Research Challenges in Internet Ad Markets: Vignettes on Complex Environments

KSHIPRA BHAWALKAR
and
GURU GURUGANESH
and
SÉBASTIEN LAHAIE
and
ANDRES PERLROTH
and
BALASUBRAMANIAN SIVAN
Google Research

Internet ad markets have been a great source of motivation and challenging problems for the Economics and Computation community. The practical design of these ad markets has also greatly benefited from advances in the Economics and Computation community. Through a series of vignettes, we aim to highlight interesting challenges in modern internet ad markets, born out of a Google perspective. In this first article, we focus on a few particular auction design and optimization questions motivated by the Sponsored Search and Display Ads markets. We hope these articles spur further progress in the research community and in industry.

Categories and Subject Descriptors: F.2.0 [**ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY**]: General

General Terms: Economics; Algorithms; Theory;

Additional Key Words and Phrases: Sponsored-Search, First-Price, Intermediaries, Ad Auctions

1. INTRODUCTION

The Internet Advertising Industry has greatly benefited from the interaction with the Economics and Computation community. Since the foundational papers in the field [Aggarwal et al. 2006; Varian 2007; Edelman et al. 2007] (which received the ACM SIGecom Test of Time Award in 2018), the market has matured and evolved in several directions. This has resulted in a new landscape arising from the maturity of the market and growing needs of the participants involved. Designing systems to address this new landscape requires new ideas and opens many new exciting research directions.

Modern ad markets have many new complexities that play an important role in determining outcomes in those markets (e.g. dealing with intermediaries rather than directly with advertisers, ad formats that go beyond the traditional single-item setting). Furthermore, there are many different spheres of optimization that contribute to the successful functioning of an ad market, such as budget pacing,

Authors' addresses: `kshipra,gurug,slahaie,perlroth,balusivan@google.com`

autobidding on behalf of advertisers, reserve price optimization for revenue maximization, and auction design that takes into account the repeated nature of ad auctions, to name a few. Through a series of articles we highlight some of the theoretical models, results, and open questions that have arisen from tackling these challenges at Google Research. This is the first article of this series, which we hope to follow up with a few more to highlight the diversity of topics. The selection of topics or the works surveyed are not comprehensive, but based on the authors' perspectives as researchers at Google.

In the first article in this series, we focus on two specific markets—the search ads and display ads markets—and discuss recent developments, some concrete models, recent results on these models, and the research challenges they present. In particular we focus on four different vignettes: Rich Ad Auctions, First Price Auctions, Auction with Intermediaries and Optimal Mechanisms for Market Intermediaries.

The first vignette in this series highlights the fact that the typical sponsored search auction is no longer a simple position auction. Ads that are sold can have different sizes and formats, advertisers can be strategic along multiple dimensions (in specifying their value for individual ads or set of ads that will allow), and there are new optimization challenges due to sheer volume of ad options and the packing constraints that limit the total space occupied by ads. Designing truthful, computationally tractable auctions in this setting is challenging. This model is very close to the more general Combinatorial Auctions model and thus also gives rise to other phenomena such as the impossibility of having allocations that are in the core. We highlight a couple of papers that address some of these challenges.

The second vignette is focused on the resurgence of interest in first-price auctions caused by the recent changes in the display ads industry where almost all ad exchanges have switched over to the first-price auction. This creates a variety of questions to study. From analyzing this complex game between competing exchanges involved in choosing their auction formats, to analyzing first-price auctions for their pratical aspects like price manipulability, and designing algorithms to optimize reserve price in first-price auctions, there are several challenges that the Economics and Computation community can contribute to. We discuss initial results on these topics and highlight several future directions.

The third vignette is focused on the role of intermediaries in the modern ad ecosystem. While conventional auction theory assumes that the advertisers directly interact with the seller through the auction, in practice advertisers are almost always represented by intermediaries. These intermediaries often have their own incentives and run downstream mechanisms to charge the advertisers, which can have negative welfare and revenue implications for the advertisers and the seller. We briefly survey the state of the art in mechanism design that tries to address this challenge.

The fourth vignette is built around the fact that the platform facilitating trade in two-sided markets like display ads exchange, ride-hailing, online retailing etc. often charges a fee for its services. It discusses the basic model for two-sided markets, and several results describing when the simple affine fee structure is optimal, how to design optimal revenue sharing schemes between the seller and the platform in repeated settings, and more generally how to design optimal mechanisms in repeated settings.

In the rest of the article, we will dive deeper into each vignette and include open questions which we hope will spur further research in these areas.

2. RICH ADVERTISING AUCTIONS

The first vignette is motivated by the area of Sponsored Search Auctions which has been the classic model to spur a lot of research in the field of algorithmic game theory; see [Lahaie et al. 2007] for a survey. Sponsored Search Auctions these days are quite different from the traditional setting with position auctions used to auction a constant number of text ads. These days advertisers can provide *Rich Ads* which are ads of variable size that include useful details for the user such as sitelinks, seller ratings and promotions. While the ads can be of varying sizes, the total ad space is limited carefully. The auctioneer's problem is thus to select the optimal set of ads that fit within the available space. This has been referred as the *Rich Advertising Auction* problem in the literature. The model of rich ads is interesting in two aspects. First, it's a special case of Combinatorial Auctions, which is a well-studied area of mechanism design with computational and economical challenges. The problem however is not as difficult as Combinatorial Auctions; we can assume reasonable structure on the buyer's valuations for the rich ads and hope to prove strong results in the special case. Second, the model captures Knapsack, thus solving for the optimal allocation is NP-hard. We will discuss a number of different attempts to handle challenges in specific models for Rich Ad Auctions.

Computationally Tractable Auctions. The rich ad auction is computationally challenging. It extends the knapsack problem which is NP-hard. While smaller instances could be solved via brute-force, serving constraints derived from traditional position auctions prohibit brute-force solutions. A recent paper by [Aggarwal et al. 2022] studies truthful auction design under computational constraints. They assume that each advertiser provides a value per click and a set of rich ads and can be strategic about either of these. Each ad has a size indicating the amount of vertical space it occupies and a probability of click. The probability of click can be arbitrary and different for different rich ads. The height and probability of click are assumed to be public. The goal of the auctioneer is to find the optimal set of ads that fit within the available space to maximize social welfare. The VCG auction is truthful in this setting but is computationally infeasible. This paper adopts a Myersonian approach—it focuses on a designing a monotone allocation rule that can be paired with Myerson's Lemma like payments to obtain a truthful auction. Note that this is a multi-parameter setting, so Myerson's Lemma does not directly apply, but an extension is proved. The main theorem of the paper is:

Informal Theorem 2.1. There is a truthful auction composed of a simple greedy monotone allocation rule and Myerson's lemma like payment that is computationally tractable and provides a 3-approximation to the optimal social welfare.

The rich-ad model studied in this paper is an instance of the *Multi-Choice Knapsack* problem. An incremental bang-per-buck heuristic can provide a 2-approximation for multi-choice knapsack. However that heuristic is not monotone: a buyer can report a subset of their rich ads to get an allocation with higher probability of click. The paper instead shows that the standard bang-per-buck heuristic provides a 3-approximation to the optimal and is monotone.

Another challenge faced in practical rich-ad auctions is deploying them in the real world where Generalized Second Price (GSP) payment rule is the norm. Switching to truthful payment rules can cause short-term revenue loss and the truthful payment rules are also harder to explain to the advertisers. Due to this GSP-like payment rules are preferred. The monotone-allocation rule from the paper can be paired with a GSP-like payment rule with the guarantee that the worst-case Price of Anarchy is not too bad. In particular,

Informal Theorem 2.2. The simple greedy allocation rule can be paired with GSP-like payments to obtain a mechanism with pure price of anarchy of at most 6 and coarse-correlated/Bayes-Nash price of anarchy of at most $6(1-1/e)$.

An earlier paper in this setting by [Ghiasi et al. 2019] studies a model where the probability of click (pCTR) for an ad depends on the integer height of the ad (captured as set of slots occupied). They assume that the pCTR is a submodular or subadditive function of the number of slots occupied, and consider the computational and auction design problem.

Informal Theorem 2.3. When pCTR as a function of number of slots occupied is a subadditive (respectively submodular) function then there is an LP rounding algorithm that obtains an $O(\log\log m/\log m)$ (respectively 1/4) approximation to the optimal social welfare allocation.

This is a purely algorithmic result and does not provide any pricing or incentive guarantees. For incentives, the authors consider a competitive-equilibrium approach. They show that using the allocation of the approximation algorithm above, one can compute prices that form a competitive equilibrium, where each agent obtains the set of slots that maximize their utility and all the slots are sold. This result is not fully general, and relies on the fact that pCTR as a function of the set of slots is the same for all buyers. It would be interesting to see if a more general result is feasible.

There are also works by [Deng et al. 2010], [Cavallo et al. 2017] that consider various algorithms and pricing models for specialized models of the rich ad auctions.

Core Allocation and Core-Competitive Auctions. Another consideration in designing rich ad auctions is that there is no auction that is a truthful and produces a core allocation [Ausubel and Milgrom 2002]. In Combinatorial Auctions there are no truthful auctions that produce allocations in the core and in fact VCG can have revenue close to zero. A standard solution is to find a core allocation that is close to truthful. [Hartline et al. 2018] adopt this approach for the rich ad auctions model. Assuming that the optimal allocation is computationally feasible, they provide an algorithm to compute a core allocation that is close to truthful. [Goel et al. 2015] instead focus on designing truthful auctions that are core-competitive; that is, a truthful auction that approximates the revenue of the least-revenue core allocation. The results are for a more stylized model termed "Text and Image". Suppose an advertiser provides a small text ad or a large image ad. The auctioneer wishes to show k text ads or a single image ad. What are the best set of ads to show and how should they be priced? In this simple model, finding the optimal allocation is computationally feasible, but unlike [Hartline et al. 2018] the focus is on finding truthful auctions. The authors prove:

Informal Theorem 2.4. For the Text-And-Image setting there is a core-competitive randomized auction that is $O(\ln \ln k)$-competitive and a core-competitive deterministic auction that is $O(\sqrt{\ln k})$-competitive.

Both of these bounds are shown to be tight.

Open Questions. There are a lot of remaining open questions in this domain. In the most general model an advertiser specifies a set of rich ads with possibly different value per click for each rich ad. How to design a computationally tractable truthful mechanism in this multi-dimensional setting is not at all understood. Even if a truthful mechanism is discovered, it would also be interesting to study if a core-competitive mechanism can be formulated. Note that VCG in this general context is truthful but not computationally tractable nor core-competitive.

In a more specialized setting, where advertisers specify only a value per click and set of rich ads, the general optimization problem faced in the Sponsored Search setting is not fully tackled. In the Sponsored Search setting, there is often a constraint on the total number of ads shown. There are also position effects in the probability of click—ads shown higher on the page will have increased probability of click. Designing truthful auctions with good approximation guarantees in this practical model remains open.

3. FIRST-PRICE AUCTIONS AND THE DISPLAY ADS MARKET

The display ads market has witnessed a sea change in recent years, which has spurred a resurgence in the study of first-price auctions. This vignette focuses on some of the challenges that these changes have motivated. Display ads are those that are shown to users visiting various webpages on the Internet, as opposed to sponsored search ads discussed in the preceding section that are shown in response to user-issued search engine queries. Display ads are sold either via contracts negotiated offline between the publisher and the advertiser, or via two-sided marketplaces called ad exchanges that bring the sellers (publishers) and buyers (advertisers or ad networks) together to a central marketplace. Our focus in this section will be on questions motivated by the recent changes in ad exchange auction formats. Articles by Muthukrishnan serve as excellent introductions to ad exchanges in general [Muthukrishnan 2009a; 2009b].

Traditionally ad exchanges have used second-price auctions as their choice of real-time auctions. However, over the past few years, there has been a rapid transition of exchanges to the first-price auction format, culminating in Google Ad Exchange making the switch in 2019 [Blog 2019]. Given this backdrop, in this note, we focus on three specific lines of exploration:

—Considering the competition between ad exchanges as a game, with the choice of auction format being each exchange's strategy, can we give a game-theoretic explanation of the recent move to the first-price auction? While there are other motivations for switching to first-price auctions like transparency (winner-pays-bid) and credibility [Akbarpour and Li 2020], we aim to investigate whether there are game-theoretic justifications as well.

—First-price auctions are clearly non-truthful. But how non-truthful are they? As the number of bidders increases, can the auction pressure automatically take care of truthfulness by inducing bidders to act as price-takers, or can bidders still have an incentive to shade their bids even in the limit as the number of bidders goes to infinity? How does this depend on the value distribution? How does the number of items traded play a role here?

—Reserve pricing in auctions is an important source of revenue improvements. However, unlike second price auctions, setting reserves in first-price auctions is considerably more challenging because first-price auctions are non-truthful, thereby making it harder to learn bidders' values. What are some good strategies to learn effective reserve prices in first-price auctions?

Competition between exchanges. Motivated by the recent display industry change where there has been a rapid transition of exchanges to the first-price auctions, [Paes Leme et al. 2020] ask whether the change to first-price auctions can be justified using game-theoretic principles. In their model there are n bidders where each bidder's value is drawn i.i.d. from a publicly known distribution F. Each exchange (or equivalently seller) j accounts for a λ_j fraction of the total number of auctions. All bidders participate in all auctions, and in particular therefore, bidders participate in auctions run by multiple different sellers, with seller j running mechanism $\mathcal{M}_j$. Bidders optimize their bids for the *aggregate auction* $\sum_j \lambda_j \mathcal{M}_j$. While optimizing their bid for each exchange is a possible strategy for the buyer, and perhaps large sophisticated buyers could devise exchange-specific strategies, smaller buyers often design a uniform strategy for the entire market. There are many practical reasons for this that the paper describes in some detail, and the cost of developing many different bidding strategies being non-trivial is an important reason. With such a bidding strategy, what is the equilibrium of the game between exchanges?

First, in a setting without reserve prices, the paper shows that the first-price auction is the unique Nash equilibrium in the game between exchanges.

Informal Theorem 3.1. If every exchange is only allowed to pick an auction where (a) the highest bidder wins, (b) winners pay no more than bid, (c) losers pay nothing and (d) there exists a pure-strategy Bayesian Nash Equilibrium among the bidders, then each exchange picking first-price auction is the unique Nash equilibrium.

This result gives a game-theoretic foundation for the switch to first-price auctions. The proof makes strong use of the revenue-equivalence principle, and indeed the proof reflects how the change in the industry itself unfolded. It says that as long there is at least one exchange not using a first-price auction, there is some exchange that would want to switch to a first-price auction and in the process gain more revenue. This cascading process ultimately results in all exchanges running a first-price auction.

In a setting where exchanges are allowed to pick auctions with reserve prices (which is the more realistic setting), the highest bidder does not always win. Therefore a straightforward application of revenue equivalence is not possible. However the paper shows that the desirability of first-price auctions is robust, by showing the following result based on a careful analysis of the segments of the bidding function.

Informal Theorem 3.2. If every exchange is only allowed to choose between a first-price auction with reserve and a second-price auction with reserve, then in a pure-strategy equilibrium, every exchange will pick a first-price auction with reserve.

Price-manipulability in first-price auctions. How non-truthful are first-price auctions? Given the many desirable properties they possess, like transparency (winners pay bid) and credibility [Akbarpour and Li 2020], and the fact that they are indeed being used in a variety of settings, it is useful to quantify the extent of non-truthfulness in first-price auctions. Adopting a metric used in the context of bitcoin fee designs, [Brustle et al. 2022] consider the following question. Suppose in an n agents k item first-price auction, where the values of bidders are drawn i.i.d. from some distribution F, we define the honest payment for agent i as $p_i^{\text{honest}} := v_i$ whenever $v_i \geq v_{(k)}$ and 0 otherwise, and $p_i^{\text{strategic}} := v_{(k+1)}$ whenever $v_i \geq v_{(k)}$ and 0 otherwise. In other words, $p_i^{\text{strategic}}$ captures the smallest payment an agent can manage to pay while still getting the good allocated. How much smaller can $p_i^{\text{strategic}}$ be when compared to p_i^{honest}? How does this depend on the number of bidders, number of items sold k, and the distribution F? In particular, the quantity of interest is $E[\delta_{\max}] = E\left[\max_i 1 - \frac{p_i^{\text{strategic}}}{p_i^{\text{honest}}}\right]$.

Interestingly, the authors exhibit surprising boundaries for this simple question. In particular, they show that the answer is not as simple as just letting $n \to \infty$ automatically makes $E[\delta_{\max}]$ vanish. First, even in the single item setting, where the pricing pressure is the highest, taking $n \to \infty$, even for relatively light-tailed distributions like α-strongly regular distributions[1] with α arbitrarily close to, but smaller than, 1 (that includes for example distributions like $F(x) = 1 - 1/x^d$ for arbitrarily large d), the quantity $\lim_{n\to\infty} E[\delta_{\max}] > 0$. But a sharp boundary is exhibited between MHR (which is 1-strongly regular) and α-strongly regular distributions for any $\alpha < 1$ by showing that for MHR distributions $\lim_{n\to\infty} E[\delta_{\max}] = 0$. This holds true for even k item auctions with k that is $\Theta(1)$. But as k grows larger as a function of n, as one might expect, the effect of pricing pressure goes down and we expect more and more distributions where $E[\delta_{\max}] > 0$. Indeed the authors show that for $k = \sqrt{n}$, even for MHR distributions $E[\delta_{\max}] > 0$. Precisely how small can k be for this to happen is an open question.

Reserve pricing in first-price auctions. The move to first-price auctions has led to a re-evaluation of many auction parameters, notably *reserve prices* (also known as floor prices). Under the original second-price auctions, the fact that these auc-

[1]The virtual valuation function $\phi(v) = v - \frac{1-F(v)}{f(v)}$ plays an important role in revenue maximization and many distribution classes are defined using this function. Regular distributions are those for which $\frac{d\phi}{dv} \geq 0$ and Monotone Hazard Rate (MHR) distributions are those for which $\frac{d\phi}{dv} \geq 1$. The class of α-strongly regular distributions interpolates between these two extremes by asking for $\frac{d\phi}{dv} \geq \alpha$ for some $\alpha \in [0, 1)$. When the virtual valuation function $\phi(\cdot)$ is not differentiable, the definition for α-strongly regular distributions is that for any $y \geq x \geq 0$, we have $\phi(y) - \phi(x) \geq \alpha(y - x)$ (see [Cole and Roughgarden 2014]).

tions are truthful means that bid distributions can be taken to represent true value distributions, and these distributions in theory do not change under reserve prices (modulo some bid censoring if values do not meet the reserves). This greatly simplifies the counterfactual evaluation of different reserve prices. Therefore, research on reserve pricing in second-price ad auctions has historically focused on other aspects, such as the relative merits of different personalized pricing policies [Paes Leme et al. 2016].

In contrast to second-price auctions, bidders in first-price auctions have an incentive to shade their bids, but bid shading strategies can vary among bidders in practice, making it difficult to infer underlying values. [Feng et al. 2021] propose a gradient-based approach where the reserve price is regularly updated based on an estimate of the current revenue objective (revenue as a function of floor price) and its gradient. This obviates the need to infer values; instead, the challenge is to apply variance reduction techniques to obtain an accurate estimate of the revenue gradient. The authors decompose the revenue objective into a demand component (amenable to parametric modeling), and a bidder response component, whose variance can be reduced by filtering certain bids. Intuitively, bids far from the reserve are little affected by changes in the reserve, under natural bidding models, so they can be omitted when computing gradients to reduce noise. The authors demonstrate how these techniques achieve improved convergence rates via simulations over display ad auction data drawn from Google Ad Exchange.

Open Questions. Each of these lines of exploration offers interesting open questions. The authors in [Paes Leme et al. 2020] introduce and analyze a new market and game, where the players are exchanges, and they compete with each other using auctions as their strategies. What other questions can be studied in this different kind of market? Concretely, the authors consider one specific buyer strategy, namely one that responds to the average auction, and describe the consequences. It would be quite interesting to study other buyer strategies when responding to multiple exchanges running different auctions, and analyze the outcomes of that game. In the price manipulability paper, [Brustle et al. 2022] show that even with an infinite number of competing bidders, price manipulability does not vanish in first-price auctions. Can this same metric be studied for other non-truthful auctions like GSP? Can we formally establish with this metric the intuitive notion that GSP is less manipulable than first-price auctions? In the paper on learning reserve prices for first-price auctions, [Feng et al. 2021] show that a faithful parametric model of demand curves (i.e., the mass of bidders whose values exceed a given threshold) leads to improved variance reduction and convergence. What parametric models are most relevant in practice?

4. INTERMEDIATION IN AUCTIONS

This vignette is motivated by a common feature in many ad auction markets: There are many cases in which advertisers do not participate directly in the auction but, instead, are represented by intermediaries. For example, it is quite common to have advertisers work with ad agencies to create their marketing campaigns and also to facilitate negotiation and buying[2]. In the AdExchange marketplace that sells advertisements on webpages, advertisers are represented by Ad Networks, who do real-time bidding on their behalf into the AdExchange auction. These Ad Networks create value for advertisers by pooling many advertisers together and thereby obtaining more favorable prices in the AdExchange auction [Muthukrishnan 2009a].

Intermediation on behalf of buyers is often neglected when designing the auction mechanisms and can lead to poor performance of such auctions. In particular, intermediaries may enable buyers to engage in collusive behavior, decreasing the auction pressure. Also, intermediaries may charge participation fees (or enact more sophisticated mechanisms) to price the buyers they represent, leading to double-marginalization effects on the auction. Both of these effects negatively impact the efficiency and revenue of the auction. We next present some relevant works that tackle this problem and present some new solutions.

One of the first works to address the question of intermediaries in Online Ad Networks was the work of [Feldman et al. 2010]. In their model, the advertisers choose (or are required) to participate in the auction via intermediaries. They consider the case where each intermediary is a revenue maximizing agent and all of them will first run (independently and in parallel) a contingent auction. This auction determines a contingent buyer and a contingent price for that intermediary (i.e., which buyer among that intermediary's captive buyers will get the good if the intermediary wins the good), and the price she will pay the intermediary. Then, the seller runs an auction for the good among intermediaries. The intermediaries bid in the main seller's auction based on the outcome in their own contingent auctions. The winning intermediary then transfers the good as determined by its contingent auction. In their work, [Feldman et al. 2010] show a symmetric equilibrium where the intermediaries choose a randomized reserve. Based on this mechanism, they also show that the optimal auction for the seller (which takes into account the intermediaries' reserves) and consider the optimal reserve that the seller can set.

Their analysis even for the case where each intermediary has only one captive buyer is very clever and intricate, and extending it is highly nontrivial. For example, [Balseiro and Candogan 2017] show that even for the case when there is only one possible intermediary, the optimal downstream mechanism between advertisers and the intermediary is the solution to a nontrivial convex program. For the case of multiple intermediaries and advertisers, characterizing the optimal contracts remains an open question. Even though characterizing the optimal downstream mechanism is hard, in practice, sellers are usually unaware about which buyers each intermediary represents and what are the intermediaries' incentives in the auction (which depend on the downstream negotiation with buyers).

Given these challenges, [Aggarwal et al. 2022] begin the quest for mechanisms that are robust to the demand structure between intermediaries and the buyers. Specifically, the authors design an intermediary-proof mechanism with two key properties:

—It is independent of the downstream market structure (which buyers are represented by which intermediaries).

[2]See [Choi et al. 2020] for a comprehensive review of intermediation in ad-auctions.

—The revenue obtained is a constant factor of the optimal welfare (the highest revenue the auctioneer could possibly obtain).

[Aggarwal et al. 2022] considers a model where the seller designs an auction to sell multiple items to unit-demand buyers. Each buyer has private valuations for the items and either participates directly in the auction or is represented by an intermediary. The model is agnostic about the downstream decisions between buyers and intermediaries and only sets some mild restrictions on the behavior of the intermediaries in the main auction.

The first result of the paper addresses the case where items are homogeneous and buyers' valuations are drawn independently according to a distribution F which is assumed to be α-strongly regular[3] with $\alpha > 0$. For this setting the authors show that:

Informal Theorem 4.1. When items are homogeneous, there exists a posted-price per item mechanism that achieves a factor $1/2 \cdot \alpha^{2/(1-\alpha)} \cdot (1 - 1/e)$ of the welfare regardless of the downstream interaction between buyers and intermediaries.

The second result of the paper tries to incorporate important features from sponsored search advertising and assumes that items are no longer identical but instead are separable. That is, the buyer's valuation for an item is now weighted by its *position normalizer*. In this setting, the authors show that a simple price per item is not longer optimal but it rather requires a more sophisticated mechanism that induces a high value buyer to get a high value item.

Informal Theorem 4.2. For the separable item case, there exists a sequential posted-price mechanism that achieves a factor $(1 - e^{-1/2\alpha^{1/(1-\alpha)}}) \cdot (1 - 1/e)$ of the welfare regardless of the downstream interaction between buyers and intermediaries.[4]

Open Questions. While the previous paper is the first to study the robustness of mechanisms in the presence of intermediation, important questions remain unanswered. First, what are simple practical mechanisms that have good revenue guarantees when buyers' valuations are not α-strongly regular? Second, in the heterogeneous item case, if we consider valuations are no longer separable but rather independent across items, is it true that a pricing mechanism can achieve good revenue guarantee regardless of the downstream market?

Aside from these two natural questions left open from the prior work, a more relevant practical question is to understand the setting where buyers are not unit-demand but may have more complex valuations for goods. Does there exist a mechanisms robust to downstream market conditions while also achieving a constant factor of the optimal revenue?

[3]See Section 3 for a formal definition of α-strongly regularity. Intuitively, this class of distributions smoothly interpolates between monotone hazard rate distributions ($\alpha = 1$) and regular distributions ($\alpha = 0$).

[4]For this result, the authors impose a more restrictive set of assumptions regarding the space of contracts between buyers and intermediaries (see Section 5 of [Aggarwal et al. 2022] for more details.)

5. OPTIMAL MECHANISMS FOR MARKET INTERMEDIARIES

This vignette focuses on the fact that in many online marketplaces, a platform acts as an intermediary to facilitate the sale between sellers and buyers. Such marketplaces include the Display Ads market's Ad Exchange marketplace mentioned in the earlier sections, online retailing, ride-hailing marketplaces, labor platforms, short-term lodging marketplaces etc. An important feature of modeling such marketplaces is that apart from the sellers and the buyers, the platform itself is another utility-maximizing participant. In this section we focus on mechanisms that align the incentives of the seller and the platform, and (approximately) optimize their objectives.

The bilateral trade model where there is two-sided private information, namely, a seller with a private cost c and a buyer with private value v is a good starting place for thinking about two-sided markets. The pioneering work of [Myerson and Satterthwaite 1983] shows that when the private values and costs are drawn from independent publicly known distributions, there is no weakly budget-balanced mechanism (i.e., ones where the intermediary, namely the platform/mechanism designer, cannot receive money from elsewhere apart from the buyer), that is Bayesian Incentive Compatible (BIC) and Interim Individually Rational (IIR), that can maximize efficiency; i.e. no such mechanism can ensure that trade happens whenever there is gains from trade, namely when buyer's private value v exceeds seller's cost c. In other words, efficient trade is impossible in the presence of two-sided private information. [Myerson and Satterthwaite 1983] also give the BIC, IIR and weakly budget-balanced mechanism that achieves the optimal gains from trade, namely, the second-best. Since then there have been a number of follow-up works that produce simple mechanisms that approximate the second-best (e.g. [Brustle et al. 2017]), simple mechanisms that approximate the first-best under some conditions (e.g. [McAfee 2008]) and more recent works that show how simple mechanisms can constant-factor approximate the first-best unconditionally ([Deng et al. 2022], [Fei 2022]).

The preceding discussion misses an important feature of the market intermediation problem: the platform has its own utility as well, namely the profit it earns. In practice, often a revenue-sharing agreement exists giving the platform a constant fraction share of the revenue in each transaction. Once a constant revenue share is established, note that the incentives of the seller and the platform are aligned. However, is a constant revenue share optimal in any formal sense – for example does it optimize platform's profits? [Loertscher and Niedermayer 2007], [Loertscher and Niedermayer 2013] initiated the study of fee-setting mechanisms: namely the ones where (a) the platform sets a fee structure as a function of the transaction price, (b) the seller observes the fee structure and sets a price, and (c) the buyers takes or leaves the offer made by the seller. They show that fee-setting mechanisms are without loss of generality, in that any allocation rule that can be implemented in a Bayesian Nash Equilibrium (BNE) can be implemented by a fee-setting mechanism as well with an appropriately chosen fee structure. They provide sufficient conditions for the platform's optimal mechanism to be implemented by an *affine fee-setting mechanism*: ones that set a fixed percentage of each transaction as their fees, plus a fixed constant: $\alpha \cdot p + k$. They show that in markets where the threshold to enter the market is very high for the buyers (so that only those with the highest value enter) and very low for the sellers (so that only sellers with the lowest cost enter), the conditional distribution upon entry converges to generalized Pareto distributions, and the latter is shown as a sufficient condition for affine fee-setting mechanisms to be optimal. [Niazadeh et al. 2014] continue this line of work and study when an affine fee structure is *approximately optimal*. They provide sufficient conditions on the buyer distribution so that affine fee-setting mechanisms get a constant fraction of the optimal platform profit for all seller distributions.

[Gomes and Mirrokni 2014] note that platforms operate in a very competitive market, and study the mechanism to be employed by the platform when operating in equilibrium with multiple other competing platforms. They show that for a large class of competition games between platforms, in equilibrium, the platform maximizes a convex combination of seller's profit and the platform's profit. The weight α assigned to the seller's profit is a measure of the degree of competitiveness. Generalizing the analysis of [Myerson and Satterthwaite 1983] they show how to derive the optimal mechanism that maximizes such a convex combination of profits. Further, they show that the popular practice of charging a constant percentage of the transaction revenue as fees is optimal if and only if the seller's cost has a cdf of the form $F(c) = (c/H)^k$ where H is the largest point in the support of the distribution.

[Balseiro et al. 2017] study revenue sharing schemes for the platform when the interaction between seller and buyer (intermediated by the platform) happens over a large number of rounds. For instance, the number of auctions intermediated by an ad exchange over the course of a day in the display ads market is very large. Exploiting the repeated nature of the interaction, they study dynamic revenue sharing schemes, namely, ones where the revenue sharing percentage has to be met in aggregate over all interactions, and not over individual interactions. In aggregate means that the payout to the seller over all interactions weakly exceeds the sum of costs in all those interactions, and the platform profit is at most α fraction of the sum of the transaction prices over all the interactions. They study the case where the platform has to run a second-price auction among all the buyers. Even with this restriction, and in a single-period setting, they show that the simple scheme of the platform running an auction with a reserve price of $c/(1-\alpha)$ and keeping an α fraction of the revenue upon sale is not always optimal. They derive the optimal reserve price for the platform. In the multi-period setting, they use Lagrangian relaxation to derive the optimal policy that pays the seller a convex combination of their cost c and $1-\alpha$ of the transaction revenue. While the obtained policy satisfies the revenue sharing and cost constraints only in expectation, they show how it can be modified so that the modified policy approaches the Lagrangian policy in the limit.

While [Balseiro et al. 2017] fixed the auction and studied how to optimize the revenue sharing schemes in repeated settings, [Balseiro et al. 2019] study the design of the optimal auction for the intermediary. They show how to design mechanisms that are asymptotically efficient and weakly budget-balanced with high probability.

REFERENCES

AGGARWAL, G., BHAWALKAR, K., GURUGANESH, G., AND PERLROTH, A. 2022. Maximizing revenue in the presence of intermediaries. In *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, M. Braverman, Ed. LIPIcs, vol. 215. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 1:1–1:22.

AGGARWAL, G., BHAWALKAR, K., MEHTA, A., MOHAN, D., AND PSOMAS, A. 2022. Simple mechanisms for welfare maximization in rich advertising auctions. *CoRR abs/2206.02948*.

AGGARWAL, G., GOEL, A., AND MOTWANI, R. 2006. Truthful auctions for pricing search keywords. In *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006), Ann Arbor, Michigan, USA, June 11-15, 2006*, J. Feigenbaum, J. C. Chuang, and D. M. Pennock, Eds. ACM, 1–7.

AKBARPOUR, M. AND LI, S. 2020. Credible auctions: A trilemma. *Econometrica 88,* 2, 425–467.

AUSUBEL, L. AND MILGROM, P. 2002. Ascending auctions with package bidding. *The B.E. Journal of Theoretical Economics 1,* 1, 1–44.

BALSEIRO, S. R. AND CANDOGAN, O. 2017. Optimal contracts for intermediaries in online advertising. *Operations Research 65,* 4, 878–896.

BALSEIRO, S. R., LIN, M., MIRROKNI, V. S., LEME, R. P., AND ZUO, S. 2017. Dynamic revenue sharing. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, Eds. 2681–2689.

BALSEIRO, S. R., MIRROKNI, V. S., LEME, R. P., AND ZUO, S. 2019. Dynamic double auctions: Towards first best. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, T. M. Chan, Ed. SIAM, 157–172.

BLOG, G. A. M. 2019. *An update on first price auctions for Google Ad Manager.*

BRUSTLE, J., CAI, Y., WU, F., AND ZHAO, M. 2017. Approximating gains from trade in two-sided markets via simple mechanisms. In *EC*. 589–590.

BRUSTLE, J., DÜTTING, P., AND SIVAN, B. 2022. Price manipulability in first-price auctions. In *Proceedings of the ACM Web Conference 2022*. 58–67.

CAVALLO, R., KRISHNAMURTHY, P., SVIRIDENKO, M., AND WILKENS, C. A. 2017. Sponsored search auctions with rich ads. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*, R. Barrett, R. Cummings, E. Agichtein, and E. Gabrilovich, Eds. ACM, 43–51.

CHOI, H., MELA, C. F., BALSEIRO, S. R., AND LEARY, A. 2020. Online Display Advertising Markets: A Literature Review and Future Directions. *Information Systems Research 31,* 2 (June), 556–575.

COLE, R. AND ROUGHGARDEN, T. 2014. The sample complexity of revenue maximization. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, D. B. Shmoys, Ed. ACM, 243–252.

DENG, X., SUN, Y., YIN, M., AND ZHOU, Y. 2010. Mechanism design for multi-slot ads auction in sponsored search markets. In *Frontiers in Algorithmics*, D.-T. Lee, D. Z. Chen, and S. Ying, Eds. Springer Berlin Heidelberg, Berlin, Heidelberg, 11–22.

DENG, Y., MAO, J., SIVAN, B., AND WANG, K. 2022. Approximately efficient bilateral trade. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, S. Leonardi and A. Gupta, Eds. ACM, 718–721.

EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review 97,* 1 (March), 242–259.

FEI, Y. 2022. Improved approximation to first-best gains-from-trade. In *To appear in WINE '22: The 18th Conference on Web and Internet Economics.*

FELDMAN, J., MIRROKNI, V., MUTHUKRISHNAN, S., AND PAI, M. M. 2010. Auctions with intermediaries. In *Proceedings of the 11th ACM conference on Electronic commerce.* 23–32.

Feng, Z., Lahaie, S., Schneider, J., and Ye, J. 2021. Reserve price optimization for first price auctions in display advertising. In *International Conference on Machine Learning*. PMLR, 3230–3239.

Ghiasi, M., Hajiaghayi, M., Lahaie, S., and Yami, H. 2019. On the efficiency and equilibria of rich ads. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, S. Kraus, Ed. ijcai.org, 301–307.

Goel, G., Khani, M. R., and Leme, R. P. 2015. Core-competitive auctions. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, T. Roughgarden, M. Feldman, and M. Schwarz, Eds. ACM, 149–166.

Gomes, R. and Mirrokni, V. S. 2014. Optimal revenue-sharing double auctions with applications to ad exchanges. In *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*, C. Chung, A. Z. Broder, K. Shim, and T. Suel, Eds. ACM, 19–28.

Hartline, J. D., Immorlica, N., Khani, M. R., Lucier, B., and Niazadeh, R. 2018. Fast core pricing for rich advertising auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, É. Tardos, E. Elkind, and R. Vohra, Eds. ACM, 111–112.

Lahaie, S., Pennock, D. M., Saberi, A., and Vohra, R. V. 2007. Sponsored search auctions. *Algorithmic game theory 1*, 699–716.

Loertscher, S. and Niedermayer, A. 2007. When is seller price setting with linear fees optimal for intermediaries? `https://www.econstor.eu/bitstream/10419/76741/1/617153094.pdf`.

Loertscher, S. and Niedermayer, A. 2013. Fee-setting mechanisms: On optimal pricing by intermediaries and indirect taxation. `http://andras.niedermayer.ch/uploads/File/LN2013-10-10.pdf`.

McAfee, R. P. 2008. The gains from trade under fixed price mechanisms. *Applied economics research bulletin 1,* 1, 1–10.

Muthukrishnan, S. 2009a. Ad exchanges: Research issues. In *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*, S. Leonardi, Ed. Lecture Notes in Computer Science, vol. 5929. Springer, 1–12.

Muthukrishnan, S. 2009b. Adx: a model for ad exchanges. *SIGecom Exch. 8,* 2, 9.

Myerson, R. B. and Satterthwaite, M. A. 1983. Efficient mechanisms for bilateral trading. *Journal of economic theory 29,* 2, 265–281.

Niazadeh, R., Yuan, Y., and Kleinberg, R. D. 2014. Simple and near-optimal mechanisms for market intermediation. In *Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*, T. Liu, Q. Qi, and Y. Ye, Eds. Lecture Notes in Computer Science, vol. 8877. Springer, 386–399.

Paes Leme, R., Pal, M., and Vassilvitskii, S. 2016. A field guide to personalized reserve prices. In *Proceedings of the 25th international conference on world wide web*. 1093–1102.

Paes Leme, R., Sivan, B., and Teng, Y. 2020. Why do competitive markets converge to first-price auctions? In *Proceedings of The Web Conference 2020*. 596–605.

Varian, H. 2007. Position auctions. *International Journal of Industrial Organization 25,* 6, 1163–1178.