Challenges in Online Ad Marketplaces

Okke Schrijvers

December 30, 2021

1 Introduction

Modern online ad marketplaces are complex ecosystems with many novel challenges: 1) ads are optimized for heterogeneous conversion events and may appear interspersed with organic content, necessitating allocation and pricing algorithms beyond the simple greedy approach in sponsored search, 2) auction code bases have grown too large to manually verify, necessitating data-driven approaches to guarantee incentive compatibility, and 3) budget management solutions are offered to advertisers that aim to optimally bid on the advertiser's behalf thousands or millions of auctions. These problems don't have off-the-shelve solutions, and require novel research to drive product impact. The Core Data Science (CDS) team at Meta partners with Ads teams across the company to do this fundamental research and help our partner teams launch innovative products.

In this note, I survey the research projects that I've personally been involved in while working at CDS. To prevent confusion, I have only cited the papers that I co-authored while working at CDS with co-workers, interns, and academic collaborators. The related work sections of the cited papers place each of the contributions in the larger academic context.

2 Auctions in a Feed

Content feeds—such as Facebook's News Feed, reddit, and Apple News—have two properties that aren't commonly captured in existing literature on online advertising: firstly they have multi-modal content, e.g. images, video, and text, which naturally leads to multi-modal ads with different associated conversion events, and secondly they intersperse organic content with ads content which makes the allocation (and thus pricing) challenging. We've studied both algorithms and equilibrium properties in this setting.

Auctions with Ad Types. Our first paper [1] proposes a formal model to reason about such auctions. An *Ad Types Problem* instance has k ad types[1], that each have their own discount curve over n slots, i.e. all ads of type θ have discount curve $\alpha_1^{(\theta)} \geq \alpha_2^{(\theta)} \geq \ldots \geq \alpha_n^{(\theta)} \geq 0$ that represents the slot-specific action-rates. All ad types agree on the order of the

[1]In economics literature, "type" sometimes refers to private information. That's not the case here: type represents the content type, e.g. video or link-click, and is publicly known.

slots. Gap rules are modeled by a $k \times k$ matrix G, which indicates for each pair of ad types (θ_i, θ_j), that after showing an ad of type θ_i, the next G_{ij} stories cannot be of type θ_j.

For the Ad Types Problem without gap rules, i.e. when all entries $G_{ij} = C$ are some constant: 1) we give an algorithm that finds the maximum weight matching that runs in $O(n^2(k+\log n))$ time for n slots and n ads of each type—cf. $O(kn^3)$ when using the Hungarian algorithm—, and 2) we show how to apply reserve prices in total time $O(n^3(k + \log n))$.

For the Ad Types Problem with gap rules, i.e. G is unconstrained, we show that the problem is hard to approximate within $k^{1-\epsilon}$ for any $\epsilon > 0$ (even without discount curves) by reduction from Maximum Independent Set. On the positive side, we show a Dynamic Program formulation that solves the problem (including discount curves) optimally and runs in $O(k \cdot n^{2k+1})$ time.

Equilibria. Our second paper [2], which was driven by intern Hadi Elzayn, studies equilibrium quality of auctions in the *Ad Types* setting with *greedy* or *optimal* allocation combined with *generalized second-price* (GSP) or *Vickrey-Clarke-Groves* (VCG) pricing. The paper makes three contributions: first, we give upper and lower bounds on the Price of Anarchy (PoA) for auctions which use greedy allocation with GSP pricing, greedy allocations with VCG pricing, and optimal allocation with GSP pricing. Second, we give Bayes-Nash equilibrium characterizations for two-player, two-slot instances (for all auction formats) and show that there exists both a revenue hierarchy and revenue equivalence across some formats. Finally, we used no-regret learning algorithms and bidding data from a large online advertising platform and no-regret learning algorithms to evaluate the performance of the mechanisms under semi-realistic conditions. For welfare, we find that the optimal-to-realized welfare ratio (an empirical PoA analogue) is broadly better than our upper bounds on PoA; For revenue, we find that the hierarchy in practice may sometimes agree with simple theory, but generally appears sensitive to the underlying distribution of bidder valuations.

3 Measuring Incentive Compatibility

For a theoretician, incentive compatibility is a property that you prove, not something that you measure, so what's going on? With auctions growing in complexity, mathematical proofs are intractable, so it's useful to be able to verify the extent to which the auction is incentive compatible independently. This is also an important problem for advertisers who cannot directly observe the auction that they participate in. We've published two methods for measuring incentive compatibility under two different informational models.

Measure IC-Regret using Online Learning. In our first paper [3], driven by intern Zhe Feng, we investigate the problem of measuring end-to-end Incentive Compatibility (IC) regret given black-box access to an auction mechanism. Our goal is to 1) compute an estimate for IC regret in an auction, 2) provide a measure of certainty around the estimate of IC regret, and 3) minimize the time it takes to arrive at an accurate estimate. We consider two main problems, with different informational assumptions: In the *advertiser problem* the goal is to measure IC regret for some known valuation v, while in the more general *demand-side platform (DSP) problem* we wish to determine the worst-case IC regret over all possible valuations. The problems are naturally phrased in an online learning model and we design algorithms for both problems. We give an online learning algorithm where for the advertiser problem the error of determining IC shrinks as $O\left(\frac{|B|}{T} \cdot \left(\frac{\ln T}{n} + \sqrt{\frac{\ln T}{n}}\right)\right)$ (where B is the finite set of bids, T is the number of time steps, and n is number of auctions per time step), and for the DSP problem it shrinks as $O\left(\frac{|B|}{T} \cdot \left(\frac{|B| \ln T}{n} + \sqrt{\frac{|B| \ln T}{n}}\right)\right)$. For the DSP problem, we also consider stronger IC regret estimation and extend our algorithm to achieve better IC regret error. We validate the theoretical results using simulations with Generalized Second Price (GSP) auctions, which are known to not be incentive compatible and thus have strictly positive IC regret.

Envy as a Proxy for IC-Regret. One of the short-comings of the previous work is that it is computationally intensive, since it requires the execution of many counterfactual experiments. In follow-up work [4] we show that similar results can be obtained using the notion of IC-Envy. The advantage of IC-Envy is its efficiency: it can be computed using only the auction's outcome. We show that this is true both for traditional position auctions and Ad Types auctions. For position auctions, we show that for a large class of pricing schemes (which includes e.g. VCG and GSP), IC-Envy $\geq$ IC-Regret (and IC-Envy = IC-Regret under mild supplementary conditions). For Ad Types auctions, we show that for a generalization of the GSP mechanism IC-Envy $\geq$ IC-Regret holds as well. Our theoretical results are completed showing that, in the position auction environment, IC-Envy can be used to bound the loss in social welfare due to the advertiser untruthful behavior.

Finally, we show experimentally that IC-Envy can be used as a feature to predict IC-Regret in settings not covered by the theoretical results. In particular, using IC-Envy yields better results than training models using only price and value features.

4 Budget Management

While advertisers can participate in individual auction directly, large online platforms generally also offer budget-management tools and auto-bidding solutions. These leads to generally different bidding dynamics, so it's important to understand how advertisers can learn to best-respond and what properties the market has in equilibrium.

Pacing Equilibrium. In the isolated auction of a single item, second price often dominates first price in properties of theoretical interest. But, single items are rarely sold in true isolation, so considering the broader context is critical when adopting a pricing strategy. In the paper with Conitzer et al. [5], we study a model centrally relevant to Internet advertising and show that when items (ad impressions) are individually auctioned within the context of a larger system that is managing budgets, theory offers surprising endorsement for using a first price auction to sell each individual item. In particular, first price auctions offer theoretical guarantees of equilibrium uniqueness, monotonicity, and other desirable properties, as well as efficient computability as the solution to the well-studied Eisenberg-Gale convex program. We also use simulations to demonstrate that a bidder's incentive to deviate vanishes in thick markets.

Online Learning with Bandit Feedback. In follow-up work [6] we study the problem of an online advertising system that wants to optimally spend an advertiser's given budget for a campaign across multiple platforms, without knowing the value for showing an ad to the users on those platforms. We model this challenging practical application as a Stochastic Bandits with Knapsacks problem over T rounds of bidding with the set of arms given by the set of distinct bidding m-tuples, where m is the number of platforms. We modify the algorithm proposed in Badanidiyuru *et al.,* to extend it to the case of multiple platforms to obtain an algorithm for both the discrete and continuous bid-spaces. Namely, for discrete bid spaces we give an algorithm with regret $O\left(OPT\sqrt{\frac{mn}{B}} + \sqrt{mnOPT}\right)$, where OPT is the performance of the optimal algorithm that knows the distributions. For continuous bid spaces the regret of our algorithm is $\tilde{O}\left(m^{1/3} \cdot \min\left\{B^{2/3}, (mT)^{2/3}\right\}\right)$. When restricted to this special-case, this bound improves over Sankararaman and Slivkins in the regime $OPT \ll T$, as is the case in the particular application at hand. Second, we show an $\Omega\left(\sqrt{mOPT}\right)$ lower bound for the discrete case and an $\Omega\left(m^{1/3}B^{2/3}\right)$ lower bound for the continuous setting, almost matching the upper bounds. Finally, we use a real-world data set from a large internet online advertising company with multiple ad platforms and show that our algorithms outperform common benchmarks and satisfy the required properties warranted in the real-world application.

Acknowledgements. Big thanks to all my awesome co-authors. Firstly, the interns I've worked with: Zhe Feng and Hadi Elzayn who drove the main technical results for [3] and [2] respectively. Secondly, my co-workers at Meta: Riccardo Colini-Baldeschi, Eric Sodomka, Chris Wilkens, Christian Kroer, Julián Mestre, Nico Stier, Vashist Avadhanula, Karthik Abinav Sankararaman, and Brian Lan. Lastly our amazing academic collaborators Stefano Leonardi, Vince Conitzer, and Debmalya Panigrahi.

References

[1] Riccardo Colini-Baldeschi, Julián Mestre, Okke Schrijvers, and Christopher A Wilkens. The ad types problem. In *International Conference on Web and Internet Economics*, pages 45–58. Springer, 2020.

[2] Hadi Elzayn, Riccardo Colini-Baldeschi, Brian Lan, and Okke Schrijvers. Equilibria in auctions with ad types, 2021.

[3] Zhe Feng, Okke Schrijvers, and Eric Sodomka. Online learning for measuring incentive compatibility in ad auctions. In *The World Wide Web Conference*, pages 2729–2735, 2019.

[4] Riccardo Colini-Baldeschi, Stefano Leonardi, Okke Schrijvers, and Eric Sodomka. Envy, regret, and social welfare loss. In *Proceedings of The Web Conference 2020*, pages 2913–2919, 2020.

[5] Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E Stier-Moses, and Chris Wilkens. Pacing equilibrium in first-price auction

markets. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 587–587, 2019.

[6] Vashist Avadhanula, Riccardo Colini-Baldeschi, Stefano Leonardi, Karthik Abinav Sankararaman, and Okke Schrijvers. Stochastic bandits for multi-platform budget optimization in online advertising. In *Proceedings of The Web Conference 2021*, 2021.