



HAL
open science

How to bid in unified second-price auctions when requests are duplicated

Benjamin Heymann

► **To cite this version:**

Benjamin Heymann. How to bid in unified second-price auctions when requests are duplicated. *Operations Research Letters*, 2020, 48 (4), pp.446-451. 10.1016/j.orl.2020.05.010 . hal-02867979

HAL Id: hal-02867979

<https://hal.science/hal-02867979>

Submitted on 15 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

How to bid in unified second-price auctions when requests are duplicated

Benjamin Heymann, Criteo

32 rue Blanche, 75009 Paris, France. b.heyman@criteo.com

Abstract

In display advertising auctions, a unique display opportunity may trigger many bid requests being sent to the same buyer. Bid request duplication is an issue: programmatic bidding agents might bid against themselves. In a simplified setting of unified second-price auctions, the optimal solution for the bidder is to randomize the bid, which is quite unusual. Our results motivate the recent switch to a unified first-price auction by showing that a unified second-price auction could have been detrimental to all participants.

Keywords: advertising, mechanism design, bidding, second-price auctions

1. Introduction

A large portion of the internet is financed by ad placements on publishers' websites. Those ad placements are sold either through guaranteed contracts or through auction markets called Real-Time Bidding (RTB) [1, 2]. While guaranteed contracts decide in advance the number of displays and the sale price, RTB markets take place in real-time - as the name suggests- via programmatic buying while the page is loading in the user's browser.

The Waterfall. Until recently [3], publishers mostly relied on the waterfall logic [4], which we describe hereafter. First, the publisher (e.g., The NY Times) sets in advance a floor price for each ad exchange. Then, when the user loads the page, the publisher's ad server calls the ad exchanges sequentially. The publisher might rely on an intermediate piece of technology called the Supply Side Platform (SSP, e.g., DoubleClick for Publishers, Rubicon Project for Sellers and MoPub) for this purpose. The ad exchanges (e.g., DoubleClick Ad Exchange and AppNexus) host internal auctions. The Demand Side platforms (DSP, e.g., MediaMath and Criteo) receive the bid requests from the ad exchanges they

<https://doi.org/10.1016/j.orl.2020.05.010>

© 2015, Elsevier. Licensed under the Creative Commons Attribution-NonCommercialNoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>

are connected to, and bid in the name of their clients: the advertisers. Once the ad exchanges have received the bids from the DSPs, they send the clearing prices back to the publisher. The waterfall logic follows a fixed priority order, and allocates the impression to the first ad exchange that proposes a bid above its floor.

Header Bidding. The waterfall logic has several drawbacks. In particular, it increases the latency of the user experience, and the allocation is not efficient (a low bid in a high ranked exchange might beat a high bid in a low ranked exchange). To tackle those drawbacks, the waterfall logic has been progressively replaced by Header Bidding, in which the ad exchanges themselves participate in a first-price auction. That is, the ad exchange with the highest bid wins the auction and pays its bid.

When Header Bidding was introduced, the most common setup was a two-step auction mixing first and second pricing rules [5, 6]: in the first step, several ad exchanges hosted a second-price auction – the highest bidder wins the auction and pays the second highest bid – and during the second step, they used their clearing price as a bid in a final first-price auction. This mechanism is not efficient [7]: the highest bidder may still not get the item. This motivated a switch to a more efficient and less obscure mechanism. The natural candidates would have been a second-price auction both in the ad exchanges and in the SSP, or a first-price auction both in the ad exchanges and in the SSP. In this paper, we provide an argument against the use of a unified second-price auction, which is describe more precisely in §3. This is an important argument, but we do not claim that this is the only reason why the unified first-price auction - in which only first-price auctions are hosted, and intermediate winning bids are send to the next stage- was eventually chosen. Despotakis *et al.* [8] provide another interesting perspective centered on the publisher and the ad-exchanges.

Requests Duplication. If several ad exchanges are connected to the publisher, the DSPs receive duplicated bid requests, and hence bid several times for the same impression. Besides, some DSPs are directly connected to the publisher’s header bidding wrapper, which might increase even more the multiplicity of the requests.

In a nutshell, when an internet user reaches a publisher page containing some display inventory, it triggers a chain of calls that ends up in the servers of the buying side: the advertiser’s Demand Side Platform (DSP). The request is then handled by a programmatic bidding agent that implements the buyer’s bidding strategy. The request travels through intermediaries before getting to the end buyer’s server. Due to the complexity of the chain of calls and the multiplicity of intermediates (in particular, more than one intermediate can be plugged to the publisher page), it may happen that a buyer is called several times for the same display opportunity (see Figure 1): the bid requests are then said to be *duplicated*.

Bid request duplication is a challenging issue for display advertising buyers [9, 10]. We can picture the duplication problem this way: it is possible that

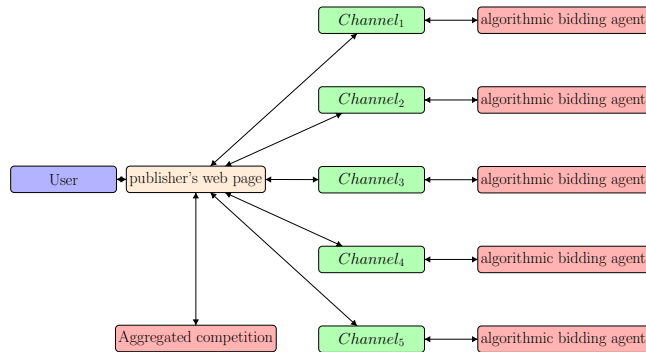


Figure 1: A user reaching a page triggers a chain of calls that end up in the buyer’s servers. At bidding time, the bidding agent might not know that other bidding agents of the same buyer are bidding for the same opportunity. Hence, his bid will be facing not only the aggregated competition, but also the bids of his fellow bidding agents. Moreover, if the agent cannot tell apart $Channel_1$ from $Channel_2$, then its bidding strategy is bound to be the same on those two channels.

several programmatic bidding agents of the same buyer receive a bid request without knowing whether the other agents have received a (duplicated) request and whether they are going to bid. Indeed, the time scale involved to answer the request is so short (less than 100 ms) that it can be technically impossible (or at least very challenging) for the buyer to synchronize the servers’ behaviors.

It is hard to assess the prevalence of duplication overall, since it is specific to the advertiser/publisher integration. But it is safe to assume that, at time of writing, it is more the rule than the exception.

Effect of Duplication on the Buyer’s Cost. This is a cause of sub-optimal bidding for two reasons: (a) the bidder mistakenly interprets a lost participation as a higher price to beat, while he may be in fact the winner of the auction on another request, which results in a bad estimation of the competitions, (b) by competing against themselves, bidders may increase their costs.

In a second-price auction, it is clear that if we bid more than once for the same opportunity, we end up paying our bid, as in first-price. This paper focuses on this *first-price effect*, which is specific to second-price auctions. Our goal is to characterize the optimal solution in a simplified setting: a unique second-price auction is resolved for the allocation of the display and the channels are not taking any margins.

Why not a Switch to Unified Second-Price?. We explain here why a unified second-price auction may make the life of the buyer much more complicated than expected. In the presence of duplications, the unified second-price auction is too complex for the bidder to be a good design choice.

In practice, there are many ad hoc business rules along the resolution of the consolidated auction that may slightly contradict with our setting. But we argue that, since the rules are evolving quickly and are not always fully transparent,

it is better to focus on one aspect of real auctions and discuss the challenges this aspect brings.

To the best of our knowledge, this is the first attempt to discuss the issue of duplicated requests in the literature, despite the fact that it has been a recurring topic in the industry, which has for now mostly relied on Supply Path optimisation (SPO) heuristics. Actions that could be taken by a bidder at the channel level mostly consists of (1) blacklisting the channel, (2) shade the bid on this channel or (3) answer to only a fraction of the bid requests (sampling) on this channel.

The optimal solution for the context we envision depends on whether the bid requests for a given opportunity can be identified by a unique feature such as the providing channel. As explained in [9, 10], this is not always the case in practice.

Contributions. The main contributions of this paper are: (1) the modeling of the bid request duplication issue in unified second-price auctions, (2) the resolution of the decentralized optimisation problem when the requests are all identical.

Agenda. After a brief introduction of the literature in the next section, we expose the bidding with duplicates problem in §4 and derive Lemma 1, a tool from which we will derive the optimal bid formula. We then characterize in §5 the solutions when the requests are all strictly identical. We follow up with a discussion.

2. Bidding in display advertising auction markets

We refer the reader to [3, 11] for general reviews on Real Time Bidding.

Our work contributes to the literature of marketplace design, but we mostly take the bidders perspective: the gap we attempt to fill relates to the last step of bidders' architecture. Hence, let us briefly indicate some pointers to the design of bidders. Chapelle *et al.* [12] provide a very precise description of an industrial architecture. Yet, the bidder they describe (1) uses a last touch attribution, (2) is operating in an incentive compatible world.

The question of attribution is a very active track of research. For instance, some researchers propose the use of multi-touch attribution [13, 14, 15] which aims at allocating conversion credits to each display. In [16], Diemert *et al.* derive from an attribution model a bid modifier that outperforms the benchmarks.

Besides, [12] does not take into account the business constraints of the advertiser (which have an impact on the value of the display opportunity). Another track of research relates to the study of the impact of the advertiser's business constraints on the optimal bid and the market. Zang *et al.* derive an optimal solution for the budget constrained bidding problem [17]. In [18], Balseiro *et al.* study the impact of budget constraints on advertiser behavior by combining queuing theory and mean field approaches. The notion of pacing algorithms

[19, 20] emerged from the need to dynamically implement such business constraints. Heymann [21] provides an analysis of the impact of CPA constraints in the buyer optimisation problem on the market outcomes.

In this work, we take for granted the fact that the bidder is able to estimate the economical value of the display opportunity. Hence, the questions of attribution and business constraints are orthogonal to ours.

Another category of research questions is the derivation of tools for transforming an economical value estimation into a bid. Ren *et al.* [22] propose to frame the problems of utility and competition estimation as well as the bid optimisation into a unique problem. Cai *et al.* [23] introduce a reinforcement learning based approach to take into account the competition landscape evolution. Nedelec *et al.* introduce a technique for the buyer to mitigate the loss of profit due to the seller’s reserve price optimisation strategy [24, 25]. Their presentation use functional analysis tool, while Tang *et al.*’s is based on quantiles [26]. Our work belongs to this category. We do not envision the use of reserve prices for the sake of simplicity. We also assume that we are able to estimate the competition.

Since our work studies a rule change on a marketplace, we should mention those studies on soft floor auctions [27], waterfall auctions [4], as well as this seminal work [28] on Generalized second-price auctions for Search. While those are relevant mostly from a general context perspective, [8] study the switch to first-price auction of display advertising markets, but with the perspectives of the ad exchanges and the publisher.

As already mentioned, we are not aware of any modeling attempts of the duplicated request problem. Because this problem can be seen as a decentralized control problem, this work owes some inspirational credits to [29], who introduced a powerful framework to model information in decentralized control.

3. The unified second-price auction

We assume the auction follows the following steps when an item is auctioned:

1. The seller (the publisher) sends a bid request to the ad exchanges
2. Each ad exchanges sends a bid request to its potential buyers
3. The potential buyers answer the ad exchanges bid requests
4. Each ad exchanges i discover its highest bid x_i and its second highest y_i
5. The x_i are sent to the seller as bid
6. The seller host a second-price auction. The highest x_i makes ad exchanges i win the item. The channel is billed the second highest $x^{(2)}$.
7. The highest bidder in ad exchanges i get the item, and is billed by the ad exchanges $\max(x^{(2)}, y_i)$

With such a payment rule, this setting is equivalent to the one drawn in Figure 1 from the bidder perspective.

4. The Duplicated Requests Bidding Problem

In this section, we construct a static analysis of the situation faced by the bidding agents in the presence of duplicated requests. We derive in Lemma 1 a formula of the buyer's expected utility that will be used in the subsequent sections.

We take the viewpoint of a buyer who values the opportunity $v > 0$ (hence, the distribution of v will not appear in the discussion). In the following, we assume everything is conditioned on the features of the opportunity: the competition, the buyer's bidding strategy, the distribution of requests we might receive for this opportunity at this point in time... In practice, the stochastic patterns of the environment are learned by machine learning algorithms that use the display opportunity features as input.

We focus on the case where the requests received at the same time for the same display opportunity are all identical. This can be caused by integration redundancy: an intermediary is plugged twice to the same inventory. In practice, two requests for the same opportunity may differ on channel related information (source), but we trade-off simplicity against generality, and focus on the closed form expression we derive.

The bidding strategy of the buyer is technically bound to be the same on requests that are identical. As we will see thereafter, the buyer needs to randomize its bid to reach optimality. We need to look for a solution in the class of distributions on $[0, v]$, which is a super-set of both the shading strategies and the sampling strategies. We denote by $K : t \in [0, v] \rightarrow [0, 1]$ the cumulative distribution of the buyer's bid b .

Let b^- be the highest bid of the competition in the consolidated auction (the price to beat), G its cumulative distribution and g its density distribution. We denote by b^i the i^{th} highest bid of the buyer. The buyer wins the auction whenever $b^1 > b^-$. In this case, by definition of the second-price auction, he will be billed $\max(b^-, b^2)$. The payoff of the buyer is its net utility

$$[b^1 > b^-] (v - \max(b^-, b^2)),$$

where for any Boolean variable X , $[X] \in \{0, 1\}$ and $[X] = 1$ when X (it is sometimes called the characteristic function of X and written $\mathbb{1}_X$).

The buyer's expected payoff maximization problem is

$$\max_K \mathbb{E}[b^1 > b^-] (v - \max(b^-, b^2)),$$

with b^- being distributed according to the cumulative distribution G , independently of the other bid.

We have now all we need to introduce an intermediate result that will be used in the subsequent sections.

Lemma 1 (Payoff). *The expected payoff of a buyer receiving n identical requests is*

$$\int_0^v K(t)^{n-1} G(t)n - K(t)^n (G(t)(n-1) + g(t)(v-t)) dt.$$

Proof. We compute separately the terms $\mathbb{E}[b^1 > b^-]$ and $\mathbb{E}[b^1 > b^-] \max(b^-, b^2)$. We remind the reader that the cumulative distribution of the maximum of n independent random variables is the product of their cumulative. For the first term we get

$$\int_0^v G(t) \frac{\partial K^n}{\partial t}(t) dt = G(v) - \int_0^v g(t) K(t)^n dt,$$

while the second term rewrites

$$\begin{aligned} & n \int_0^v k(t) \int_0^t u (\partial_u G K^{n-1})(u) du dt = \\ & n \int_0^v k(t) (tG(t)K^{n-1}(t) - \int_0^t G(u)K(u)^{n-1} du) dt. \end{aligned}$$

Then, by applying an integration by part on the two terms of the previous expression, we get

$$\begin{aligned} n \int_0^v k(t) t G(t) K^{n-1}(t) dt &= \int_0^v t G(t) \partial_t K(t)^n dt \\ &= vG(v) - \int_0^v (tg(t) + G(t)) K(t)^n dt \end{aligned}$$

and

$$n \int_0^v k(t) \int_0^t G(u) K(u)^{n-1} du dt = \int_0^v G(t) n (K(t)^{n-1} - K(t)^n) dt$$

We get the result by summing everything.

We combine regularization of K by convolution with smooth functions and the continuity of (1) with respect to K for the L^∞ norm to extend to the case where K has a discontinuity. \square

5. Optimal Bid

In this section, we assume a channel is calling the buyer several times (say n times) with exactly the same request, for the same opportunity. From a technical perspective, it may not be possible for the buyer to select one request to answer to, as the involved time scales are too short to allow for a grouping of all the requests for the same opportunity at bidding time. Thus, the buyer is restricted to apply the same strategy to all requests. The main results in this section tell us that the standard strategies (such as shading or sampling) may fail to find the optimal bid when requests are duplicated.

The following result can be deduced from the previous one.

Lemma 2. *We can restrict the bid (without loss of optimality) to be valued in the support of g .*

And now comes one of the main result.

Theorem 1 (Optimal Strategy for Identical Requests). *For a bidder receiving n identical requests, let*

$$H(t) = \frac{(n-1)G(t)}{g(t)(v-t) + (n-1)G(t)},$$

for t in the support of g . If H is non-decreasing, then it maximizes the buyer's expected payoff.

Proof. Let K be a maximizer of the buyer's payoff. For any $\epsilon \in [0, 1]$, $t \in [0, v]$, we define

$$K_{t,\epsilon} := (1-\epsilon)K(t) + \epsilon H(t).$$

Observe that for $i \in \mathbb{N}$

$$K_{t,\epsilon}^i - K_t^i = i\epsilon K(t)^{i-1}(H(t) - K(t)) + o(\epsilon).$$

Therefore the increment of payoff when replacing K by $K_{\cdot,\epsilon}$ is:

$$\begin{aligned} & \int_0^v \epsilon K(t)^{n-2}(H(t) - K(t))(n(G(t)(n-1) \\ & \quad + g(t)(v-t))K(t) - (n-1)G(t))ndt \\ & \quad + o(\epsilon) = \\ & \quad \epsilon n \int_0^v K(t)^{n-2}(G(t)(n-1) \\ & \quad + g(t)(v-t))(H(t) - K(t))^2 dt + o(\epsilon). \end{aligned}$$

So for ϵ small enough, this quantity is strictly positive if $K \neq H$ (on a non zero measure set), and $K_{\cdot,\epsilon}$ is admissible which is in contradiction with the optimality of K . QED \square

This result is counter intuitive: we insist that the optimal solutions identified here are not in the same class of functions as the one traditionally used to solve bidding problems. This result puts into perspective the intuition that shading or sampling is the right thing to do in the presence of duplicated identical requests. Moreover, it shows that despite the supposedly simplifying second-price rule, the computation of the optimal bidding strategy is quite complex.

What happens when H is not increasing? We build an example using the following probability density function: $g(t) = 0.01$ for $t < 0.25$ and $t > 0.75$, $g(t) = 1.99$ else. We plot the resulting functions g , G and H in Figure 2. Here is an intuition from optimal control theory: assuming the existence of the bid density distribution $k(t)$, the buyer's payoff is

$$n \int_0^v K(t)^{n-1} G(t) (1 + k(t)(v-t) - K(t)).$$




Figure 2: It is easy to build a distribution g (in red) so that the condition on H for Theorem 1 is not met. Indeed, in this example H cannot be a cumulative distribution (not monotone). However, the solution identified in Theorem 1 can be adapted (see discussion in the text).

Then, observe that optimizing the payoff is not that different from solving an optimal control problem with $k(t) \in [0, k_{max}]$ as a control and $K(t)$ as a state (one need to add a final constraint on $K(v)$, but it does not matter for our conclusion). Because the Hamiltonian of the system is affine in the control k the Pontryagin's Maximum Principle (PMP) indicates that either k is bang-bang (so we either have a Dirac of bid or no bid) or the term in factor of k in the Hamiltonian cancels. If we denote by p the costate, this implies that $nK^{n-1}Gk(v-t) = -p$. By the PMP, the time derivative of this quantity is equal to $n(n-1)K^{n-2}G(t)(1-k(t)(v-t)) - n^2K^{n-1}(t)$, which implies that $K(y) = H(t)$. Hence on an interval over which the control is not bang-bang, we shall have $K(t) = H(t)$.

To illustrate Theorem 1, take $v = 1$, $G(t) = t$, and $n = 2$. Then $K(t) = t$. One can check that the buyer's expected payoff is $1/3$ when using the randomized strategy, while he would only get $1/4$ by applying the optimal shading strategy.

We now pinpoint an easy extension:

Theorem 2 (Extension to stochastic number of requests). *If we receive n requests with probability p_n , then we can adapt Theorem 1 by setting*

$$H(t) = \frac{\sum_n p_n G(t)(n-1)}{\sum_n p_n g(t)(1-t) + G(t)(n-1)}$$

Proof. The adaptation of the proof of Theorem 1 is straightforward. □

6. Discussion

When the features of the bid requests allow the bidder to identify them individually, the bidder can then design one bidding strategy per type of request. We were not able to derive a closed-form solution in this situation.

probability \ channel	c_1	c_2	c_3
1/3	1	1	1
1/3	1	0	1
1/3	0	1	1

Figure 3: Here we suppose that we can define channel-specific strategies, and we illustrate on an example why greedy blacklisting approaches may fail to find the best blacklisting strategy. Suppose $v = 2$, the price to beat is 1 (constant), and there are 3 channels, c_1 , c_2 and c_3 . We receive a request from the three channels with probability 1/3. Channel c_1 is missing with probability 1/3 and c_2 is missing with probability 1/3. Observe that if we bid on only one channel, we might lose the opportunities that were not sent through this channel. On the other hand, if we bid everywhere, we might end-up second pricing ourselves. We suppose that we start with all channels turned on, and apply iteratively blacklisting/whitelisting decisions to improve our payoff. If we start with the three channels c_i turned on, then blacklisting channel c_1 or c_2 brings 1, while blacklisting c_3 brings 2. However, the best solution is to keep only c_3 . Hence a greedy search might miss the optimal blacklisting solution.




Figure 4: Here we suppose that we can define channel specific strategies. The expected payoff as a function of the bids on channel 1 and channel 2. We set the probability of being called by only one channel equal to 0.3 each, and the probability to be called by both at the same time equal to 0.4 and the competition is uniform on $[0,1]$ and $v = 1$. We note that even in this simple situation, the expected payoff is not convex. This illustrates the fact that finding an optimal shading strategy might be complicated in practice.

One important special case is the following. If requests across channels have different features, and hence there is one bidding strategy per request, then there is no need for the bidder to randomize. A shading strategy is enough. In this case the optimisation problem is not convex (cf. Figure 4). Even if the bidder restricts himself to blacklisting strategies, he may still miss the optimal. This is illustrated by the example in Figure 3. Hence while A/B tests appear quite attractive to decide which channels to turn on or off, they may fail to provide satisfying answers, as illustrated with the example. (More generally solutions with greedy approaches or local approaches are unlikely to succeed in finding the optimal channel selection)

We observe that if the bidder bids optimally, the randomization implies a loss of social welfare. The buyer with the highest valuation may not get the item.

Moreover, if the bidder uses a shading strategy, it may also be detrimental for the seller. If we take for instance a bidder facing a uniform distribution on $[0,1]$ for an opportunity valued at 1 and receiving two requests, then: (1) if the bidder does not react to the duplication, he will be paying 1 to the seller, (2) if the bidder randomizes according to $K(t) = t$, he would be paying $1/3$ on average. By comparison, (3) if the bidder were offered a real second-price rule, the bidder would be paying 0.5 on average, hence, in the presence of an informed buyer, it is not in the interest of the seller to send duplicated requests.

To our knowledge, this is the first academic work on bid request duplications. We show that standard strategies (shading, sampling, incremental tests) may fail to find the optimal bid in a unified second-price auction because of requests duplication. In particular, the truth-fullness property of the standard second-price auction is lost (even in a non-repeated setting). This undesired complexity is an argument in favor of the recent market move to unified first-price auctions.

Acknowledgement

I would like to thank Henry Jantet, Nicolas Chrysanthos, Aloïs Bissuel, Gilles Legoux, Jeremie Mary, Laure Alexandre and Colin Reingewirtz as well as the anonymous reviewers for their feedbacks and comments.

References

- [1] V. Krishna, Auction theory, Academic press, 2009. doi:10.1016/B978-0-12-426297-3.X5026-7.
- [2] T. Roughgarden, Twenty lectures on algorithmic game theory, Cambridge University Press, 2016.
- [3] H. Choi, C. F. Mela, S. Balseiro, A. Leary, Online display advertising markets: A literature review and future directions, Columbia Business School Research Paper (18-1) (2017).

- [4] B. Kveton, S. Mahdian, S. Muthukrishnan, Z. Wen, Y. Xian, Waterfall Bandits: Learning to Sell Ads Online (2019) 1–21arXiv:1904.09404.
URL <http://arxiv.org/abs/1904.09404>
- [5] M. Necheporenko, Second-Price Auctions And The Potential For Gaming The System — AdExchanger (2017).
URL <https://adexchanger.com/the-sell-sider/second-price-auctions-potential-gaming-system>
- [6] PubMatic, Understanding Auction Dynamics, a primer (2017).
URL <https://pubmatic.com/wp-content/uploads/2017/08/PubMatic-UnderstandingAuctionDynamics>
- [7] S. Sluis, Google Switches To First-Price Auction — AdExchanger (2019).
URL <https://adexchanger.com/online-advertising/google-switches-to-first-price-auction/>
- [8] S. Despotakis, R. R. Ravi, A. Sayedi, First-price auctions in online display advertising, 2019.
- [9] B. Epstein, Conflicting Truths About Auction Duplication — AdExchanger (2019).
URL <https://adexchanger.com/ad-exchange-news/conflicting-truths-about-auction-duplication/>
- [10] S. Sluis, DSPs To SSPs: ‘Clean Up Or We Cut You Off’ — AdExchanger (2017).
URL <https://adexchanger.com/platforms/dsps-ssps-clean-cut-off/>
- [11] J. Wang, W. Zhang, S. Yuan, Display advertising with real-time bidding (rtb) and behavioural targeting, Foundations and Trends® in Information Retrieval 11 (4-5) (2017) 267–435. doi:10.1561/15000000049.
URL <http://dx.doi.org/10.1561/15000000049>
- [12] O. Chapelle, Modeling delayed feedback in display advertising, Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2014) 1097–1105doi:10.1145/2623330.2623634.
URL <https://dl.acm.org/citation.cfm?id=2623634>
- [13] Y. Zhang, Y. Wei, J. Ren, Multi-touch Attribution in Online Advertising with Survival Theory, Proceedings - IEEE International Conference on Data Mining, ICDM 2015-Janua (January) (2015) 687–696. doi:10.1109/ICDM.2014.130.
URL <https://ieeexplore.ieee.org/abstract/document/7023386/>
- [14] W. Ji, X. Wang, D. Zhang, A Probabilistic Multi-Touch Attribution Model for Online Advertising, dl.acm.org (2016) 1373–1382doi:10.1145/2983323.2983787.
URL <https://dl.acm.org/citation.cfm?id=2983787>

- [15] X. Shao, L. Li, Data-driven multi-touch attribution models, in: Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2011, pp. 258–264.
- [16] E. Diemert, J. Meynet, P. Galland, D. Lefortier, Attribution Modeling Increases Efficiency of Bidding in Display Advertising, Tech. rep. (2017). URL <https://arxiv.org/pdf/1707.06409.pdf>
- [17] W. Zhang, S. Yuan, J. Wang, Optimal real-time bidding for display advertising, Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '14 (2014) 1077–1086doi:10.1145/2623330.2623633. URL <http://dl.acm.org/citation.cfm?doid=2623330.2623633>
- [18] S. R. Balseiro, O. Besbes, G. Y. Weintraub, Repeated auctions with budgets in ad exchanges: Approximations and design, Management Science 61 (4) (2015) 864–884.
- [19] V. Conitzer, C. Kroer, D. Panigrahi, O. Schrijvers, E. Sodomka, N. E. Stier-Moses, C. Wilkens, Pacing equilibrium in first-price auction markets, arXiv preprint arXiv:1811.07166 (2018).
- [20] V. Conitzer, C. Kroer, E. Sodomka, N. E. Stier-Moses, Multiplicative Pacing Equilibria in Auction Markets, Springer (2017). URL <http://arxiv.org/abs/1706.07151>
- [21] B. Heymann, Cost per action constrained auctions, in: Proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation, ACM, 2019, p. 1.
- [22] K. Ren, W. Zhang, K. Chang, Y. Rong, Y. Yu, J. Wang, Bidding Machine: Learning to Bid for Directly Optimizing Profits in Display Advertising, IEEE Transactions on Knowledge and Data Engineering (2018). doi:10.1109/TKDE.2017.2775228.
- [23] H. Cai, K. Ren, W. Zhang, K. Malialis, J. Wang, Y. Yu, D. Guo, Real-Time Bidding by Reinforcement Learning in Display Advertising, in: Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, ACM, 2017, pp. 661–670. doi:10.1145/3018661.3018702. URL <http://arxiv.org/abs/1701.024903018661.3018702>
- [24] T. Nedelec, M. Abeille, C. Calauzènes, N. E. Karoui, B. Heymann, V. Perchet, Thresholding the virtual value: a simple method to increase welfare and lower reserve prices in online auction systems, arXiv preprint arXiv:1808.06979 (2018).
- [25] T. Nedelec, N. E. Karoui, V. Perchet, Learning to bid in revenue-maximizing auctions (2 2019). URL <http://arxiv.org/abs/1902.10427>

- [26] P. Tang, Y. Zeng, The Price of Prior Dependence in Auctions, dl.acm.org (2018) 485–502doi:10.1145/3219166.3219183.
URL <https://dl.acm.org/citation.cfm?id=3219183>
- [27] R. Zeithammer, Soft floors in auctions, Management Science 65 (9) (2019) 4204–4221. doi:10.1287/mnsc.2018.3164.
URL <http://pubsonline.informs.org/journal/mnsc/http://orcid.org/0000-0002-8388-247X>
- [28] B. Edelman, M. Ostrovsky, M. Schwarz, Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords, American economic review 97 (1) (2007) 242–259.
- [29] H. S. Witsenhausen, On Information Structures, Feedback and Causality, SIAM Journal on Control 9 (2) (2005) 149–160. doi:10.1137/0309013.
URL <http://epubs.siam.org/doi/10.1137/0309013>