

Vulnerabilities of Single-Round Incentive Compatibility in Auto-bidding: Theory and Evidence from ROI-Constrained Online Advertising Markets*

Juncheng Li and Pingzhong Tang

Tsinghua University

{lijuncheng13, kenshinping}@gmail.com

Abstract

Most of the work in the auction design literature assumes that bidders behave rationally based on the information available for every *individual* auction, and the revelation principle enables designers to restrict their efforts to *incentive compatible* (IC) mechanisms. However, in today’s online advertising markets, one of the most important real-life applications of auction design, the data and computational power required to bid optimally are only available to the platform, and an advertiser can only participate by setting performance objectives and constraints for its proxy *auto-bidder* provided by the platform. The prevalence of auto-bidding necessitates a review of auction theory. In this paper, we examine the markets through the lens of ROI-constrained value-maximizing campaigns. We show that second price auction exhibits many undesirable properties (computational hardness, non-monotonicity, instability of bidders’ utilities, and interference in A/B testing) and loses its dominant theoretical advantages in single-item scenarios. In addition, we make it clear how IC and its runner-up-winner interdependence contribute to each property. We hope that our work could bring new perspectives to the community and benefit practitioners to attain a better grasp of real-world markets.

1 Introduction

Auto-bidding has become a corner stone of modern advertising markets. For better end-to-end performance and customer experience, platforms now provide algorithmic agents to set fine-grained bids for advertisers, who only need to submit campaign-level optimization objectives and constraints. As a result, the community has seen a surge of publications on auto-bidding in recent years.

One of the most notable features of the auto-bidding paradigm is the change of roles played by advertisers and platforms. Traditionally advertisers are assumed to bid rationally for each individual ad slot, and thus the real-time bid-

ding (RTB) literature focuses on developing algorithms for advertisers to maximize their objectives subject to different constraints and auction rules, while the auction design literature, anticipating the best response of advertisers or RTB algorithms, explores new auction rules to optimize various goals of the platform. However, in auto-bidding markets, most of the technical components are under the management of the platform: auto-bidders provided by the platform will compete with each other, based on valuations predicted by the platform, under auction rules that are also designed by the platform. Such a greater control over the market imposes a more diverse set of requirements upon the designer. The auction design literature has long been focusing on incentive compatibility and welfare/revenue guarantee, but the desiderata of an auto-bidding mechanism are far beyond these two. Google AdSense’s *partial* shift at 2021 from second to first price auction¹ might serve as an exemplary demonstration of this perplexity.

Despite the evolution of ecosystems, first and second price auction remain as the dominant mechanisms used in practice. In this paper, we study the mathematical model abstracted from real-world first and second price auction markets with ROI-constrained auto-bidders. In addition to theoretic interests, our work are also motivated by real-world observations. Traditional interpretations of some phenomena may confuse both sides of the market and lead to business choices detrimental in the long run. Our goal is thus to develop a deeper understanding of auto-bidding at the market scale, and provide practitioners a more holistic view to facilitate decision making. We will show, through a series of theoretical and empirical results, that the dominant advantage of second price auction over first price is, in many regards, reversed in the world of auto-bidding. It will be made clear throughout the journey how IC, instead of simplifying the reasoning of both bidders and auctioneers, unnecessarily complicates the game in a profound way.

1.1 ROI-Constrained Auto-bidding Markets and Related Works

Previously much effort was spent on the study of auto-bidders with budget constraints [Karande *et al.*, 2013; Charles *et*

¹In November 17, 2021, Google moves the AdSense auction for Content, Video, and Games from second price auction to first price, while keeping Search and Shopping as before [Google, 2021].

*The full version is at <https://arxiv.org/abs/2210.06107>.

et al., 2013; Balseiro and Gur, 2019; Gao and Kroer, 2020; Chen *et al.*, 2021a; Balseiro *et al.*, 2021b; Conitzer *et al.*, 2022a; Conitzer *et al.*, 2022b], but Return-on-Investment (ROI), another widely adopted auto-bidding option, received much less attention until very recently [Golrezaei *et al.*, 2021; Balseiro *et al.*, 2021c; Deng *et al.*, 2021]. For campaigns with ROI-constraints, advertisers should submit either a target ROI, a target Return-on-Ad-Spend or a target Cost-per-Action (tROI/tROAS/tCPA).² The objective of the auto-bidder is then to maximize the acquired value while keeping the average spend for each unit of value below the target threshold. ROI-constrained auto-bidders are dominating in market share in many regions of the world. There is also recent empirical evidence [Golrezaei *et al.*, 2021] from Google AdX that advertisers are indeed ROI-constrained.

An advertiser's value for an ad slot (to which we will refer as a generic *good* throughout the rest of the paper) is typically given by the product of the conversion-rate (predicted by the platform) and the value of each conversion. For ROI-constrained campaigns, the latter part is the amount of money that the advertiser is willing to pay the platform for each (unit of) conversion.² Truthful bidding for each individual auction is ex-post optimal for non-constrained quasi-linear utility-maximizing bidders. But for ROI-constrained value-maximizers, it is possible to raise bids above values to win more while keeping the average spend of each conversion below the threshold. In this paper, auto-bidders will take values as given and be restricted to the *multiplicative pacing* strategy,² wherein the bids of each bidder could only be generated through scaling the values for all goods by a common multiplier of its choice. There is always a multiplier that is ex-post bidder-optimal in second price auction, and the strategy is one of the most implemented in the industry regardless of auction formats (in particular, it remains popular in first price auction³).

Table 1 positions our market model within the growing literature on auto-bidding. Existing works differ in the modeling of advertisers' valuations and utilities. Traditionally the valuation of each bidder is modeled as being drawn from a stochastic process independently of each other. We adopt a framework first studied by Conitzer *et al.* [2022b; 2022a], where the valuation is given as the (fixed) input and can be viewed as a discretized density function or a realized sample of an arbitrary joint distribution. In particular, the correlation prevalent in advertising can thus be captured. For utilities, previous research mainly focuses on budget-constraints, and our paper complements the recent line of work on ROI-constraints. Babaioff *et al.* [2021] also consider ROI, but they model bidding behaviors with respect to the marginal ROI of each individual auction rather than the average ROI across.

Deng *et al.* [2021] and its follow-up work [Balseiro *et al.*, 2021a] study a model that largely coincide with ours. Their

goal is to design mechanisms having revenue and welfare guarantees when the designer has fairly accurate signals on the valuation. Though we will also report some results on revenue and welfare, we emphasize that, for today's large scale auto-bidding systems, auction (combined with auto-bidding strategies) acts more like an efficient distributed algorithm to match demand with supply and compute market-clearing prices. With this mentality, our work covers a broader range of properties and focuses heavily on their possible practical impacts on advertisers and platforms. We also differ in the adopted solution concept. We allow fractional allocation and incorporate the tie-breaking rule into the solution concept, which is not only well-motivated, but also guarantees the existence of equilibrium even though the market is discrete and discontinuous. In contrast, Balseiro *et al.* [2021a] break ties lexicographically and more complex auctions like VCG and GSP are considered there. So they choose a weaker solution concept called undominated bids and avoid the discussion of existence. Nonetheless, this distinction diminishes for large markets as the result of a single auction becomes negligible.

Auto-bidding or RTB algorithms have been studied for a long time. Such works typically assume a *stationary* environment and optimize various objectives for a *single* advertiser. They fail to capture other bidders' responses invoked by the action of the focal agent, and the resulting equilibrium outcome may not fulfill the initial design goal if all the bidders implement the same strategy (see, e.g., Appendix L.1). One notable exception is the work by Aggarwal *et al.* [2019], who were aware of this problem and tried to prove the existence of an equilibrium. But their treatment of equilibrium is incomplete (see a discussion in Appendix C).

Researchers have shown that the computation of equilibrium bears some inherent hardness with budget-constrained multiplicative pacing bidders. For computing any equilibrium, we improve previous result to show that it is PPAD-hard to approximate within *constant* parameters (for budget-constraints, it was shown to be hard to approximate within *polynomially small* parameters [Chen *et al.*, 2021a]). Our result is built on a more concise reduction quite distinct from the previous one. We also improve the NP-hardness result of optimizing revenue/welfare [Conitzer *et al.*, 2022b] to APX-hardness. Moreover, the source of these hardness and the vulnerabilities of IC are demonstrated more clearly with our constructions, which we believe could help both researchers and practitioners better extrapolate our techniques and insights.

1.2 Single-round Incentive Compatibility

The classic revelation principle ensures us that, when designing single-item auctions, any implementable allocation rule could be implemented in an incentive compatible way by directly eliciting bidders' private information (in most cases only the valuations to the item to be sold are needed). By focusing on IC mechanisms, the designer loses nothing while bidders could be prevented from strategic behaviors. In comparison, the characterization and computation of equilibrium in first price auction is notoriously hard (see, e.g., [Filos-Ratsikas *et al.*, 2021] and the survey therein).

Though second price auction is only IC for single-item auctions, it possesses another fascinating property in auto-

²See Appendix B for more notes on terminology.

³Possibly surprisingly, we will show in Section 6 that it will *not* bring incentive issues for first price auction and it is in the interests of a platform to enforce so. In contrast, both bidders and sellers have incentives to deviate from it in second price auction (Appendix J.3).

Utility model	Valuations	
	Deterministic, correlated	Stochastic, independent
ROI-constrained value-maximizer	Our model, Balseiro et al. [2021a]	Balseiro et al. [2021c], Golrezaei et al. [2021]
Budget-constrained utility-maximizer	Conitzer et al. [2022a; 2022b], Chen et al. [2021a], Gao and Kroer [2020]	Balseiro and Gur [2019], Balseiro et al. [2021b]
ROI&budget-constrained value-maximizer	Deng et al. [2021], Aggarwal et al. [2019]	

Table 1: Common auto-bidder types and valuation assumptions in the literature.

bidding markets: from a single bidder’s perspective, each individual auction comes with a *winning price* independent of its own bid (which is essentially an equivalent statement of IC). As a consequence, the (ex-post/offline) task of each auto-bidder is a linear program for bidders with linear constraints like ROI and budget. The optimal solution can be well approximated by the multiplicative pacing strategy, which is simple to implement and performs well even in the online setting [Balseiro and Gur, 2019; Balseiro *et al.*, 2022]. This provides second price auction an illusory strategyproofness that could be called *ex-post IC*: every advertiser could truthfully report its tROI and happily accept the equilibrium bids given by its proxy auto-bidder as deviating unilaterally will not bring extra profit *at the moment*.

Nonetheless, even long before the advent of the auto-bidding era, experiments have revealed that bidders’ behaviors in second price auction are far from truthful (see, e.g., [Kagel and Levin, 2011]). Recall that, in second price auction, the price is *set* by the runner-up, but *paid* by the winner. In some senses, both the winner and the runner-up care *little* about the absolute magnitudes of their *own* bids: only the ranking (first and second) is important. This can also be seen from the linear program optimizing utility for the auto-bidder (see, e.g., [Aggarwal *et al.*, 2019]), where no decision variable denoting bids appears and the solution only prescribes whether each auction should be won or not. However, the bid of the runner-up always means *a lot* to the winner. A well-known enemy of IC due to this *runner-up-winner interdependence* is externality, i.e., the utility of a bidder depends on the allocation and payment of not only itself, but also others.

It should not be surprising that externality makes manipulation worthwhile since IC is not designed for the job. In our auto-bidding markets, the objective of each auto-bidder is defined clearly without externality.⁴ However, auto-bidders adjust bids based on the overall performance across *all* auctions. Even though at equilibrium changing bids unilaterally could not benefit the manipulator immediately, it may trigger a cascade of responses that shift the whole market state. This creates a kind of externality that is *internalized* into the outcome of the shifted equilibrium. It is worthwhile to point out that, for some results in this paper, it seems to be the multiplicative pacing strategy that leads to a property. Actu-

⁴Another common scenario where IC fails is that bidders do not have complete knowledge of their valuations. This is not an issue either in our case.

ally it is irrelevant of the specific bidding strategy or even the ROI-constraint: for simultaneous auctions with single-round IC, bidders will always bid strategically across all auctions,⁵ which is enough to establish those results.

Letting the (bids of) opponents determine the payment of the winner is the key to establish IC for single-item auctions, but it is also the key to open the Pandora’s box in auto-bidding markets, as will be detailed in the remainder of this paper.

1.3 Contributions

In this paper, we consider the auto-bidding market where ROI-constrained value-maximizing bidders compete with each other in simultaneous auctions. Our main focus is second price auction: behaviors of auto-bidders within first price auction markets will be discussed at the very end.

We start by formulating our own solution concept of the market (and its approximate version), named *auto-bidding equilibrium*, since a pure Nash equilibrium may not always exist. The equilibrium reasonably captures the expected steady-state that the auto-bidders intend to reach collectively, and its guaranteed existence puts our study on a solid theoretical footing. Our most important results are as follows:

- It is PPAD-hard to find an approximate auto-bidding equilibrium within constant parameters.
- It is APX-hard to optimize revenue or welfare over all auto-bidding equilibria.
- Non-monotonicity: an advertiser who raises its tROI/tROAS (or equivalently, lowers the tCPA) at the equilibrium could end up with a higher revenue after deviation (through a natural equilibrium transition process to resolve multiplicity).

Besides their own significance, the constructions used in establishing these results reveal many crucial structures of the market and highlight the role played by IC. They serve as the foundation to comprehend and interpret other characteristics of the market.

We proceed to explore several practical concerns of great consequence. For advertisers, we show that the market suffers severe utility instability and input sensitivity issues. For platforms, we demonstrate that biases exist broadly in the widely-used A/B testing. Finally, we give a comprehensive

⁵Even a non-constrained utility-maximizer could manipulate if one of its opponents has constraints. Truthful reporting is secure only when all the other bidders are insensitive to each other’s bid.

comparison between first and second price auction, and show that first price auction is generally exempted from the above undesirable properties and performs better. As an application of our results, we give our guess about why Google AdSense moves from second to first price auction in a partial manner.

2 Markets and Equilibria

We consider a market where a set of bidders $N = \{1, \dots, n\}$ compete for a set of divisible goods $M = \{1, \dots, m\}$. Without loss of generality, the tROIs of all bidders are set to zero (equivalently, tROAS of one),⁶ i.e., each bidder's spend should be no more than its acquired value. We use $v_{i,j}$ to denote the value of bidder i to good j . For each good j , there is at least one bidder i such that $v_{i,j} > 0$. The platform simultaneously runs a single-item second price auction for every good. Auto-bidders are restricted to apply multiplicative pacing strategies: the action space of bidder i is the set of undominated multipliers $\alpha_i \in [1, A]$,⁷ and its bid for good j is $\alpha_i v_{i,j}$. Multiplicative pacing is ex-post bidder-optimal, as shown in Proposition 1. The omitted proof follows from a linear programming formulation of the ROI-constrained value maximization problem, and the result is widely known in both the literature and the industry.

Proposition 1. *Suppose that bidders can bid arbitrarily across auctions. Holding all other bidders' bids, each bidder has a best response wherein bids are generated by scaling its valuations of all goods by a uniform multiplier, given that it could freely choose to win any fraction of a good of which it is a tied winner.*

To complete the picture, one's first intuition may be to specify a tie-breaking rule, make the allocation $x \in [0, 1]^{n \times m}$ (where $x_{i,j}$ is the fraction of good j allocated to bidder i) uniquely determined by α , define bidder's utility as the ROI-constrained valuation:

$$u_i(\alpha) = \begin{cases} \sum_j x_{i,j} v_{i,j}, & \text{if } \sum_j x_{i,j} p_j \leq \sum_j x_{i,j} v_{i,j}; \\ -\infty, & \text{otherwise;} \end{cases}$$

and study the pure Nash equilibrium (PNE) of the game. However, in Appendix D, we will give an example where, no matter how ties are broken, a PNE does not exist. On closer inspection, we find that the steady-state of the market can be captured instead by the solution concept defined below.

Definition 1. *An auto-bidding equilibrium (α, x) consists of multipliers $\alpha \in [1, A]^n$ and allocation $x \in [0, 1]^{n \times m}$ such that*

- *bidders with the highest bid win the good: if $x_{i,j} > 0$, $\alpha_i v_{i,j} = \max_k \alpha_k v_{k,j}$ for all i, j ;*
- *winner pays the second price: if $x_{i,j} > 0$, then $p_j = \max_{k \neq i} \alpha_k v_{k,j}$ for all i, j ;*
- *full allocation of goods: $\sum_i x_{i,j} = 1$ for all j ;*

⁶See Appendix D on why this is WLOG. Later we may sometimes change the tCPA of a bidder by a factor λ , by which we mean all the valuations of this bidder is scaled by λ .

⁷Fixing other bidders' bids, $\alpha_i = 1$ always weakly dominates any $\alpha_i < 1$. The cap A can safely be replaced by $+\infty$ for most cases. See more discussion in Appendix D.

- *ROI-feasible: $\sum_j x_{i,j} p_j \leq \sum_j x_{i,j} v_{i,j}$ for all i ;*
- *maximal pacing: unless $\alpha_i = A$, $\sum_j x_{i,j} p_j = \sum_j x_{i,j} v_{i,j}$ for all i .*

The auction rules and ROI-constraints are directly imposed by the first four conditions, while the best responses among bidders are encoded in the maximal pacing condition in a less straightforward way. To see this, note that, from bidder i 's perspective, given the winning price of each good, α_i acts as a marginal-ROI threshold: it will win all goods in whole with a marginal ROI strictly larger than $\frac{1}{\alpha_i} - 1$ and lose those with ROI strictly lower. At any auto-bidding equilibrium, if $\alpha_i > 1$, α_i (and the resulting vector of bids with $b_i = \alpha_i v_{i,j}$) is exactly a best response since the marginal ROI of any good lost or tied is strictly lower than zero, and winning anymore will definitely violate the ROI-constraint. If, however, $\alpha_i = 1$ and bidder i is only allocated a fraction of some good as in the example given in Appendix D, α_i is technically not a best response (for the normal form game) but the maximal pacing condition is still satisfied. For such bidders, their opponents could always oscillate their multipliers around the equilibrium point to achieve the corresponding stable allocation.

Theorem 1. *An auto-bidding equilibrium always exists.*

The proof is in Appendix F. In addition, the definition and existence result can be extended to incorporate *reserve prices* and *additive boosts* (see Appendix E), both of which are common practice in the industry.

The relaxed version of the equilibrium, named (η, δ) -approximate auto-bidding equilibrium, is defined as follows:

- bidders with bids close enough to the highest can win the good: if $x_{i,j} > 0$, $\alpha_i v_{i,j} \geq (1 - \eta) \max_k \alpha_k v_{k,j}$ (if there is a reserve price r_j , winner's bid should also be no less than $(1 - \eta)r_j$);
- winner pays the second price (even if it is the bidder with the second highest bid; if there is a reserve price, the price is the maximum of the reserve price and the second price);
- full allocation of goods;
- approximately ROI-feasible: $\sum_j x_{i,j} p_j \leq (1 + \delta) \sum_j x_{i,j} v_{i,j}$ for all i ;
- approximately maximal pacing: unless $\alpha_i = A$, $\sum_j x_{i,j} p_j \geq (1 - \delta) \sum_j x_{i,j} v_{i,j}$ for all i .

3 Computational Complexities

In this section, we demonstrate two different kinds of intractability or unpredictability of the market. Our first result shows that, in general, it is hard for a market to reach a stable state. But even for cases where an equilibrium is easy to achieve, the difference among equilibria may be large, and the second result tells us that, in general, it is hard to determine how large the difference is.

At a high level, the interdependence that the (bid of) runner-up determines the payment of the winner endows the market with a structure similar to Boolean operators (electronic components and conductive wires). In digital electronics, a functionally complete set of operators can be assembled

goods	w_1	w_2	w_3	$\underline{u}$	$\bar{u}$
input bidder w_1	1/14				
input bidder w_2		1/14			
input bidder w_3			1/14		
gate bidder u	1/3	1/3	1/3	1/2	1/4
reserve price	1/7	1/7	1/7	1	1

 Table 2: Valuation profile to encode gate u .

to compute any Boolean function. In our model, it is PPAD-hard to find an equilibrium since it can encode any stable state of a circuit that is continuous. When optimizing revenue or welfare, the circuit structure collapses to discrete choices that correspond to equilibrium selection, and the problem becomes NP-hard (and APX-hard since the correspondence is almost exact). Note that hardness does not only exist in the family of instances constructed in the reduction: for a generally hard problem, it is possible but usually non-trivial to identify a meaningful subset of instances that are computationally tractable.

Our focus here is the many complex structures brought into the market by IC. Nonetheless, to explore equilibrium properties quantitatively, we develop two algorithms (see Appendix I for details) to compute equilibrium. Their own properties are beyond the scope of this paper.

3.1 Complexity of Finding Any Equilibrium

Theorem 2. *It is PPAD-hard to find an (η, δ) -approximate auto-bidding equilibrium for some constant $\eta, \delta > 0$.*

The full proof is in Appendix G. We will use equilibria in a properly constructed market to encode feasible states of a circuit, which is one of the most fundamental and frequently used objects in complexity theory [Chen *et al.*, 2009; Rubinstein, 2018]. Papadimitriou and Peng [2021] show that a circuit consisting only of a continuous version of NAND (NOT-AND) gate is enough to capture PPAD-hardness, in analogy to the *functional completeness* of NAND gate in digital electronics. The continuous gate computes the function that sums all (at most 3) inputs and inverts the result: for a gate u , given all the input values y_w of its incoming gates $w \in N_u$, its own value y_u should satisfy that

$$y_u \in \begin{cases} \{0\}, & \text{if } \sum_{w \in N_u} y_w > 0.5; \\ \{1\}, & \text{if } \sum_{w \in N_u} y_w < 0.5; \\ [0, 1], & \text{otherwise.} \end{cases}$$

An assignment y of values to gates is feasible if the above constraints are satisfied for all gates. The key construction of our reduction is given in Table 2. (This is a simplified version: in the full proof we will remove the reliance on reserve prices and take approximation into account.) We will associate each gate u with a bidder of the same name and use its multiplier α_u to encode the gate's value y_u (with $y_u = \alpha_u/2 - 1$). The valuation profiles are calibrated such that bidder u would always win all input goods w_1, w_2, w_3 but at varying prices, determined by (multipliers of) their corresponding input bidders. The correspondence between bidder u and gate u is almost straightforward:

- If the sum of input prices is too high, to satisfy its ROI-constraint, bidder u could not even win good $\underline{u}$ in whole and thus won't raise α_u above 2.
- If the sum is too low, good $\bar{u}$ is required to satisfy bidder u 's appetite for more valuation, which forces α_u to be fixed at 4.
- If input prices sum to exactly 0.5, bidder u can freely choose α_u as long as it is allocated $\underline{u}$ in whole but not $\bar{u}$ at all.

Intuitively, the hardness of finding a feasible state of a circuit lies in the coordination among the interconnected gates. Given any value assignment, for each unsatisfied gate u , the corresponding α_u in the constructed market is either too large (input goods are too expensive and ROI-constraint is violated) or too small (input goods are too cheap and bidder u has the incentive to win $\bar{u}$). Consider hypothetically that we apply a naive search algorithm⁸ where, for each non-equilibrium assignment y , we choose some unsatisfied u , lower α_u by a small amount if it is too large, and raise it if it is too small. As we adjust α_u (or y_u) for a bidder (gate), the payments (sums of inputs) of its outgoing neighbors change accordingly, which may also change their directions of adjustment (e.g., a bidder goes from ROI-feasible to infeasible). As gates can be assembled arbitrarily, we can imagine how hard it is to find a state that satisfies the constraints of all bidders/gates.

In retrospect, IC requires the payment of the winner to be determined *externally* by other bidders. As a result, nothing is local in the market and we can connect bidders in a way that encodes any circuit perfectly. You may think that the market constructed in the reduction can be simplified, e.g., by setting a reserve price for the input good such that its price would no longer be affected by the input bidder and an equilibrium could be easier to find. However, this requires much knowledge of the specific market *a priori* such as an upper bound of a bidder's multiplier, which is typically impossible. In general markets, the aforementioned chasing behavior is even more complex: e.g., if a gate bidder lowers its multiplier, the consequence is simply lowering payments for its outgoing neighbors and increase their ROIs, but in general markets it may also lose goods it previously won and instead decrease ROI for the opponent who now wins the item with a negative marginal ROI. See the non-monotonicity instance in Section 4 for an example of this complex cascading phenomenon.

3.2 Complexity of Finding Revenue or Welfare Optimal Equilibrium

Theorem 3. *It is APX-hard to find the optimal revenue or welfare over all auto-bidding equilibria.*

⁸Theoretically, any problem in PPAD can be reduced to the generic End-Of-The-Line problem (by which the class is defined), where we are given (1) a directed graph consisting solely of non-intersecting directed paths (lines) and (2) a vertex with no predecessor (the start of a line). The task is to find a vertex with no successor (the end of a line). There is a natural algorithm (inevitably inefficient if $\text{PPAD} \neq \text{P}$) that simply searches along any path. The search procedure we depict here shares a similar spirit, but it may (or may not) circulate and is only used to give some intuition.

valuation	good 1	good 2	goods that both bidders will never win
bidder 1	2	1	ϵ
bidder 2	1	2	ϵ

Table 3: A symmetric market with extremely asymmetric equilibria.

The full proof is in Appendix H. When dealing with optimization problems, besides encoding a *feasible* circuit state into equilibrium, we also need to relate the *objective* (revenue/welfare) of the market owner to one that is hard to optimize within the circuit. Interestingly, it is still the encoding of feasibility that reveals more distinctive structures of the problem.⁹ The latter step only involves simple operations (max and sum) that are naturally embedded in the auction rules (winner pays the *largest* non-winning bid) and bidder’s rationale (*aggregating* outcomes across all auctions).

The problem to be reduced is of discrete nature, and we will encode a 0-1 *choice* using equilibrium *multiplicity* as shown in Table 3. Here the market is symmetric in valuation but has equilibria that represents two extremes in allocation: bidder i (1 or 2) wins both good 1 and 2 with $\alpha_i \geq 2$ and $\alpha_{-i} = 1$. There is another equilibrium where $\alpha_1 = \alpha_2 = 2$ and each bidder gets the good it values the most, which is natural, fair and revenue-optimal within the sub-market consisting of good 1 and 2. But we will see in the full proof that it is often advantageous for the seller to choose the asymmetric equilibrium as it frequently dominates the symmetric one in revenue *of the whole market*. The reason lies in the last column of Table 3: though bidder 1 and 2 will never win those goods, they are price-setters and it is usually profitable to enforce a high multiplier on one, rather than letting them share the sub-market fairly but both bid at a moderate level.

The reduction clearly demonstrates the “externality” created by IC: each pair of bidders in the sub-market characterized by Table 3 determines the clearing prices of some other auctions they will never win (though such knowledge is hard to acquire a priori in practice). On the other hand, if the seller is able to engage in the choice of equilibrium, it may favor those unfair outcomes that, though less profitable locally, could drive up revenue from goods outside these sub-markets. See Appendix J.2 and L for further discussion on externality among sub-markets and Appendix J.3 on how sellers could prevent those unwanted equilibria by actively elevating the bid landscape without breaking single-round IC.

4 Non-monotonicity

With auto-bidding built into the mechanism, advertisers are effectively playing a meta-game through reporting tROIs. In single-item first or second price auctions, raising one’s bid will never decrease its winning probability. In real-world markets, advertisers also expect ROI monotonicity, i.e., lowering tROI/tROAS (raising tCPA) should bring them more valuation. Due to equilibrium multiplicity, however, it is not

⁹Here feasibility should be encoded in a way such that all the equilibria are known and easily enumerable. The hardness comes from finding the best (not any) one.

clear how to define utility functions for the advertiser game, let alone monotonicity. We avoid this technicality by examining a proper equilibrium transition process, from which we can see how the runner-up-winner interdependence triggers the chain reaction that is complex and counter-intuitive.

The deviation of the manipulator and the transition of equilibrium work as follows. At round 0, the manipulator i_0 changes its tCPA to a fraction $r < 1$ of the original, resulting in a valuation profile v' such that $v'_{i_0,j} = rv_{i_0,j}, \forall j$. Each bidder then applies an iterative method (see Appendix I.2) to optimize their utilities and collectively find the new equilibrium. The behavior of the algorithm is very intuitive: if the current ROI (aggregated over a moving window of recent rounds) is too high, lower the multiplier, and vice versa. To make the transition smooth, α_{i_0} is divided by r right after the tCPA modification such that the fine-grained bids of the manipulator are kept unchanged at the moment.

The detail of the non-monotone example is in Appendix J.1. Here we give a high-level description of the process. At the old equilibrium, good 1 is the only good won by i_0 , and it is shared between i_0 and another bidder i_1 . i_0 initiates the dynamics by lowering α_{i_0} , since its ROI-constraint is now violated after the update of tCPA. However, i_1 does not want to win good 1 completely, otherwise its ROI-constraint will also be violated. So i_1 lowers its multiplier as well, which further triggers the same behavior for bidder i_2 . As a result, i_0 , i_1 and i_2 reach an almost perfect coordination where the multiplicative ratios among their multipliers remain nearly constant all the way through the transition. There is another bidder, i_3 , who pays less due to the lowered second prices set by i_0 , i_1 and i_2 . Therefore it tries to win more goods by gradually raising its multiplier. During the process, i_2 and i_1 pay more for goods whose second prices are set by i_3 , and thus they have to give up goods of which they are one of the tied winners (these goods have the lowest marginal ROIs): i_2 gives up good 2 to i_1 , and i_1 wins more good 2 but loses good 1 to i_0 to balance its deficit, which contributes to the success of the manipulation of i_0 . In the end, i_3 takes a fraction of good 2 away from i_1 to bind its ROI-constraint, and i_1 compensates this by taking a fraction of good 1 from i_0 . Nonetheless, i_0 still benefits from lowering its tCPA.

5 Practical Properties of the Market

5.1 Utility Instability for Advertisers

Besides high-quality value estimation and bid optimization, platforms are also trying to serve many other needs of their clients, among which utility (i.e., the total acquired value) stability stands out because (1) advertisers expect a smooth experience, and more importantly (2) utility is the most prominent feedback on how successful their advertising campaigns are. As a result, utility instability may bring confusions and put many good campaigns at the risk of being forfeited prematurely. Our experiments (see Appendix K.1) show that, in markets generated from several different stochastic processes, a large utility gap between the worst and the best equilibrium for an advertiser is quite often to be observed, and it is fairly common that an advertiser wins nothing in some equilibrium but acquires a significant positive value in others. The gap

seems to reduce for thicker markets, but large-scale realistic instances suffer another type of instability: sensitivity to input valuations (see Appendix K.2).

Instability differs in degree market-by-market and we will give more analysis in Appendix K.3. To get a basic idea, consider a two-bidder market that is symmetric in the sense that goods appear in pair, of which one is valued v_1 and v_2 and the other is valued v_2 and v_1 by bidder 1 and 2, respectively (note that the key construction in the proof of APX-hardness shares a similar structure). There is always an equilibrium where bidder 1 wins all goods, and one where bidder 2 wins all. Depending on specific valuation profiles, there may also be many intermediate ones. From a dynamic point of view, committing a higher multiplier would make the opponent pay more, and the bidder who quits the price war first would lower its multiplier to satisfy its ROI-constraint (and also the opponent's) but lose the market share.

In addition, the above prototypical example distinguishes two sources of instability: an intensely competitive landscape and the IC property. First price auction also suffers high-sensitivity if it holds for a large percentage of goods that the values of top bidders are extremely close. However, in first price auction, competition is local and direct, i.e., bid or value perturbations only affect the auctions in which they happen. But in second price auction, any fluctuation will propagate to the whole market through the runner-up-winner interdependence and the impact is more widespread and unpredictable.

5.2 Interference in A/B Testing for Platforms

A/B testing is an indispensable tool to evaluate new technologies and assist business decisions. In a typical setup, users in experiment are randomly assigned to either a treatment or a control variant (e.g., different reserve pricing strategies), and metrics are aggregated within each group to compare and see which variant is better. The same idea can also be applied to randomize advertisers. An ideal experiment requires the Stable Unit Treatment Value Assumption (SUTVA) to hold, which generally means that there should be no interference between the treatment and the control group. Ad-side experiments (regardless of auction formats) clearly violate SUTVA since all ads compete for the same set of goods, and user-side violation (in non-auction scenarios) is also common in practice. We show empirically that an unpredictable¹⁰ bias exists broadly in naive implementations of both user-side and ad-side A/B testing in second price auction markets. The bias comes from the fact that bidders' behaviors in the (counterfactual) A/A, A/B and B/B tests are all different. Experiment details and more discussion can be found in Appendix L, where we also propose a simple approach to discerning biases and designing less-biased experiments.

¹⁰It is known that A/B testing in two-sided markets suffers from *cannibalization bias* [Blake and Coey, 2014; Liu *et al.*, 2020], which often enlarges the estimated advantages of the better variant. Such a bias may actually increase the experimental power since practitioners care more about whether a treatment is better, rather than how much better. In contrast, the interference introduced by IC is complex and it may lead to wrong decisions easily.

6 First Price Auction versus Second Price Auction: Within and Beyond Auto-bidding

With first price auction, no auto-bidding is needed and the platform simply allocates goods to bidders with the highest (tROI-discounted) valuations, which are also charged as payments. The best possible revenue (see Appendix J.3) is naturally achieved. Quasi-linear utility makes no difference for advertisers and it is a dominant strategy to report their tROIs truthfully. Even if they prefer spending less with the same acquired value, the incentive to deviate diminishes as the market becomes thicker.¹¹ First price auction does not provide advertisers an *ex-post optimal* outcome. However, advertisers should happily accept it since it is still *fair/envy-free* as it relates closely to the classic *competitive equilibrium* in Fisher markets (see Appendix N).

Besides revenue-optimality, strategyproofness and fairness, first price auction also dominates in almost all the other aspects studied in this paper: (1) market outcome is unique; (2) computation is straightforward; (3) ROI-monotone; (4) no interference among sub-markets; (5) competition is direct and local; even if the competition is so intense that the utility becomes unstable, it can easily be smoothed by actively applying small perturbations or probabilistic allocations [Borgs *et al.*, 2007]. It shares with second price auction the problem of biases in ad-side A/B testing since it is rooted in the setup itself, irrelevant to auction formats. Budget-constrained markets share similar results [Conitzer *et al.*, 2022a]: (1) the equilibrium is unique; (2) computation is convex and tractable; (3) budget-monotone. Since budget should still be paced, there remains interference among sub-markets and the competition is less direct than within ROI-constrained markets. Nonetheless, with first price auction, the market is convex and more predictable, in contrast to the complex combinatorial structure of second price auction that is difficult to deal with.

In Appendix O, we extrapolate from our model to the case where advertisers may submit tROIs for several different sub-markets (e.g., via targeting in practice). We argue that, if optional sub-markets are coarse-grained and advertisers do not have enough knowledge to differentiate them, the market could still be well captured by our model and first price auction mostly retains the upper hand. As an application of our results, we also give our guess on why Google moves from second to first price auction for only Content, Video and Games, but not Search and Shopping.

7 Conclusion

In this paper, we study a model that abstracts several most influential features of present auto-bidding markets. We try our best to pinpoint how the IC property contributes to every theoretical or empirical phenomenon such that readers could better extrapolate our results. We hope that our work could bring new perspectives to the community, and inspire practitioners to pay closer attention to the IC property and attain a better grasp of real-world markets.

¹¹This technically deviates from our model. Rigorous treatment is given in Appendix M.

References

- [Aggarwal *et al.*, 2019] Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. Autobidding with constraints. In *International Conference on Web and Internet Economics*, pages 17–30. Springer, 2019.
- [Akbarpour and Li, 2020] Mohammad Akbarpour and Shengwu Li. Credible auctions: A trilemma. *Econometrica*, 88(2):425–467, 2020.
- [Ausiello *et al.*, 2012] Giorgio Ausiello, Pierluigi Crescenzi, Giorgio Gambosi, Viggo Kann, Alberto Marchetti-Spaccamela, and Marco Protasi. *Complexity and approximation: Combinatorial optimization problems and their approximability properties*. Springer Science & Business Media, 2012.
- [Babaioff *et al.*, 2021] Moshe Babaioff, Richard Cole, Jason Hartline, Nicole Immorlica, and Brendan Lucier. Non-quasi-linear agents in quasi-linear mechanisms. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [Balseiro and Gur, 2019] Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.
- [Balseiro *et al.*, 2021a] Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Balseiro *et al.*, 2021b] Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. Budget-management strategies in repeated auctions. *Operations Research*, 2021.
- [Balseiro *et al.*, 2021c] Santiago R Balseiro, Yuan Deng, Jieming Mao, Vahab S Mirrokni, and Song Zuo. The landscape of auto-bidding auctions: Value versus utility maximization. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 132–133, 2021.
- [Balseiro *et al.*, 2022] Santiago R Balseiro, Haihao Lu, and Vahab Mirrokni. The best of many worlds: Dual mirror descent for online allocation problems. *Operations Research*, 2022.
- [Blake and Coey, 2014] Thomas Blake and Dominic Coey. Why marketplace experimentation is harder than it seems: The role of test-control interference. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 567–582, 2014.
- [Borgs *et al.*, 2007] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*, pages 531–540, 2007.
- [Charles *et al.*, 2013] Denis Charles, Deeparnab Chakrabarty, Max Chickering, Nikhil R Devanur, and Lei Wang. Budget smoothing for internet ad auctions: a game theoretic approach. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 163–180, 2013.
- [Chen *et al.*, 2009] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.
- [Chen *et al.*, 2021a] Xi Chen, Christian Kroer, and Rachitesh Kumar. The complexity of pacing for second-price auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 318–318, 2021.
- [Chen *et al.*, 2021b] Xi Chen, Christian Kroer, and Rachitesh Kumar. Throttling equilibria in auction markets. *arXiv preprint arXiv:2107.10923*, 2021.
- [Conitzer *et al.*, 2022a] Vincent Conitzer, Christian Kroer, Debmalaya Panigrahi, Okke Schrijvers, Nicolas E Stier-Moses, Eric Sodomka, and Christopher A Wilkens. Pacing equilibrium in first price auction markets. *Management Science*, 68(12):8515–8535, 2022.
- [Conitzer *et al.*, 2022b] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. Multiplicative pacing equilibria in auction markets. *Operations Research*, 70(2):963–989, 2022.
- [Dasgupta and Tsui, 2004] Sudipto Dasgupta and Kevin Tsui. Auctions with cross-shareholdings. *Economic Theory*, 24(1):163–194, 2004.
- [Debreu, 1952] Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.
- [Deng *et al.*, 2021] Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference 2021*, pages 3965–3973, 2021.
- [Deng *et al.*, 2022] Yuan Deng, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo. Efficiency of the first-price auction in the autobidding world. *arXiv preprint arXiv:2208.10650*, 2022.
- [Despotakis *et al.*, 2021] Stylianos Despotakis, R Ravi, and Amin Sayedi. First-price auctions in online display advertising. *Journal of Marketing Research*, 58(5):888–907, 2021.
- [Fan, 1952] Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences of the United States of America*, 38(2):121, 1952.
- [Filos-Ratsikas *et al.*, 2021] Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. On the complexity of equilibrium computation in first-price auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 454–476, 2021.
- [Gao and Kroer, 2020] Yuan Gao and Christian Kroer. First-order methods for large-scale market equilibrium compu-

- tation. *Advances in Neural Information Processing Systems*, 33:21738–21750, 2020.
- [Glicksberg, 1952] Irving L Glicksberg. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [Golrezaei *et al.*, 2021] Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. Auction design for roi-constrained buyers. In *Proceedings of the Web Conference 2021*, pages 3941–3952, 2021.
- [Google, 2021] Google. Faqs about adsense moving to a first-price auction. <https://support.google.com/adsense/answer/10858748#faqs>, 2021. Accessed: 2024-05-08.
- [Gürtler and Gürtler, 2012] Marc Gürtler and Oliver Gürtler. Inequality aversion and externalities. *Journal of Economic Behavior & Organization*, 84(1):111–117, 2012.
- [Ha-Thuc *et al.*, 2020] Viet Ha-Thuc, Avishek Dutta, Ren Mao, Matthew Wood, and Yunli Liu. A counterfactual framework for seller-side a/b testing on marketplaces. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2288–2296, 2020.
- [Kagel and Levin, 2011] John H Kagel and Dan Levin. Auctions: A survey of experimental research, 1995-2010. *Handbook of experimental economics*, 2:563–637, 2011.
- [Karande *et al.*, 2013] Chinmay Karande, Aranyak Mehta, and Ramakrishnan Srikant. Optimizing budget constrained spend in search advertising. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 697–706, 2013.
- [Kimbrough and Reiss, 2012] Erik O Kimbrough and J Philipp Reiss. Measuring the distribution of spitefulness. 2012.
- [Leme *et al.*, 2012] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. Sequential auctions and externalities. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 869–886. SIAM, 2012.
- [Liaw *et al.*, 2022] Christopher Liaw, Aranyak Mehta, and Andres Perlroth. Efficiency of non-truthful auctions under auto-bidding. *arXiv preprint arXiv:2207.03630*, 2022.
- [Liu *et al.*, 2020] Min Liu, Jialiang Mao, and Kang Kang. Trustworthy online marketplace experimentation with budget-split design. *arXiv preprint arXiv:2012.08724*, 2020.
- [Paes Leme *et al.*, 2020] Renato Paes Leme, Balasubramanian Sivan, and Yifeng Teng. Why do competitive markets converge to first-price auctions? In *Proceedings of The Web Conference 2020*, pages 596–605, 2020.
- [Papadimitriou and Peng, 2021] Christos Papadimitriou and Binghui Peng. Public goods games in directed networks. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 745–762, 2021.
- [Rabin, 1993] Matthew Rabin. Incorporating fairness into game theory and economics. *The American economic review*, pages 1281–1302, 1993.
- [Rubinstein, 2018] Aviad Rubinstein. Inapproximability of nash equilibrium. *SIAM Journal on Computing*, 47(3):917–959, 2018.
- [Zhang *et al.*, 2016] Weinan Zhang, Yifei Rong, Jun Wang, Tianchi Zhu, and Xiaofan Wang. Feedback control of real-time display advertising. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, pages 407–416, 2016.