ads are shown to a user, while for the second type of ads, the advertisers acquire value from some specific user actions on their websites (e.g., buying some goods or services, subscribing to the website, etc.). In other words, for the second type of ads, the advertiser only has an estimation for her value of an impression (or a click) to a user, while her exact value depends on the actual actions taken by the user.

Under such a setup, we assume that both estimation F_t^i and the exact value v_t^i are private information of buyer i . Although in general, the estimation could rely on both the seller's private information, such as the characteristics of the user, and the buyer's private information, such as the prior knowledge of certain users' behaviors. However, such an information structure is not the main focus of this paper. Hence, to make the model clean, we assume that the estimation F_t^i is private information of buyer i . With such a simplification, we can ensure a more robust incentive guarantee for the buyers. In fact, our results also easily extend to the general settings.

Throughout this paper, we focus on "direct mechanisms", where the seller requires the buyers to report all their private information and determine the allocation of the item and the payments from the buyers. Although in general, reporting the entire estimation F_t^i is required, for most practical mechanisms, reporting the expected value $\mathbb{E}[F_t^i]$ would suffice.

The following abstract auction outlines the common structure of a direct mechanism under this setup:

Mechanism 2.1 (Abstract direct mechanism). *For each period $t \in [T]$:*

1. *Each buyer has a private estimation F_t^i for the t -th item and reports an estimation \hat{F}_t^i to the seller;*
2. *The seller determines the allocation $\mathbf{x}_t \in [0, 1]^n$ of the t -th item based on the reported estimations $\hat{\mathbf{F}}_t$ and all the historical reports, i.e., $\hat{\mathbf{F}}_{1..t-1} := \hat{\mathbf{F}}_1, \dots, \hat{\mathbf{F}}_{t-1}$ and $\hat{\mathbf{v}}_{1..t-1} := \hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_{t-1}$;*
3. *The winner j receives the item, realizes her exact value $v_t^j \sim F_t^j$, and reports $\hat{v}_t^j \in \mathbb{R}_+$ to the seller;*
4. *The seller determines the payment \mathbf{p}_t based on the reported estimations and values from the current and past periods, i.e., $\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t}$.*

Since only the winner reported her exact value in each period, hence \mathbf{v}_t (or $\hat{\mathbf{v}}_t$) is a vector consists of one $v_t^{j_t}$ (or $\hat{v}_t^{j_t}$) and $n - 1$ empty elements \perp :

$$\mathbf{v}_t = \perp, \dots, v_t^{j_t}, \dots, \perp \quad \hat{\mathbf{v}}_t = \perp, \dots, \hat{v}_t^{j_t}, \dots, \perp,$$

where j_t is the index of the winner in period t .⁵ In addition,

⁵If there is no winner in period t , then $\mathbf{v}_t = \hat{\mathbf{v}}_t = \perp, \dots, \perp$.

we will adopt the short notation $a_{1..t}$ for a_1, \dots, a_t throughout this paper.

In such a direct mechanism, in each period t , each buyer acquires a quasi-linear utility u_t^i :

$$u_t^i(\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t}; v_t^i) = v_t^i \cdot \mathbf{x}_t^i(\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t-1}) - p_t^i(\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t}).$$

Her cumulative utility is the sum of her utility in each period:

$$U^i(\hat{\mathbf{F}}_{1..T}, \hat{\mathbf{v}}_{1..T}; v_{1..T}^i) = \sum_{t \in [T]} u_t^i(\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t}; v_t^i).$$

Finally, the revenue of the seller, REV, is the cumulative payments collected from the buyers, and the social welfare, WEL, is the sum of REV and all the cumulative buyer utilities:

$$\text{REV} = \sum_{t \in [T]} \sum_{i \in [n]} p_t^i$$

$$\text{WEL} = \text{REV} + \sum_{i \in [n]} U^i = \sum_{t \in [T]} \sum_{i \in [n]} v_t^i \cdot \mathbf{x}_t^i.$$

We assume the buyers are risk-neutral and self-interested, hence the best strategy of each buyer i must maximize her cumulative utility U^i . For the seller we do not fix any specific goal now, since we mainly focus on the implementability.

A direct mechanism is *truthful* or *incentive compatible*, if in each period t , reporting $\hat{F}_t^i = F_t^i$ and $\hat{v}_t^i = v_t^i$ (if she won the item) is the dominant strategy of buyer i , regardless of the strategies of others and all historical bids. In this paper, we focus on *dominant-strategy incentive compatibility* because we do not make any common knowledge assumptions.

Let $(\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t})$ be any reporting profile, and $(\hat{\mathbf{F}}_{1..t}^{*i\tau}, \hat{\mathbf{v}}_{1..t}^{*i\tau})$ be the reporting profile where buyer i reports truthfully in periods τ, \dots, t , while the other entries remain the same:

$\hat{\mathbf{F}}_{1..t}^{*i\tau} = \hat{\mathbf{F}}_{1..t-1}, \hat{\mathbf{F}}_{\tau..t}^{-i}, F_{\tau..t}^i$ $\hat{\mathbf{v}}_{1..t}^{*i\tau} = \hat{\mathbf{v}}_{1..t-1}, \hat{\mathbf{v}}_{\tau..t}^{-i}, v_{\tau..t}^i$, where superscript $-i$ indicates the quantities for all buyers other than i .

To make the definition compact, we define the following notations. Let I_{it} be the expected utility of buyer i reporting truthfully since the first step of period t with arbitrary fixed historical reports and reports by other buyers:

$$I_{it} = \mathbb{E}_{v_{1..t}^i \sim F_{1..t}^i} \left[U^i(\hat{\mathbf{F}}_{1..T}, \hat{\mathbf{v}}_{1..T}^{*it}; v_{1..T}^i) \right].$$

Let Π_{it} be the similar utility of buyer i , except that she reports truthfully since the *second step* of period t while her report for the *first step* is fixed:

$$\Pi_{it} = \mathbb{E}_{v_{1..t}^i \sim F_{1..t}^i} \left[U^i(\hat{\mathbf{F}}_{1..T}^{*i,t+1}, \hat{\mathbf{v}}_{1..T}^{*it}; v_{1..T}^i) \right].$$

Let I'_{it} and Π'_{it} be the expected utility for buyer i misreporting solely in the first step and the second step of period t :

$$I'_{it} = \mathbb{E}_{v_{1..t}^i \sim F_{1..t}^i} \left[U^i(\hat{\mathbf{F}}_{1..T}^{*i,t+1}, \hat{\mathbf{v}}_{1..T}^{*it}; v_{1..T}^i) \right];$$

$$\Pi'_{it} = \mathbb{E}_{v_{1..t}^i \sim F_{1..t}^i} \left[U^i(\hat{\mathbf{F}}_{1..T}^{*i,t+1}, \hat{\mathbf{v}}_{1..T}^{*i,t+1}; v_{1..T}^i) \right].$$

Formally, we have the following definition:

Definition 2.2 ((Dynamic) Dominant-Strategy Incentive Compatible). *A direct mechanism is dominant-strategy incentive compatible, if truthfully reporting is the best strategy for each buyer i in each period t (both before and after the allocation of the item):*

$$\begin{aligned} \forall i, t, \hat{\mathbf{F}}_{1..T}, \hat{\mathbf{v}}_{1..T}, F_{1..T}^i, v_{1..T}^i, \quad I_{it} &\geq I'_{it}; \\ \forall i, t, \hat{\mathbf{F}}_{1..T}, \hat{\mathbf{v}}_{1..T}, F_{1..T}^i, v_{1..T}^i, \quad \Pi_{it} &\geq \Pi'_{it}. \end{aligned} \quad (\text{DIC})$$

Besides the incentive compatibility, another important guarantee is *individual rationality*, which means that the util-

ity of each buyer must be non-negative. Formally,

Definition 2.3 (Ex-ante Individually Rational). *A direct mechanism is ex-ante individually rational, if each buyer's expected utility in each period t is non-negative, $\forall \hat{F}_{1..t}^{*it}, \hat{v}_{1..t}^i$:*

$$\mathbb{E}_{v_t^i \sim F_t^i} \left[u_t^i(\hat{F}_{1..t}^{*it}, \hat{v}_{1..t}^{*it}; v_t^i) \right] \geq 0. \quad (\text{EX-ANTE IR})$$

Ex-post individual rationality is a stronger condition in the sense that the utility of each buyer in each period must be non-negative without taking expectation over v_t^i :

Definition 2.4 (Ex-post Individual Rational). *A direct mechanism is ex-post individual rational, if the utility of each buyer i is non-negative in each period t for any v_t^i :*

$$\forall i, t, \hat{F}_{1..t}^{*it}, \hat{v}_{1..t}^{*it}, u_t^i(\hat{F}_{1..t}^{*it}, \hat{v}_{1..t}^{*it}; v_t^i) \geq 0. \quad (\text{EX-POST IR})$$

3 Mechanisms

In this section, we present our first main result, the credit account mechanism, which is a general reduction framework that reduces any (DIC) and (EX-ANTE IR) mechanism to an mechanism that achieves approximate incentive compatibility (to be defined later) and (EX-POST IR). In particular, as we will show in the next section, the new mechanism implements the same allocation rule as the original one with high probability when the buyers report truthfully.

3.1 An Efficient Auction with Ex-ante IR

Let's start with a simple auction that maximizes the social welfare WEL and satisfies (DIC) and (EX-ANTE IR). Then we show how we can strengthen the IR from (EX-ANTE IR) to (EX-POST IR) with a tolerable loss in the guarantee of IC.

The following auction in each period is in fact a second price auction with $\mathbb{E}[\hat{F}_t^i]$ as each buyer's bid. Note that the payment p_t^i is independent of her reported exact value \hat{v}_t^i , so the mechanism satisfies (DIC). On the other hand, since only the winner is charged and for each winner, $\mathbb{E}_{v_t^i \sim F_t^i} [u_t^i] = x_t^i \cdot (b_t^i - b_t^{(2)}) \geq 0$, the mechanism also satisfies (EX-ANTE IR). However, in general, it does not satisfy (EX-POST IR) since if buyer j wins in period t and $v_t^j = 0$, $u_t^j = -b_t^{(2)}$ is negative.

Mechanism 3.1 (Ex-ante IR Auction). *In each period t :*

1. Ask each buyer to report \hat{F}_t^i and define her bid as $b_t^i = \mathbb{E}[\hat{F}_t^i] := \mathbb{E}_{v \sim \hat{F}_t^i}[v]$, or equivalently, ask each buyer to report such b_t^i directly.
2. Allocate the item to the buyer with the highest bid, break ties arbitrarily:

$$x_t^i = \begin{cases} 1 & \forall j \neq i, b_t^i > b_t^j \\ 0 & \exists j \neq i, b_t^i < b_t^j \\ \text{by the tie-breaking rule} & \text{otherwise} \end{cases}$$

3. Ignore the winner i 's exact value and simply charge her the second highest bid: $p_t^i = x_t^i \cdot b_t^{(2)}$, where $b_t^{(2)} = \max_{j \neq i} b_t^j$ is the second highest bid.

Lemma 3.2. *Mechanism 3.1 maximizes the social welfare and is (DIC) and (EX-ANTE IR), but not (EX-POST IR).*

3.2 The Credit Account Mechanism

We now present the credit account mechanism. As we mentioned previously, it reduces a mechanism \mathcal{M} to another mechanism $\tilde{\mathcal{M}}$ that is *approximate incentive compatible* (see Definition 3.3 below) and ex-post individually rational.

Definition 3.3. *A direct mechanism is ϵ -approximate incentive compatible, if for each buyer i , no strategy could outperform truthfully reporting by a margin larger than ϵ :*

$$\forall F_{1..T}^i, v_{1..T}^i, \hat{F}_{1..T}, \hat{v}_{1..T}, \mathcal{U}^i \geq \mathcal{U}^{i'} - \epsilon, \quad (\text{APPROX IC})$$

where $\mathcal{U}^i = \mathbb{E}[U^i(F_{1..T}^i, \hat{F}_{1..T}^{-i}, v_{1..T}^i, \hat{v}_{1..T}^{-i}; v_{1..T}^i)]$ and $\mathcal{U}^{i'} = \mathbb{E}[U^i(\hat{F}_{1..T}, \hat{v}_{1..T}; v_{1..T}^i)]$ are buyer i 's expected utility for truth-telling and any given strategy, respectively.

Credit account mechanisms To enforce the ex-post IR constraint, the major challenge is that the winner is incentivized to report her exact value as 0, because her payment will be restricted to 0 by (EX-POST IR). In other words, if the strict IC is also enforced, the payment must be always 0 and cannot implement general allocation rules.

The credit account mechanism circumvents the contradiction by slightly relaxing the IC constraint as well as tolerating a small probability that the allocation of the resulting mechanism $\tilde{\mathcal{M}}$ is different from the original \mathcal{M} .

The core concept of the reduction framework, is a *credit account* for each buyer. Intuitively, each credit account maintains the ‘‘reliability’’ of the corresponding buyer, i.e., it tracks the difference between the total price the buyer should pay under the rule of the original mechanism \mathcal{M} and the price she actually paid so far in the resulting mechanism $\tilde{\mathcal{M}}$. Since the draws of each $v_t^i \sim F_t^i$ are independent over time, by the law of large numbers, the magnitude of the credit account balance grows at the order of \sqrt{t} with high probability.

Based on the observation, there are two major changes in the resulting mechanism comparing with the original:

- The payment rule is modified in the way such that i) it never exceeds the reported exact value \hat{v}_t^i and ii) its expectation equals the payment rule of the original mechanism if the buyer reports her exact value truthfully;
- A gradually growing *credit quota*, q_t , is set for the buyers. Once a buyer's credit balance exceed the quota, she will be labeled as ‘‘unreliable’’ and the item won't be allocated when she wins. Hence her overall benefit of misreporting is essentially bounded by the quota.

Formally, the credit account mechanism $\tilde{\mathcal{M}} = \langle \tilde{x}, \tilde{p} \rangle$ for any given $\mathcal{M} = \langle x, p \rangle$ is defined as follows:

Mechanism 3.4 (Credit Account Mechanism). *Initially, let the credit account of each buyer be $c_1^i = 0$ and q_t be the quota at period t all buyers.*

For each period t :

1. Ask each buyer i to report \hat{F}_t^i and define α_t^i as:⁶

$$\alpha_t^i = \min\{c_t^i + q_t, \mathbb{E}[\hat{F}_t^i]\} / \mathbb{E}[\hat{F}_t^i]; \quad (1)$$

2. Allocate the item according to

$$\tilde{x}_t^i(\mathbf{c}_t, \hat{\mathbf{F}}_t) = \alpha_t^i \cdot x_t^i(\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t-1}); \quad (2)$$

3. Winner j reports her exact value \hat{v}_t^j and pays

$$\tilde{p}_t^j(\mathbf{c}_t, \hat{\mathbf{F}}_t, \hat{\mathbf{v}}_t) = \frac{\alpha_t^j \cdot \hat{v}_t^j}{\mathbb{E}[\hat{F}_t^j]} \cdot \tilde{p}_t^j, \quad (3)$$

where \tilde{p}_t^j is her expected payment under the original mechanism \mathcal{M} with value v' sampled from \hat{F}_t^j :⁷

$$\tilde{p}_t^j = \mathbb{E}_{v' \sim \hat{F}_t^j} [p_t^j(\hat{\mathbf{F}}_{1..t}, \hat{\mathbf{v}}_{1..t-1}, \hat{\mathbf{v}}_t^{-j}, v')]; \quad (4)$$

4. Increase the credit account of the winner by her payment and reduce the credit account of the winner by \tilde{p}_t^j , i.e.,

$$c_{t+1}^j = c_t^j + \tilde{p}_t^j - \alpha_t^j \cdot \tilde{p}_t^j; \quad (5)$$

5. For all other (not winning) buyers, charge them nothing and keep their credit accounts unchanged:

$$\forall i \neq j, \tilde{p}_t^i(\mathbf{c}_t, \hat{\mathbf{F}}_t, \hat{\mathbf{v}}_t) = 0, \quad c_{t+1}^i = c_t^i. \quad (6)$$

4 Truthfulness vs Efficiency

With the general reduction, we can construct a dynamic cost-per-action auction that is approximate efficient, achieving at least $(1 - \delta)$ fraction of the maximum social welfare, and approximate IC with $\epsilon = 2q_T + \bar{v}$ (see Proposition 4.3).

In this section, we present the second main result of this paper, which is how the q_t determines the trade-off between implementation tolerance δ and the truthfulness tolerance ϵ .

To establish such a trade-off, we prove some even stronger properties of the credit account mechanisms with any general \mathcal{M} satisfying (DIC) and (EX-ANTE IR). Intuitively:

- *Implementation:* The resulting credit account mechanism $\tilde{\mathcal{M}}$ implements the same allocation rule as \mathcal{M} with high probability $1 - \delta$ (Theorem 4.1);
- *Truthfulness:* The resulting credit account mechanism $\tilde{\mathcal{M}}$ is ϵ -APPROX IC and EX-POST IR (Theorem 4.2).

The trade-off between the failure probability δ and the approximate IC ϵ is roughly as follows:

$$\epsilon \leq 4\bar{v}\sqrt{t(\ln t + \ln n + \ln \delta^{-1})}.$$

We remark that having ϵ sub-linear in t is especially meaningful for the application of online advertising, where the value in each period \bar{v} is typically small, whereas T is quite large.

Formally, we have the following two theorems, whose proofs are deferred to Section 4.1:

Theorem 4.1. *For any DIC and EX-ANTE IR mechanism \mathcal{M} , there exists an equivalent credit account mechanism $\tilde{\mathcal{M}}$, such that when all buyers report truthfully, for all possible $\mathbf{F}_{1..T}$ and $\mathbf{v}_{1..T} \sim \mathbf{F}_{1..T}$, the allocation $\tilde{\mathbf{x}}_{1..T}$ are the same as in \mathcal{M} with high probability $(1 - \delta)$, where the quota $q_t \leq 2\bar{v}\sqrt{t(\ln t + \ln n + \ln \delta^{-1})}$.*

⁶ α_t^i is set in the way that ensures the buyer i 's credit balance never exceeds the quota in the worst case.

⁷Note that v' is not the private value realized by the winner j . It is just sampled from her reported estimation \hat{F}_t^j .

Theorem 4.2. *For any DIC and EX-ANTE IR mechanism \mathcal{M} , the corresponding credit account mechanism $\tilde{\mathcal{M}}$ defined according to Mechanism 3.4 satisfies both (APPROX IC) and (EX-POST IR), with $\epsilon \leq 2q_T + \bar{v} = \tilde{O}(\sqrt{T})$ sub-linear in T .*

Taking Mechanism 3.1 as \mathcal{M} , we get a direct consequence:

Proposition 4.3. *Let \mathcal{M} be Mechanism 3.1, then the corresponding credit account mechanism $\tilde{\mathcal{M}}$ is APPROX IC with $\epsilon = 2q_T + \bar{v} = \tilde{O}(\sqrt{T})$ and EX-POST IR. Meanwhile, under truthful reporting, $\tilde{\mathcal{M}}$ can achieve at least $(1 - \delta)$ fraction of the maximum social welfare, where δ can be as small as any polynomial of $1/T$, i.e., $\Omega(T^{-k})$ for any constant k .*

In particular, in each period, each buyer only needs to report her expected value $\mathbb{E}[\hat{F}_t^i]$ (instead of her full estimation) and the winner also needs to report her realized value \hat{v}_t^i .

4.1 Formal Proofs

Proof of Theorem 4.1. According to the construction of Mechanism 3.4, the allocation of the constructed $\tilde{\mathcal{M}}$ is almost the same as in \mathcal{M} . In particular, if $\forall i, t, \alpha_t^i = 1$, the allocations, $\mathbf{x}_{1..T}$ and $\tilde{\mathbf{x}}_{1..T}$, will be exactly the same. In the rest of the proof, we show it holds with high probability:

$$\Pr[\forall i, t, \alpha_t^i = 1] \geq 1 - \delta. \quad (\diamond)$$

Note that $\alpha_t^i < 1$ only if $c_t^i + q_t \leq \mathbb{E}[F_t^i] \leq \bar{v}$, hence $\Pr[\forall i, t, \alpha_t^i = 1] \geq \Pr[\forall i, t, c_t^i \geq \bar{v} - q_t]$. According to the construction of $\tilde{\mathcal{M}}$, for any $i, \{c_t^i\}_{t=1}^T$ is a martingale. To prove this, it suffices to show $\mathbb{E}[c_{t+1}^i | c_t^i] = c_t^i$.

If buyer i is not the winner in period t , then $c_{t+1}^i = c_t^i$. Otherwise, buyer i is the winner and by the construction (5), $\mathbb{E}[c_{t+1}^i | c_t^i] = \mathbb{E}[c_t^i + \tilde{p}_t^i - \alpha_t^i \cdot \tilde{p}_t^i | c_t^i] = c_t^i + \mathbb{E}[\tilde{p}_t^i - \alpha_t^i \cdot \tilde{p}_t^i | c_t^i]$.

According to the construction (3), for any fixed F_t^i ,

$$\begin{aligned} \mathbb{E}[\tilde{p}_t^i | c_t^i, F_t^i] &= \mathbb{E}\left[\frac{\alpha_t^i \cdot v_t^i}{\mathbb{E}[F_t^i]} \cdot \tilde{p}_t^i | c_t^i, F_t^i\right] \\ &= \alpha_t^i \cdot \frac{\mathbb{E}_{v_t^i \sim F_t^i}[v_t^i]}{\mathbb{E}[F_t^i]} \cdot \mathbb{E}[\tilde{p}_t^i | c_t^i, F_t^i] = \alpha_t^i \cdot \tilde{p}_t^i, \end{aligned}$$

where the second equality is from the independence between v_t^i and the randomness of \tilde{p}_t^i and the last equality is from the construction (4) (\tilde{p}_t^i is fixed once c_t^i and F_t^i are given).

Hence $\mathbb{E}[\tilde{p}_t^i - \alpha_t^i \cdot \tilde{p}_t^i | c_t^i] = 0$ and $\mathbb{E}[c_{t+1}^i | c_t^i] = c_t^i$, so $\{c_t^i\}_{t=1}^T$ is a martingale. Note that $|c_{t+1}^i - c_t^i| \leq \bar{v}$, then by the Azuma-Hoeffding inequality,

$$\Pr[c_t^i - c_0^i < \bar{v} - q_t] \leq \exp\left(-\frac{(\bar{v} - q_t)^2}{2t\bar{v}^2}\right).$$

By $q_t = \bar{v} + \bar{v}\sqrt{2t} \cdot \sqrt{2 \ln t + \ln n + 2 \ln \pi - \ln 6 + \ln \delta^{-1}} = \tilde{O}(\sqrt{t})$ and the union bound, we conclude that⁸

$$\begin{aligned} \Pr[\exists i, t, c_t^i < \bar{v} - q_t] &\leq \sum_{i \in [n], t \in [T]} \Pr[c_t^i - c_0^i < \bar{v} - q_t] \\ &\leq n \sum_{t \in [T]} \frac{6}{\pi^2} \cdot \frac{\delta}{nt^2} \leq \delta \cdot \frac{6}{\pi^2} \sum_{t=1}^{\infty} \frac{1}{t^2} = \delta. \end{aligned}$$

Therefore, we complete the proof of (\diamond) : $\Pr[\forall i, t, \alpha_t^i = 1] \geq 1 - \Pr[\exists i, t, c_t^i < \bar{v} - q_t] \geq 1 - \delta$ \square

Proof of Theorem 4.2. EX-POST IR:

We first prove that $\tilde{\mathcal{M}}$ is (EX-POST IR). Note that the original mechanism \mathcal{M} is (EX-ANTE IR), thus we have for the

⁸The famous Basel problem is used here: $1/1^2 + 1/2^2 + \dots + 1/n^2 + \dots = \pi^2/6$. The problem is first solved by Euler in 1734.

winner j in period t

$$\begin{aligned} \tilde{p}_t^j &= \mathbb{E}_{v_t^j} [p_t^j(\hat{F}_{1..t}, \hat{v}_{1..t-1}, \hat{v}_t^{-j}, v_t^j)] \\ &\leq \mathbb{E}_{v_t^j} [v_t^j \cdot x_t^j(\hat{F}_{1..t}, \hat{v}_{1..t-1})] \leq \mathbb{E}[F_t^j] \cdot x_t^j(\hat{F}_{1..t}, \hat{v}_{1..t-1}). \end{aligned}$$

Then by construction (3), $\tilde{p}_t^j = \frac{\alpha_t^j \cdot \hat{v}_t^j}{\mathbb{E}[F_t^j]} \cdot \tilde{p}_t^j \leq \alpha_t^j \cdot x_t^j \cdot \hat{v}_t^j = \hat{v}_t^j$.

$\tilde{x}_t^j \implies$ (EX-POST IR) In particular, for $i \neq j$, the payment $\tilde{p}_t^i = 0$ by construction (6), hence (EX-POST IR) holds.

APPROX IC:

Next, we prove that $\tilde{\mathcal{M}}$ is (APPROX IC). We first upper bound one's expected cumulative utility under $\tilde{\mathcal{M}}$.

One key observation is that in Mechanism 3.4, for each buyer i and any fixed $\hat{F}_{1..T}^i, \hat{v}_{1..T}^i, F_{1..T}^i$, even if buyer i is using the best strategy in $\tilde{\mathcal{M}}$, her expected cumulative utility $\mathbb{E}[\tilde{U}_*^i]$ under the credit account mechanism $\tilde{\mathcal{M}}$ cannot be much larger than $\mathbb{E}[U_*^i]$ under the original mechanism \mathcal{M} while using the same strategy. Formally,

$$\mathbb{E}[\tilde{U}_*^i] \leq \mathbb{E}[U_*^i] + q_T. \quad (\star)$$

To show this, note that by the construction (5) and (6), the cumulative payments under $\tilde{\mathcal{M}}$ of each buyer should be close to her expected cumulative payments under \mathcal{M} :

$$\sum_{t \in [T]} \tilde{p}_t^i = c_{T+1}^i + \sum_{t \in [T]} \alpha_t^i \cdot \tilde{p}_t^i.$$

In the meanwhile, the credit account mechanism, in fact, guarantees that $c_{t+1}^i + q_t \geq 0$. Because according to (5), $c_{t+1}^i = c_t^i + \tilde{p}_t^i - \alpha_t^i \cdot \tilde{p}_t^i \geq c_t^i - \alpha_t^i \cdot \mathbb{E}[\hat{F}_t^i] \geq c_t^i - (c_t^i + q_t) = -q_t$ where the last inequality is from (1) that $\alpha_t^i \cdot \mathbb{E}[\hat{F}_t^i] = \min\{c_t^i + q_t, \mathbb{E}[\hat{F}_t^i]\}$. Therefore, we have

$$\sum_{t \in [T]} \tilde{p}_t^i = c_{T+1}^i + \sum_{t \in [T]} \alpha_t^i \cdot \tilde{p}_t^i \geq -q_T + \sum_{t \in [T]} \alpha_t^i \cdot \tilde{p}_t^i.$$

Then $\mathbb{E}[\tilde{U}_*^i]$ under $\tilde{\mathcal{M}}$ can be bounded:

$$\begin{aligned} \mathbb{E}[\tilde{U}_*^i] &= \mathbb{E} \left[\sum_{t \in [T]} v_t^i \cdot \tilde{x}_t^i - \tilde{p}_t^i \right] \\ &\leq \mathbb{E} \left[\sum_{t \in [T]} \alpha_t^i \cdot v_t^i \cdot x_t^i - \alpha_t^i \cdot \tilde{p}_t^i \right] + q_T \\ &= \sum_{t \in [T]} \mathbb{E} \left[\alpha_t^i (v_t^i \cdot x_t^i - \tilde{p}_t^i) \right] + q_T \\ &= \sum_{t \in [T]} \mathbb{E} \left[\alpha_t^i \mathbb{E}_{v_t^i} [v_t^i \cdot x_t^i - \tilde{p}_t^i] \right] + q_T, \end{aligned}$$

where the last equation is because α_t^i, x_t^i , and \tilde{p}_t^i are all independent of v_t^i . In particular, since \mathcal{M} is EX-ANTE IR, $\mathbb{E}_{v_t^i} [v_t^i \cdot x_t^i - \tilde{p}_t^i] \geq 0$. Combining with $\alpha_t^i \leq 1$, we have

$$\mathbb{E}[\tilde{U}_*^i] \leq \sum_{t \in [T]} \mathbb{E}[v_t^i \cdot x_t^i - \tilde{p}_t^i] + q_T = \mathbb{E}[U_*^i] + q_T \implies (\star).$$

On the other hand, since \mathcal{M} is dominant-strategy IC, by (DIC), we know that $\mathbb{E}[U_*^i]$ cannot be more than the utility of truthfully reporting $\mathbb{E}[U^i]$, hence

$$\mathbb{E}[\tilde{U}_*^i] \leq \mathbb{E}[U_*^i] + q_T \leq \mathbb{E}[U^i] + q_T. \quad (\star\star)$$

We then provide a lower bound on one's expected cumulative utility in $\tilde{\mathcal{M}}$ when reporting truthfully, denoted as $\mathbb{E}[\tilde{U}^i]$.

According to the proof of Theorem 4.1, we know that with high probability the credit c_t^i is always close to zero, i.e.,

$$\Pr[\forall t, c_t^i + q_t \geq 0 \text{ and } c_{T+1}^i \leq q_T] \geq 1 - \delta.$$

In this case, $\forall t \in [T], \alpha_t^i = 1$ and hence buyer i 's allocations are identical with truthfully reporting under the original mechanism \mathcal{M} , $\tilde{x}_{1..T}^i = x_{1..T}^i$. Again, by (5) and (6),

$$\sum_{t \in [T]} \tilde{p}_t^i = c_{T+1}^i + \sum_{t \in [T]} \tilde{p}_t^i \leq q_T + \sum_{t \in [T]} \tilde{p}_t^i,$$

where the last inequality is implied by $\alpha_T^i = 1$.

Hence, in this case,

$$\begin{aligned} \mathbb{E}[\tilde{U}^i] &= \mathbb{E} \left[\sum_{t \in [T]} v_t^i \cdot \tilde{x}_t^i - \tilde{p}_t^i \right] = \mathbb{E} \left[\sum_{t \in [T]} v_t^i \cdot x_t^i - \tilde{p}_t^i \right] \\ &\geq \mathbb{E} \left[\sum_{t \in [T]} v_t^i \cdot x_t^i - \tilde{p}_t^i \right] - q_T = \mathbb{E}[U^i] - q_T. \end{aligned}$$

Otherwise, if $\exists t \in [T]$ such that $\alpha_t^i < 1$ or $c_{T+1}^i > q_T$, since $\tilde{\mathcal{M}}$ is ex-post individual rational, by (EX-POST IR), buyer i 's cumulative utility must be non-negative.

Combining the cases above, we conclude that

$$\begin{aligned} \mathbb{E}[\tilde{U}^i] &\geq (1 - \delta)(\mathbb{E}[U^i] - q_T) \\ \implies \mathbb{E}[\tilde{U}^i] &\geq \mathbb{E}[U^i] - \left(\frac{\delta}{1 - \delta} \mathbb{E}[\tilde{U}^i] + (1 - \delta)q_T \right). \end{aligned}$$

Applying the upper bound ($\star\star$),

$$\mathbb{E}[\tilde{U}^i] \geq \mathbb{E}[\tilde{U}_*^i] - \left(\frac{\delta}{1 - \delta} \mathbb{E}[\tilde{U}^i] + (2 - \delta)q_T \right).$$

By choosing $\delta = 1/(T + 1)$, q_T is still in $\tilde{O}(\sqrt{T})$, and $\delta/(1 - \delta) = 1/T$. Since $\mathbb{E}[\tilde{U}^i] \leq T\bar{v}$, we have

$$\mathbb{E}[\tilde{U}^i] \geq \mathbb{E}[\tilde{U}_*^i] - (2q_T + \bar{v}).$$

Therefore, the credit account mechanism $\tilde{\mathcal{M}}$ is approximately truthful with $\epsilon = 2q_T + \bar{v} = \tilde{O}(\sqrt{T})$. \square

5 Impossibility Result

As a complement to the constructed APPROX IC and EX-POST IR credit account mechanism, the third main result shows that DIC and EX-POST IR cannot be achieved simultaneously, unless the mechanism is trivial, i.e., the allocation and payment are constant functions.

Theorem 5.1. *No non-trivial mechanism that achieves dominant-strategy incentive compatibility (DIC) and ex-post individual rationality (EX-POST IR) at the same time.*

Proof. To satisfy (DIC), reporting truthfully must be the best action of a buyer for any realization of $F_{1..T}$ and $v_{1..T}$.

Let \mathcal{M} be any EX-POST IR mechanism. Consider any period t . Suppose i is the winner in period t . For any $t' \in (t, T]$, let $F_{t'}^i$ be the distribution that $v_t^i = 0$ with probability 1, while other buyers have positive expected values.

Therefore, buyer i has no incentive to win after period t . So buyer i 's best action is always to report 0 after getting the item in period t , which results in 0 payment.

Thus, in order for \mathcal{M} to be DIC, \mathcal{M} must charge the winner 0 at period t . Otherwise the above realization of $F_{t'}^i$ becomes a counter-example. Since t is any period, it follows that \mathcal{M} charges the winner 0 in every period. So any mechanism that satisfies the two properties simultaneously must be a trivial mechanism. \square

Acknowledgements

The authors thank the anonymous reviewers for their helpful comments. The authors benefit from discussions with Balasubramanian Sivan and Renato Paes Leme. This paper is supported in part by the National Natural Science Foundation of China Grant 61561146398, a China Youth 1000-talent Program, an Alibaba Innovative Research Program, the Shanghai Sailing Program (Grant No. 18YF1407900) and the Fundamental Research Funds for the Central Universities.

References

- [Agarwal *et al.*, 2009] Nikhil Agarwal, Susan Athey, and David Yang. Skewed bidding in pay-per-action auctions for online advertising. *The American Economic Review*, 99(2):441–447, 2009. 2
- [Amin *et al.*, 2013] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. In *NIPS*, pages 1169–1177, 2013. 2
- [Amin *et al.*, 2014] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. In *NIPS*, pages 622–630, 2014. 2
- [Bachrach *et al.*, 2014] Yoram Bachrach, Sofia Ceppi, Ian A Kash, Peter Key, and David Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *EC’14*, pages 75–92. ACM, 2014. 2
- [Balseiro *et al.*, 2017] Santiago Balseiro, Max Lin, Vahab Mirrokni, Renato Leme, and Song Zuo. Dynamic revenue sharing. In *NIPS*, pages 2681–2689, 2017. 2
- [Bergemann and Morris, 2013] Dirk Bergemann and Stephen Morris. Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308, 2013. 2
- [Bergemann and Välimäki, 2017] Dirk Bergemann and Juuso Välimäki. Dynamic mechanism design: An introduction. 2017. 2
- [Cai *et al.*, 2012] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *STOC*, pages 459–478. ACM, 2012. 2
- [Daskalakis *et al.*, 2013] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *EC’13*, pages 269–286. ACM, 2013. 2
- [Edelman *et al.*, 2007] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007. 1, 2
- [Epasto *et al.*, 2018] Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Song Zuo. Incentive-aware learning for large markets. In *WWW*, 2018. 2
- [Hajaj *et al.*, 2015] Chen Hajaj, John P Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *AAAI*, pages 921–928, 2015. 2
- [Hartline and Roughgarden, 2009] Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *EC’09*, pages 225–234. ACM, 2009. 2
- [Hu *et al.*, 2015] Yu Hu, Jiwoong Shin, and Zhulei Tang. Incentive problems in performance-based online advertising pricing: cost per click vs. cost per action. *Management Science*, 62(7):2022–2038, 2015. 2
- [Kanoria and Nazerzadeh, 2014] Yash Kanoria and Hamid Nazerzadeh. Dynamic reserve prices for repeated auctions: Learning from bids. In *International Conference on Web and Internet Economics*, pages 232–232. Springer, 2014. 2
- [Klemperer, 1998] Paul Klemperer. Auctions with almost common values: The wallet game and its applications. *European Economic Review*, 42(3):757–769, 1998. 2
- [Liu *et al.*, 2014] Yuan Liu, Jie Zhang, Han Yu, and Chunyan Miao. Reputation-aware continuous double auction. In *AAAI*, pages 3126–3127, 2014. 2
- [Mirrokni *et al.*, 2016a] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Dynamic auctions with bank accounts. In *IJCAI*, 2016. 2
- [Mirrokni *et al.*, 2016b] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Optimal dynamic mechanisms with ex-post ir via bank accounts. *arXiv preprint arXiv:1605.08840*, 2016. 2
- [Mirrokni *et al.*, 2018a] Vahab Mirrokni, Renato Paes Leme, Rita Ren, and Song Zuo. Dynamic mechanism design in the field. In *WWW*, 2018. 2
- [Mirrokni *et al.*, 2018b] Vahab Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Non-clairvoyant dynamic mechanism design. In *EC’18*. ACM, 2018. 2
- [Myerson, 1981] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981. 2
- [Nazerzadeh *et al.*, 2013] Hamid Nazerzadeh, Amin Saberi, and Rakesh Vohra. Dynamic pay-per-action mechanisms and applications to online advertising. *Operations Research*, 61(1):98–111, 2013. 2
- [Shen and Tang, 2017] Weiran Shen and Pingzhong Tang. Practical versus optimal mechanisms. In *AAMAS*, pages 78–86, 2017. 2
- [Spencer, 2007] Stephan Spencer. Google deems cost-per-action as the ‘holy grail’. *CNET News*, 2007. 1
- [Tang and Wang, 2017] Pingzhong Tang and Zihe Wang. Optimal mechanisms with simple menus. *Journal of Mathematical Economics*, 69:54–70, 2017. 2
- [Tang and Zeng, 2018] Pingzhong Tang and Yulong Zeng. The price of prior dependence in auctions. In *EC’18*. ACM, 2018. 2
- [Varian, 2007] Hal R Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007. 2
- [Wang and Tang, 2014] Zihe Wang and Pingzhong Tang. Optimal mechanisms with simple menus. In *EC’14*, pages 227–240. ACM, 2014. 2
- [Yao, 2014] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *SODA*, pages 92–109. SIAM, 2014. 2
- [Yao, 2017] Andrew Chi-Chih Yao. Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *EC’17*, pages 3–20. ACM, 2017. 2