Online Display Advertising Markets:

A Literature Review and Future Directions

Hana Choi* Carl Mela† Santiago Balseiro‡ Adam Leary§

November 13, 2017

Abstract

This paper summarizes the display advertising literature, organizing the content by the agents in the display advertising ecosystem, and proposes new research directions. In doing so, we take an interdisciplinary view, drawing connections among diverse streams of theoretical and empirical research in marketing, economics, operations, and computer science. By providing an integrated view of the display advertising ecosystem, we hope to bring attention to the outstanding research opportunities in this economically consequential and rapidly growing market.

Keywords: Display Advertising, Literature Review, Real-Time Buying, Programmatic Buying, Real-Time Bidding

JEL Classification Codes: M3, L1

*Hana Choi is Ph.D. student at the Fuqua School of Business, Duke University (email: hana.choi@duke.edu, phone: 734-834-0699).

†Carl F. Mela is the T. Austin Finch Foundation Professor of Business Administration at the Fuqua School of Business, Duke University (email: mela@duke.edu, phone: 919-660-7767, web: faculty.fuqua.duke.edu/˜mela).

‡Santiago R. Balseiro is Assistant Professor of Business Administration at the Graduate School of Business, Columbia University (email: srb2155@columbia.edu, web: www.columbia.edu/˜srb2155).

§Adam Leary is VP, Data Science Services at CBS Interactive (email: adam.leary@cbsinteractive.com, phone: 415-344-2645).

The authors thank Garrett Johnson for helpful comments and suggestions.

1 Introduction

According to the Interactive Advertising Bureau (IAB herein) report, display advertising (ad herein) totaled \$31.7 billion in the US in 2016, up from the \$24.9 billion reported in FY 2015.[1] This double-digit growth rate is fueled by the upswing in mobile browsing, video ad formats, and the developments in targeting technology. Digital display ad spending is forecasted to surpass even search ad spending and will continue its rapid ascent in share and significance.[2]

Given the substantial economic importance of the display ad market, the objective of this paper is to summarize the existing literature from the perspective of each player involved in the display ad ecosystem. In doing so, we take an interdisciplinary view spanning diverse streams of theoretical and empirical research in marketing, economics, operations, and computer science. We focus on literature that explicitly emphasizes operations of the display ad market as opposed to a more general discussion on related topics, including the large body of literature on ad, sponsored search ad, auctions, mechanism design, and online algorithms.

A handful of papers survey various aspects of online ad in general, covering search, display, classified, and other forms of ads. Ha 2008 reviews online advertising research published specifically in six advertising journals.[3] Evans 2008, 2009 discusses the evolution of internet ad and provides industry perspectives, while Bucklin and Hoban 2017 focus on marketing models developed to improve decision-making in internet advertising. Due to the broad scope of these papers, the distinguishing characteristics of the display ad market (such as the issues arising from the coexistence of guaranteed and non-guaranteed selling channels) and its ecosystem are precluded.

Focusing on the display ad market, existing survey papers either discuss issues facing advertisers (demand side), publishers (supply side), or intermediaries. On the advertiser side, Tucker 2012b and Goldfarb 2014 review the economic literature on targeting and privacy concerns, and highlight their trade-off in ad effectiveness. On the publisher side, Korula et al. 2016 cover a broad range of publishers' decisions in selling display ad, and review studies on allocating impressions and designing contracts/auctions. On the intermediary side, Muthukrishnan 2009 considers issues faced by ad exchanges, which are online platforms that match advertisers with publishers. One exception

[1] https://www.iab.com/wp-content/uploads/2016/04/IAB_Internet_Advertising_Revenue_Report_FY_2016.pdf

[2] https://www.forrester.com/report/US+Digital+Marketing+Forecast+2016+To+2021/-/E-RES137095

[3] Ha 2008 reviews articles published in *Journal of Advertising, Journal of Advertising Research, Journal of Current Issues and Research in Advertising, International Journal of Advertising, Journal of Marketing Communications*, and *Journal of Interactive Advertising.*

is Wang et al. 2016, who discuss recent algorithms and challenges in computational advertising from both advertisers' (bidding strategies in real-time buying) and publishers' (pricing and auction mechanism) perspectives, focusing on the non-guaranteed display ad selling channel.

Our survey differs in that we provide a comprehensive view of the display ad ecosystem, including both guaranteed and non-guaranteed selling channels, and delineate the research progress gained from the eyes of all the players involved, including advertisers, publishers, and intermediaries. Furthermore, we take an interdisciplinary view and draw connections across disciplines. By providing an integrated view and bringing attention to the outstanding research opportunities, we hope to provide additional motivation for subsequent studies in this rapidly growing market.

The remainder of the paper is organized as follows. Section 2 overviews the ecosystem and highlights key unique features of the display ad market. Next, we organize existing literature by advertisers (Section 3), publishers (Section 4), and intermediaries (Section 5). In each sub-section, we address key managerial decisions and discuss related research issues and opportunities. Before concluding in Section 7, issues related to transparency are briefly discussed in Section 6.

2 The Display Advertising Ecosystem

This section enumerates the players, the selling channels, and the key characteristics of the display ad market. Players' objectives, roles, and available information are briefly discussed to place in context the research topics discussed in subsequent sections.

2.1 Players

The display ad market is a two-sided market: on one side, advertisers procure ad impressions on publishers' websites to reach the right segment of consumers; and on the other side, publishers with consumers' impressions purvey their ad inventory to advertisers with the highest valuations for those impressions. In between, intermediaries facilitate the match between advertisers and publishers by managing data and providing optimization tools and algorithms for serving ads.

2.2 Selling Channels

Ad exposures in display markets are sold via guaranteed and non-guaranteed selling channels. The guaranteed selling channel (so-called direct sale) involves the advance sale of a number of impressions at a fixed price. Non-guaranteed sales occur in real time via ad exchanges (so-called

Real-Time Buying, RTB herein).[4] The non-guaranteed selling channel is estimated to be 34% ($\approx 78\% * 44\%$) of the total display ad revenue in 2017.[5]

The main characteristics distinguishing the guaranteed selling channel from RTB are (i) the pricing mechanism (fixed price in the guaranteed vs. auction in RTB), (ii) the information available about the impressions (thus targeting ability), and (iii) the players involved (especially the level of intermediation). These distinguishing characteristics are described for each selling channel next.

2.2.1 Guaranteed Selling Channel

The guaranteed selling channel involves the buying and selling of bundles of impressions through a guaranteed contract, prior to users' arrivals on the publisher's site. In the guaranteed selling channel, an advertiser and a publisher negotiate a fixed price pertaining to when, where, and how the ads will be displayed. These contractual arrangements guarantee the number of impressions to be delivered satisfying certain targeting criteria, at a negotiated fixed price (cost-per-mille, CPM), during a specified time period (e.g., 1M impressions to female users in SF during July for $5000).[6]

Advertisers often work with ad agencies to facilitate media planning/buying and the negotiation process. On the publisher side, there often is an in-house sales team who negotiates the terms and prices of the contract and manages the relationship with the advertisers. Recently, many aspects of the contracting process are being automated with programmatic buying technologies.[7] The so-called *programmatic direct* (or programmatic guaranteed, automated guaranteed) enables publishers to open up their guaranteed deals even to the advertisers with smaller budgets who otherwise would not meet the minimum ad spend required to allocate the limited resources of the sales team. The advent of *programmatic direct* thus decreases the role of sales teams and lowers advertisers' search costs associated with finding prices from multiple publishers, which in turn would affect advertisers' optimal ad buying strategies across selling channels (§3.2.1).[8]

Because guaranteed contracts are drawn in advance, both advertisers and publishers rely on

[4] We use "non-guaranteed selling channel" and "RTB" interchangeably in this article.
[5] https://www.emarketer.com/Article/eMarketer-Releases-New-Programmatic-Advertising-Estimates/1015682
[6] Instead of guaranteeing number of impressions to be delivered, contracts may alternatively be written to guarantee delivery dates or share of impressions (Danaher et al. 2010) at a fixed price. Although less common, publishers may guarantee the number of clicks to be delivered based on cost-per-click (CPC).
[7] The terminology "programmatic buying" should not be confused with non-guaranteed RTB. Programmatic buying consists of a wide range of technologies that *automate* the buying, selling, and matching of ad.
[8] Depending on the publisher, negotiation is altogether precluded in *programmatic direct* to facilitate the automated contracting process. For example, the quoted price in Facebook's "Reach and Frequency Buying" is a take-it-or-leave-it fixed price.

their expectations about future impressions and associated user characteristics (of those who will arrive during the planned campaign period) when making contractual decisions. Because targeting criteria are typically coarse, advertisers end up buying a "bundle" of impressions at a fixed price. For an advertiser, important questions are when/how should this guaranteed bundle of impressions be bought and which targeting criteria to use (§3.2). For a publisher, the questions are whether to sell impressions in advance as a bundle and how to price the guaranteed inventory (§4.2).

2.2.2 Non-Guaranteed Selling Channel

The non-guaranteed selling channel involves the buying and selling of single impressions in real time via programmatic buying. It is considered RTB, because an advertiser's buying decision is made immediately after the user arrives on the website. It is non-guaranteed, because an auction takes place for the incoming impression and the advertiser's ad is served to the end user only if the advertiser wins the auction. The ecosystem in RTB is more complex than the guaranteed selling channel, with many specialized intermediaries lying in between. These intermediaries provide the fundamental infrastructure required for selling, buying, and serving ads in real time, typically within milliseconds.[9] Although little research exists to date on the role of these intermediaries, the overall welfare of the market may improve by allocating impressions to advertisers with higher valuations and by reducing search costs in price discovery (§5).

In the pursuit of monetizing (unsold) impressions and enhancing targeting, *ad exchanges* (e.g., DoubleClick Ad Exchange, OpenX), have been introduced where impressions can be bought and sold in real time. Prior to the advent of RTB, *ad networks* (e.g., Google Display Network) assumed a mediating role by aggregating (unsold) inventories across publishers and packaging them into effective targeting and large-scale audience-buying opportunities for the advertisers. Similar to those of the ad exchanges, many ad networks also provide real-time bidding features nowadays.[10]

Ad exchanges typically conduct a second-price, sealed-bid auction for each available impression in real time. As such, publishers can monetize beyond the guaranteed contracted volume. Likewise, advertisers can now buy impressions based on user-specific (e.g., cookie) information. An important decision faced by ad exchanges relates to the choice of auction mechanism, which can influence the allocation of impressions and the share of revenues between advertisers and publishers (§5).

[9] http://data.iab.com/ecosystem.html

[10] https://www.exchangewire.com/blog/2015/06/29/the-decline-of-the-ad-network/

Intermediaries have been diversifying their roles to meet the technological needs of advertisers and publishers in RTB. *Demand side platforms* (DSPs herein; e.g., MediaMath, DataXu, Turn, Rocket Fuel, Adobe Advertising Cloud), for example, help advertisers by facilitating the real-time bidding process. Recognizing that bid optimization is a difficult managerial problem for the advertisers, DSPs specialize in calculating and submitting bids based on the users' behavioral data and on targeting criteria provided by the advertiser. Questions advertisers (and DSPs) face in determining the optimal bid include how to calculate the value of an impression, how to leverage users' past behavioral data, and how to react to competitors' bidding strategies (§3.2.3).

Similarly on the publisher side, *Supply side platforms* (SSPs herein; e.g., DoubleClick for Publishers) facilitate publishers' inventory management and yield optimization. Publishers seek to maximize revenue by efficiently allocating incoming impressions to advertisers. SSPs, on behalf of the publishers, decide whether an impression should be allocated to an advertiser with a guaranteed contract or to an advertiser bidding in the RTB. Moreover, they seek to determine how to allocate impressions across advertisers in the guaranteed selling channel, given that the contracts differ in prices, campaign schedules, targeting criteria, and delivery status. Questions remain open regarding the implications of different allocation mechanisms in the presence of competing advertisers and publishers who can enter both the guaranteed selling channel and RTB (§4.4).

Both advertisers and publishers strive to manage and leverage consumer data efficiently for targeting, but have access to different information. Advertisers, for example, can augment their own proprietary data (e.g., purchase patterns on the advertiser's website) with additional data bought from third parties (e.g., income level, job history, home ownership, monthly car payment) in order to build targeting audience profiles. Publishers seek to integrate their own proprietary data (e.g., user registration information and ad viewing patterns on the publisher's website) with third-party data to build user segments and to offer better targeting options to advertisers at potentially higher prices. To this end, *Data Management Platforms* (DMPs herein; e.g., Salesforce DMP, Adobe Audience Manager, MediaMath, Neustar, Oracle DMP) provide an accessible interface to import data from multiple sources and to build audience segments from the integrated data, which are then fed into DSP/SSP toward optimizing downstream decisions. For both advertisers and publishers, important research questions pertain to the optimal level of targeting granularity (to buy or to sell) and to which additional information to buy from third parties and at what price

(§3.3.3, §4.3). Asymmetry of information also raises questions regarding how much information to share and the role of intermediaries in resolving this asymmetry (§5).

2.2.3 Interrelation of Guaranteed and Non-Guaranteed Selling Channels

Advertisers and publishers jointly consider the two selling channels (guaranteed and non-guaranteed) when making display ad decisions. The price and volume guaranteed in advance will inevitably affect the supply, demand, and equilibrium outcome in RTB, and vice versa. A question of fundamental theoretical interest is why dual selling channels co-exist in equilibrium. From the advertiser point of view, why would an advertiser buy impressions in advance at a fixed price, when RTB through an auction offers higher targeting abilities? (§3.2.1) From the publisher point of view, would selling ads through both selling channels yield higher profits than selling through one? (§4.2.3, §4.4) In terms of intermediaries who enable the existence of RTB, what are their welfare implications on advertisers, publishers, and the market as a whole? For example, does the introduction of RTB reduce prices in the guaranteed selling channel, thus benefiting advertisers? Or does it benefit publishers by better matching advertisers and users, and encouraging competition among advertisers through auctions? (§5) We address these and other questions next.

3 Advertisers

Media planning and execution constitute the demand side of the display ad market. Advertisers (i) set objectives (§3.1) and (ii) buy a set of impressions toward maximizing the objectives given budget constraints (§3.2). Throughout the ad campaign, marketers (iii) measure return on investment (ROI herein) and use it as a reference point in setting future objectives and budget (§3.3).

3.1 Objective

Characterizations of advertising have suggested twofold goal types (e.g., Zhu and Wilbur 2011). One type involves building long-term brand equity. The other type is to generate short-term direct responses similar to coupons in print ads. As display ad objective is not directly observed by the researchers in most datasets, integrated studies about what factors drive this twofold objective and the relative importance of goal types are scant. As display ad objective drives advertisers' downstream decisions (such as budgeting and ad buying) and is antecedent to modeling, understanding this twofold objective is a pertinent area of future research.

Having the advertising objective in mind, both advertisers and researchers rely on key performance indicators (KPIs herein) in evaluating campaign performances or setting the (implementable) objective function for optimization algorithms. With the advancement in tracking technology, a large number of KPIs are defined and monitored nowadays, including impressions (Danaher et al. 2010), clicks (Chatterjee et al. 2003), visits (Dalessandro et al. 2015), and purchases (Manchanda et al. 2006). Consequent to the advertising objective of increasing brand equity and direct responses, the question arises of which metrics should be evaluated and optimized. For example, should advertisers contract to buy a guaranteed number of impressions or clicks? Do the time spent on a page and the number of page views constitute good measures for advertisers focusing on long-term branding? Should survey-based metrics such as brand awareness, attitude, recognition, or recall be collected pre- and/or post-campaign (Goldfarb and Tucker 2011a; Bart et al. 2014), and if so, how can this data be leveraged in the optimization process? Another interesting question relates to which KPIs facilitate effective learning about ad responses, especially with advertisers' wider use of A/B testings and machine learning algorithms toward achieving optimal ad buying decisions. Future research can further explore the relationship among advertisers' display ad objectives, optimization processes, and various KPIs evaluated in practice.

3.2 Ad Buying

After the advertising objective is set, advertisers face the following ad buying decisions: (i) whether to buy guaranteed or non-guaranteed, (ii) ad design and inventory characteristics, (iii) user characteristics (targeting), and (iv) scheduling (frequency and timing).

3.2.1 Selling Channels: Guaranteed vs. Non-Guaranteed

Most advertisers engage both selling channels by allocating a portion of their budget in the guaranteed selling channel in advance, and then spending the rest in RTB. Unfortunately, the existing literature on display ad is largely silent on why/when one selling channel is/should be chosen over the other. Below we discuss some plausible explanations, which future research can explore.

Some reasons why advertisers may favor RTB include finer targeting and lower search costs. When buying in real time, the advertiser can submit an appropriate bid reflecting the expected value of the impression based on the revealed user-specific information. Further, procuring a guaranteed contract entails time-consuming negotiation processes with potentially multiple publishers. These search costs are significantly reduced in RTB, as advertisers can reach a large number of publishers through a centralized intermediary exchange market.

On the other hand, guaranteed contracts are often favored when advertisers have a long-standing relationship with the publisher that facilitates customization and the negotiation process as a whole. If advertisers are risk averse, paying a premium to buy a guaranteed inventory in advance helps mitigate uncertainty in either the auctions' outcomes or the amount of impressions that will be available in RTB on specific dates. Lastly, advertisers highly concerned with ensuring brand safety will choose guaranteed contracts so their ads appear on high-quality, reputable websites.

Athey et al. 2017 demonstrate that with consumers' multi-homing across publishers and imperfect tracking technology, advertisers seeking broader reach would favor larger publishers to avoid inefficient duplication, which can create an incentive for contracting guaranteed deals.[11] As a supply side explanation, Sayedi 2017 explores dynamic allocation (see §4.4), under which having both selling channels yields higher profits than selling in the guaranteed selling channel or RTB alone.

A large number of positive and normative research questions remain regarding dual selling channels. Positively speaking, an issue of interest is what the rationales are behind advertisers' observed selling channel choices and their relative impacts. For example, why would some advertisers invest more heavily in the guaranteed contracts and others more heavily in RTB? Normatively speaking, when would it be beneficial for the advertisers to choose one selling channel over the other, and what will be the optimal buying strategy?

3.2.2 Ad Design and Inventory Characteristics

In addition to the selling channels, advertisers make decisions regarding ad design and inventory characteristics.[12] These involve choices about (i) websites including traditional media (Danaher 2007; Danaher et al. 2010; Aksakallı 2012), blogs, and social media (Zubcsek and Sarvary 2011; Bakshy et al. 2012; Tucker 2012a; Lee et al. 2017); (ii) positions; (iii) sizes; (iv) creatives (Urban et al. 2013; Bruce et al. 2017); (v) formats including banner, text (e.g., Google AdSense), rich media, digital video (Tucker 2014b), and native ads (Wojdynski and Evans 2016); and (vi) devices including desktop, mobile (Okazaki et al. 2007; Ghose et al. 2012; Bart et al. 2014)[13], and tablets.

Past research has typically evaluated ad design and inventory characteristics independently

[11] In their working paper version (Athey et al. 2014), Section 7 explores publishers offering guaranteed versus non-guaranteed deals. Their analysis provides a valuable insight into the mechanism in which advertisers may have an incentive to pursue guaranteed contracts to reach a wider audience.

[12] See Ha 2008 and references therein for earlier research on advertisers' ad design and inventory decisions.

[13] See Grewal et al. 2016 for a survey on mobile advertising, including display ads.

to ascertain their respective effects on traditional brand equity measures (e.g., awareness, recall, recognition, attitude) and/or direct response measures (e.g., CTR, purchase intent). However, relatively little effort has been apportioned toward assessing the inter-relationship among these various inputs or their efficacy with respect to cost (Johnson and Lewis 2015). As ad design and inventory characteristics work jointly, optimization entails taking a combination of websites, positions, sizes, ad creatives, formats, and devices into account, and weighting the (incremental) benefits and costs of these options toward achieving the display ad goal given the budget constraint.

Furthermore, the choices of ad design and inventory characteristics are closely tied with the choices of selling channels. In RTB, standardization is necessary to ensure automation and seamless delivery within milliseconds. As a result, ad exchanges generally support inventories that conform to the IAB standard guideline, and customizability in terms of ad position, size, format, or content environment is limited.[14] For example, highly customized ad positions and environments, such as homepage takeover or conquest ads, are generally supported only via guaranteed contracts.[15] Goldfarb and Tucker 2014 provide evidence that ad recognition declines as the use of IAB standard formats rises, presumably because standard format ads attract less attention and are harder to distinguish from competitors' ads. This raises the question of how the level of standardization or the demand for customizability affect ad buying decisions and market outcomes.

From an implementation perspective, scalability, computational efficiency, and accuracy are key considerations in building optimization algorithms. An optimization algorithm taking both selling channels into account is imperative, as advertisers now have an additional option of going to RTB and bidding for each impression with user-specific information. For computational efficiency, advertisers might be making these decisions hierarchically in practice (e.g., allocating budget across selling channels first, then across websites, then devices and formats, etc.) and an easy-to-implement optimization algorithm might be considered upon these heuristics. Finally, as A/B testing becomes more affordable, incorporating A/B testing into the optimization algorithms in a test-and-learn fashion is an interesting avenue for future research.

[14] https://www.iab.com/guidelines/

[15] Best Buy, for example, can *takeover* CNET's entire front page on Black Friday and display its creatives with unique design and architecture. Or Samsung can place *conquest* ads next to publisher's review article on Apple's newly launched iPhone, to reach those considering a competitor's product.

3.2.3 Target User Characteristics

Advertisers often target ads to users whose purchase might be imminent (generating direct responses) or who might potentially become valuable customers in the future (building brand equity or customer lifetime value). Central to this goal is the identification of targets conditioned upon the available information. In *demographic targeting,* users are segmented based on age, gender, and geography. *Contextual* advertising, on the other hand, targets based on users' interests; for example, by matching ad product category to the website content the user is browsing. In *behavioral targeting*, users are sought after based on their past behavioral patterns; for example, users who visited an advertiser's website in the past. We begin with targeting in the guaranteed setting, then move to the discussion of targeting in RTB.

Targeting Under Guaranteed Selling Channel For behavioral targeting, user profiles are generated based on users' topic interests (Ahmed et al. 2011; Trusov et al. 2016), conversion intents (Chung et al. 2010; Pandey et al. 2011; Aly et al. 2012), or brand affinity (Provost et al. 2009), and the preferred users are selected as the target audience. Contracting upon such preferred user segments several months in advance via guaranteed contracts, however, is often difficult as users' characteristics (e.g., topic interests) vary considerably over time. From a publisher standpoint, committing to deliver a certain number of impressions of users demanded by the advertiser may increase the risk of under-delivery, as the publisher may not have enough information to make good predictions on those users' future impressions. Although behavioral targeting based on the advertiser-side information (e.g., users identified as having high purchase intent based on previous visits to advertiser's website) is not common, some publishers offer behavioral targeting based on the publisher-side information (e.g., users with frequent visits to publisher's website). Presumably the publisher tries to price discriminate and increase profits by offering those user segments whose future impressions are well predicted by the publisher based on its own data, and whose exposures are highly valued by the advertisers.

In terms of demographic and contextual targeting, advertisers often have good insights on which demographic groups and contents are compatible with the ad product or service. For example, the advertiser may want to target users living close to their brick-and-mortar store locations (Johnson et al. 2016a). The question advertisers often face is whether it is profitable to pay for an additional layer of targeting. Although not much research has studied advertisers' user targeting decisions in the guaranteed selling channel, this is an eminent area for future research, especially in conjunction with RTB. For example, when targeting in guaranteed contracts is limited, advertisers might opt to buy inventory in RTB auctions where targeting is finer. The shift in bargaining power between advertisers and publishers will affect how much information (targeting) is provided and the price and demand for such targeting in the guaranteed selling channel.

Targeting Under Non-Guaranteed Selling Channel In RTB, advertisers have considerably more information and control over targeting. Each impression is sold separately in real time, meaning that the value of an impression can more readily be evaluated by matching the incoming user with the advertiser's own data on users' past behavior. In addition to the demographic and contextual targeting, advertisers greatly enjoy behavioral targeting in RTB.

There is a growing stream of literature regarding optimizing bids in RTB. In order to determine optimal bids, two components are required: (i) the advertiser's own valuation and (ii) distribution of the highest bids among competing advertisers. Taking one step further, advertisers take the overall utility from a campaign into account when calculating the optimal bids. That is, the advertiser's objective function will be an aggregation of each impression's utility to reflect the advertiser's specific goal. We first review papers evaluating advertisers' valuations and the distribution of others' bids, then move to the discussion on the optimal bidding strategy in consideration with the advertiser's display ad goal and budget constraint.

1) **Advertiser Valuation**s

Understanding valuations entails predicting the value of an impression in real time as an auction is conducted for an arriving impression. Toward this end, prior research focuses upon the prediction of direct ad responses (e.g., CTR, conversion rates) and the building of machine learning algorithms to minimize prediction errors. Some computational challenges include that (i) the dimensionality of user attributes space is high and data is sparse across these dimensions, (ii) conversion events occur rarely, and (iii) cold-start situations with limited data history hinder reliable predictions. In general, the literature has evolved to address these challenges and to lower computational costs while improving predictions. Linear models (Chen et al. 2009; Agarwal et al. 2010; Agarwal et al. 2014; McMahan et al. 2013; Chapelle et al. 2015) are widely used in practice as they are easy to implement and perform effectively even with large-scale datasets. As it is difficult to capture higher order interactions with linear models, hybrid methods are being developed, incorporating factorization (Menon et al. 2011; Oentaryo et al. 2014a; Juan et al. 2016) or decision trees (He et al. 2014), and deep learning methods are being applied to further explore latent patterns (Zhang et al. 2016). The learnt user characteristics from user profiling (Zhang et al. 2014; Perlich et al. 2014), and features additionally inferred from social data (Bagherjeiran and Parekh 2008; Liu and Tang 2011; Goel and Goldstein 2013) are also leveraged in improving ad response predictions.

One promising future avenue is the incorporation at scale of additional predictors such as location and social surroundings. For example in the mobile SMS ad context, users are shown to be responsive to an offer more in a crowded than a non-crowded environment (Andrews et al. 2015) and the interaction between proximity to the focal venue and the time a promotion is received affects purchase likelihood (Luo et al. 2013).

2) **Distribution of the Highest Bids**

The distribution of competing advertisers' highest bids governs winning probabilities and expected payments. A bidder's winning probability is also linked to the expected number of impressions delivered, and this number affects the *reach* metric for a given campaign. Challenges in predicting competing advertisers' highest bids include that (i) impressions possess demographic, contextual, and behavioral attributes for which competing advertisers have heterogeneous valuations, and (ii) because competitors' bids are not disclosed, the highest of the competing bids (and the reservation price set by the publisher) is only observed from the winning payment, and the lower bound of competing advertisers' highest bids is inferred upon losing from the advertiser's own bid. Toward addressing the first point, Cui et al. 2011 assume the winning bid distribution to be a mixture of log-normal distributions, where the mixture weights reflect different targeting features. The second point can be addressed by extending a censored regression model (Wu et al. 2015) or survival model (Wang and Zhang 2016) to correct for the bias in learning due to the difference in training data (observed, censored winning price history) and the testing data.

Limited data available to advertisers raises a question of how much information should be shared between intermediaries and advertisers (and also publishers). For example, ad exchanges, who have information on all bids and associated targeting criteria, may act as an information aggregator and provide bid landscape reports to advertisers (and/or publishers). Disclosing bid information can range from revealing all bids and identity of the bidders to revealing simple summary statistics as quantiles. Advertisers may or may not consent to such disclosure and receiving more information about competitors. Most of the recent debate in display space has focused on transparency, and more studies on economic consequences of information asymmetry is called for (§5).

3) **Optimal Bidding Strategy**

In theory, the weakly dominant strategy for an advertiser in a private value, second-price auction of a single object is bidding truthfully. However, bid calculation becomes more complex as additional practical constraints are considered: the advertiser (i) faces a budget constraint for a given campaign in repeated auctions (Balseiro et al. 2015; Balseiro and Gur 2017), (ii) learns own and/or other's true valuations over time (Iyer et al. 2014; Cai et al. 2017), (iii) sets a number of impressions to attain (Ghosh et al. 2009b), and (iv) sets pacing options so that the budget is spent smoothly over a specified time period (Lee et al. 2013; Yuan et al. 2013; Xu et al. 2015).[16] With such considerations, the optimal bidding strategy is not necessarily truth telling. For example, Balseiro et al. 2015 show that the optimal bidding strategy for an advertiser facing a (binding) budget constraint is to shade values to account for the option value of future opportunities.

When designing the optimal bidding algorithm, different assumptions are made about the advertiser's objective function and the bidding environment, including the ways predictions are made about valuations and the distribution of highest bids. One common approach predicts statistics of interest (e.g., CTR for ad responses, moments for the distribution of highest competing bids) as a preliminary step toward minimizing prediction errors, and then using these predictions as inputs in calculating the bids (Perlich et al. 2012; Zhang et al. 2014; Xu et al. 2016). However, when advertisers are highly uncertain about ad responses or the distribution of the highest bid, incorporating the prediction step into the entire bidding optimization task is desirable to enhance learning about the overall utility from a campaign (Hummel and McAfee 2015; Chapelle 2015; Ren et al. 2016).

Because training the model is computationally expensive, traditionally training has been done offline to find the optimal bidding strategy, which is then applied to the real-time data. As the campaign progresses, advertisers receive additional data on winning outcomes with corresponding payments and user responses. In order to incorporate such new information or to dynamically

[16] Spending the budget relatively smoothly over the campaign period (or over a given day or part of a day) is considered to help advertisers attain a wider range of audience. However, the commonly offered pacing option by the DSP, *uniform* pacing, may not be optimal as spending the budget absolutely evenly does not take into account the changes in the traffic, the number of competing advertisers, and the value of impressions over the course of time.

adapt to the changes in the competition environment, while moderating computational cost, research considers shortening the successive offline optimization cycles by applying simple algorithms that can myopically re-optimize (Lang et al. 2012) or by tuning the bidding parameters dynamically according to the current performance feedback (Cai et al. 2017).

Substantial progress has been made on optimizing bids in RTB, especially with regards to maximizing expected short-term revenues driven by immediate direct responses such as clicks. Unfortunately, much less attention has been devoted to achieving branding goals, beyond maximizing the number of impressions given the budget constraint. An ideal optimal bidding strategy will maximize (incremental) expected profits over a long period by predicting customers' lifetime value.

Research Opportunities We emphasize two additional research questions. One is the market for data. Advertisers have different types of proprietary data in hand and often consider buying third-party data for additional user information. For example, Amazon has access to users' browsing and purchase histories on its site, whereas P&G might have a harder time linking online users to their offline purchases. Having different levels of user information raises the question of when it will be beneficial for the advertisers to buy additional data and at what price (Bhawalkar et al. 2014; Dalessandro et al. 2014). Related questions of interest pertain to the types of third-party data that can best augment the advertiser's own proprietary data toward improving user targeting.

Another under-attended area of research is the participation and/or transaction cost in operationalizing targeting and bidding in RTB. At the selling channel level, advertisers face entry costs in terms of understanding the complexity in RTB and setting up an appropriate bidding process. At the auction level, there are transaction costs such as monetary fees charged by the intermediaries and/or cognitive effort costs in bidding. Although we are aware of no research quantifying the participation and transaction costs in RTB, these costs will considerably affect advertisers' participation decisions, competitive landscape, and resulting auction outcomes.

3.2.4 Scheduling

We next consider advertisers' ad scheduling decisions; that is, when to display ads and how often. Finding optimal frequency is important not only to attain the highest ad effectiveness (Johnson et al. 2016a) but also to optimally allocate limited resources. Timing is important due to the long-term carry-over or spacing effect of ads (Braun and Moe 2013; Sahni 2015; Sahni et al. 2017).

In the guaranteed selling channel, individual-level scheduling is quite limited and publishers mostly control the ad delivery schedule within the campaign period. Commonly offered decision variables are frequency capping and pacing options (as opposed to frequency/timing at the individual-level).[17] An implementable algorithm that can choose these common decision variables (e.g., how to set frequency capping, does uniform pacing result in better exposure timing) toward achieving optimal reach and overall frequency/timing levels will be valuable.

In RTB, where each impression can be bought separately based on the user-specific information, advertisers have better control over frequency and timing at the individual level, conditioned upon users' browsing behaviors. To this end, one possible direction for future research is building an optimal bidding algorithm that takes the effect of frequency and timing into account (e.g., carry-over effect of ad) in evaluating own valuations of an impression and calculating the bid. Having a greater control over individual-level frequency and timing might be an additional rational for advertisers to allocate budget toward RTB. As winning is probabilistic in RTB and the bid landscape changes rapidly, ad scheduling is an interesting dynamic control problem.

3.3 Measuring Value of Ad Spend

Below we discuss challenges in measuring the value of ad spend (i.e., causal effect) and suggest some directions for future research.

3.3.1 Attribution

Assessing and quantifying the effectiveness of each ad is remarkably difficult, as the consumer is exposed to an advertiser multiple times via various channels, and as these "multi-touches" jointly influence consumer behavior. For example, display and TV ads simultaneously affect online and offline sales. How advertisers attribute the increase in desired outcome to each ad exposure affects how advertisers would update budget allocation across media channels and also the buying decisions in the display ad market (Jordan et al. 2011; Shao and Li 2011; Abhishek et al. 2012; Dalessandro et al. 2012; Berman 2015; Li and Kannan 2014; Xu et al. 2014; Kireyev et al. 2016). With the advancement in tracking technology, richer datasets become available that connect ad exposures not only across communication channels (display, search, TV) but also across devices (desktop,

[17] Frequency capping restricts the number of times a particular campaign is displayed to a user, and pacing is the rate at which the budget is spent during the campaign. Uniform pacing involves spending the budget evenly, whereas no pacing refers to spending as fast as possible.

mobile, tablets). Individual-level online activities are further being matched to offline behaviors, for example via reward cards. Ongoing progress is being made to capture partial aspects of these cross-channel effects (e.g., effect of display ad on online and offline sales in Lewis and Reiley 2014; effect of search and display ads on online and offline purchases in Abraham 2008; effect of display ads on consumer searches in Papadimitriou et al. 2011, Lewis and Nguyen 2015, Ghose and Todri-Adamopoulos 2016; effect of mobile and desktop exposures on mobile and desktop conversions in Ghose et al. 2013) and future research is warranted encompassing the entire customer journey (i.e., effects of multi-channel exposures on multi-channel behaviors) to better understand attribution and the resulting implications on the equilibrium outcomes.

3.3.2 Econometric Issues and Experimental Design

Putting aside the attribution and cross-channel issues, measuring the causal effect in general poses several challenges (see Lewis et al. 2015 and references therein), and observational variation, even with the most sophisticated estimation methods, can fail to yield reliable measurement of ROI (Gordon et al. 2016). To overcome the econometric issues and to better understand the value of ad spend, both advertisers and researchers are turning attention to randomized field experiments. The key question is how to design experiments to obtain clean estimates with statistical power while reducing experimental cost (Kohavi et al. 2009; Barajas et al. 2016; Johnson et al. 2016b, 2017a). Design features include the randomization method (randomize before targeted users are selected or after, sample sizes for the control and treatment groups,[18] type of ads to be shown to the control group [e.g., placebo ad or competing advertiser's ad],[19] length of the experiment(s), and sample data period [i.e., data collection period prior and/or post experiment]).[20]

Using the results from the experiments, advertisers ultimately want to infer the relative effectiveness of different ad buying strategies and optimally re-allocate resources. For example, Hoban and Bucklin 2015 calculate marginal effects and elasticities along the purchase funnel, and suggest ways

[18] Equal-sized A/B testings improve statistical power (Lewis et al. 2015), while dynamically assigning more users toward the better performing choices can potentially reduce experimental costs (Scott 2010; Schwartz et al. 2017). Optimally balancing this trade-off is an important area for future research.

[19] Johnson et al. 2016b show that adding "control ads" to the control group, which can record those users in the control group who would have received the treatment ad had they been in the treatment group, can improve statistical precision. Johnson et al. 2017a develop "ghost ads" methodology that can reduce experimental cost while attaining the precision when running control ads.

[20] Lewis et al. 2015 point out the trade-off in picking a cut-off date for the sample data. While a wider window of time can potentially capture the long-term effect, the amount of noise tends to increase faster than the signal (treatment effect) as one moves further from the campaign date.

to re-allocate impressions across users at different stages. Similarly, Barajas et al. 2016 compare performance-based campaign with CPM-based campaign and leverage the results in re-selecting target users. Presumably advertisers seek to test different ad designs and inventory characteristics, targeting types, targeting algorithms, frequency and timing of ad exposures, and so forth. Based on the result from the previous experiment, how should the next experiment be designed? Related questions are in which order these decision variables need to be tested and which variables should be grouped together in the testing environment to reduce the experimental cost. There exists an increasing demand for automated A/B testing solutions, and more work would be welcome in the context of designing and effectively learning through repeated testings toward finding the global optimal ad buying decision. Repeated feedback cycles (i.e., testing → analyzing results → updating → re-testing) become more crucial when the display ad environment changes rapidly over time and when competitors' ad strategies change in response to focal advertiser's re-optimization.

3.3.3 Advertising Effect and Privacy

More precise targeting generally increases ad effectiveness by delivering ads that are in consumers' interests (Beales 2010; Goldfarb and Tucker 2011b; Farahat and Bailey 2012; Lambrecht and Tucker 2013; Bleier and Eisenbeiss 2015). However, finer targeting requires some degree of privacy intrusion, and may backlash as customers' sense of vulnerability and privacy concerns increase (Goldfarb and Tucker 2011a; Tucker 2014a; Aguirre et al. 2015).[21] More stringent privacy policies (opt-in < opt-out < tracking ban) negatively impact profitability of advertisers and publishers, and there is still an on-going debate over regulatory controls (Johnson 2013; Johnson et al. 2017b). How privacy policies should be set in consideration with consumers' welfare (benefits from increased protection of privacy, minus loss from being exposed to less relevant ads) and changes in welfares of advertisers, publishers, and intermediaries is an interesting area for future research.[22] Another important area for future research concerns how much data to use in targeting and the value of different types of data bought from third parties, when considering the trade-off between finer targeting and privacy concerns. For example, it is unclear whether and when advertisers benefit from acquiring information on users' social graphs and leveraging it in behavioral targeting.

[21] Readers are referred to Tucker 2012b, Goldfarb 2014, and Acquisti et al. 2016 for reviews on economics of privacy and targeted advertising.

[22] On a related note, progress is being made from a computational perspective in developing behavioral targeting systems that are privacy friendly. For example, Toubiana et al. 2010 build a behavioral profiling and targeting system that does not leak user information outside the web browser.

4 Publishers

A publisher on the supply side of the display ad ecosystem faces the following decisions: (i) setting the pricing scheme, (ii) deciding what information to share, and (iii) scheduling the ad delivery. These decisions are contingent upon publishers' inferences about advertisers' valuations.

4.1 Understanding Advertisers' Valuations

Ascertaining the value of impressions to advertisers is a prelude to publishers' ad pricing, sharing information, and ad delivery decisions. It is challenging even for the advertisers to estimate the effect size and understand their own valuations, let alone for the publishers who often do not have access to advertisers' proprietary data (e.g., retail revenues). Understanding advertisers' valuations is closely related to what information the publisher has and may affect optimal pricing scheme.

Estimating advertisers' valuations as part of providing managerial insights for the publishers is relatively scarce. One example is Wu 2015, where advertisers' valuations are estimated by leveraging data obtained from an ad network (containing advertisers' retail revenue referred from ads). His counterfactual analysis reveals an important insight: a publisher would be better off with an auction (fixed) pricing system when the advertiser valuation can (cannot) be well elicited and matched. Eliciting advertisers' valuations based on the information available to publishers is an important area for future research. For example, the bid amount in the auction relative to the fixed price paid in the guaranteed contracts and the volume bought in each selling channel may provide some information about advertisers' valuations.

4.2 Setting Prices

Although the optimal pricing would require joint consideration of guaranteed selling channel and RTB, most existing literature explores pricing decisions conditional on the selling channel.

4.2.1 Guaranteed Selling Channel

In the guaranteed selling channel, publishers and advertisers come to an agreement on a fixed price pertaining to when, where, and how the ads will be displayed. We first discuss the most common CPM pricing, then explore other forms of pricing mechanisms.

CPM Pricing The difficulty in committing to guaranteed contracts in advance is that publishers face uncertain demand (advertisers' arrivals) and ad inventory (viewers' arrivals) (Feige et al. 2008).

To address these uncertainties, existing research typically constructs a revenue management model in which advertisers arrive sequentially with desired quantities of impressions and valuations, and in which the publisher makes contractual decisions prior to knowing other demands arriving in the future. The publisher either (i) determines whether or not to accept arriving advertiser's proposed price and quantity, where the accepted request may be canceled at a cost (Babaioff et al. 2009; Constantin et al. 2009; Roels and Fridgeirsdottir 2009), or (ii) dynamically determines a price to quote to an arriving advertiser's demand (Fridgeirsdottir and Najafi-Asadolahi 2016).

The actual guaranteed contractual process more closely resembles a bargaining process, where the outcome may depend on several factors including advertiser's willingness to pay, publisher's outside option, competition, business relationship, and negotiation skills. All these factors suggest interesting questions for research. As an example, little is known about the role sales people play in selling guaranteed contracts. Sales people might lower the search costs of advertisers in finding the right website, or convey additional information and value about the impressions to the advertisers. With the increase in programmatic direct, the role and value of sales people may change.

Finally, advertisers have different valuations for varying inventory characteristics (e.g., size, format). Publishers can bundle some of these ad units (e.g., desktop+mobile) to better price discriminate, and which ad units to bundle at which fixed bundle price is an open empirical question.

Other Pricing Schemes Although CPM pricing is predominant in display ad, publishers can also consider other pricing schemes such as CPC (Mangani 2004; Fjell 2009; Najafi-Asadolahi and Fridgeirsdottir 2014) or subscription fees (Baye and Morgan 2000; Kumar and Sethi 2009). Kumar and Sethi 2009, for example, develop a dynamic hybrid revenue model that determines the optimal subscription fee level and amount of ad on a website over time, in consideration with the cost of serving each customer, cost of presenting ads, and cost of changing the subscription fee. Another interesting approach is taken in Goldstein et al. 2015, in which a time-based pricing scheme is considered because duration of an ad affects memory retention. Future research is warranted in exploring and comparing different pricing mechanisms that further look at the effect of advertisers' competition, publishers' competition, and consumers' multi-homing behaviors across websites.

4.2.2 Non-Guaranteed Selling Channel

In RTB, advertisers can more readily evaluate the value of each impression for targeting. The impressions become highly differentiated, valuations are heterogenous across advertisers, and the distribution of advertisers' valuations often violates the regularity conditions in Myerson 1981.[23] Moreover, it is likely that advertisers have correlated valuations for certain types of impressions (Abraham et al. 2016). With such characteristics, the Revenue Equivalence Theorem fails to hold and publishers can improve their revenues by designing better mechanisms. Existing literature considers combinations of posted price and auctions (Celis et al. 2014), hybrid auctions where advertisers are allowed to choose between CPM or CPC payment (Zhu and Wilbur 2011), or the incentive problems in CPM, CPC, and other performance-based cost-per-action (CPA) payments (Dellarocas 2012; Hu et al. 2016). In practice, preferred deals are becoming increasingly popular in which impressions are sold at a fixed CPM price (but non-guaranteed) prior to being auctioned. These preferred deals can alleviate adverse selection concerns and improve publisher's revenue (Chen and Wang 2015; Mirrokni and Nazerzadeh 2017). In addition, adding reserve prices can substantially increase publishers' revenue, especially when asymmetric, heterogeneous advertisers compete in thin auctions (Yuan et al. 2014; Balseiro et al. 2015; Paes Leme et al. 2016). When the distribution of bids is not known by the publisher, Amin et al. 2013, Medina and Mohri 2014, and Cesa-Bianchi et al. 2015 explore how to learn optimal reserve prices in a non-parametric fashion.

With the advent of header bidding technology, which gathers bids from multiple ad exchanges simultaneously, competition is increasing among ad exchanges and advertisers participating in the auctions.[24] One interesting avenue for future research is looking at the optimal pricing mechanism with the header bidding technology enabled, and its impact on the market equilibrium and the welfare of advertisers, publishers, and intermediaries.

4.2.3 Interrelation of Guaranteed and Non-Guaranteed Selling Channels

Studies discussed so far take the selling channel as given in studying publishers' decision variables. However, as advertisers can enter in both selling channels and publishers can sell in both, these selling channels need to be optimized jointly (Chen et al. 2014). When publishers are considering

[23] In analyzing the second-price auction data from Microsoft Advertising Exchange, Celis et al. 2014 find that there is often a large gap between the highest and second-highest bid, and that the gap is bigger than the winning payment price on average.
[24] https://www.sovrn.com/blog/header-bidding-grows-up/

the inter-related benefits and costs of selling in both channels, it remains unclear what the optimal reservation prices are or which mechanism yields higher profits for publishers. Other pricing mechanisms might potentially further price discriminate based on advertisers' valuations or uncertainty. For example, a publisher can consider selling call-options that can be transferred among advertisers, who may pay a premium to hedge against fluctuating prices in real-time auctions.

4.3 Information

What information to share with the advertisers and at what price are important questions for the publishers and warrant further studies. In the guaranteed selling channel, the publisher faces a trade-off in providing more information to the advertisers and making more granular level targeting options available. On the one hand, the publisher can facilitate price discrimination as advertisers are willing to pay a premium for an additional layer of targeting. On the other hand, the publisher can choose to disclose less information by bundling impressions together. When the publisher does not know advertisers' heterogenous valuations, bundling reduces information asymmetry. The diminution in this asymmetry can, in turn, lead to higher revenues for the publisher. Additionally, coarser targeting criteria can reduce the risk of not meeting the contractual terms and paying a penalty in the event that the publisher falls short of the impressions.[25] Thus, the publisher accounts for the benefits and costs of providing more information to the advertiser when underwriting the targeting terms in guaranteed contracts.

In RTB, the publisher may benefit from providing more information as advertisers bid higher amounts conditional on participating in the auctions (De Corniere and De Nijs 2016; Hummel and McAfee 2016). At the same time, providing more information decreases the number of participating bidders and creates thin markets, as fewer advertisers are interested in a given impression with highly differentiated attributes (Levin and Milgrom 2010; Fu et al. 2012; Chen and Stallaert 2014). This also results in displaying ads that highly coincide with users' recent behaviors, which may increase privacy concerns and reduce ad responses (Lu and Yang 2015). Furthermore, advertisers will be able to use this shared information in targeting the same user on other, cheaper websites (Ghosh et al. 2015). Because of these trade-offs, a publisher's decision to share (Emek et al. 2014) or to price and sell his/her own proprietary data is more complex than that of a third party selling

[25] For example, some advertisers may value and want to target first-time home buyers. First-time home buyers in their 30s visiting The New York Times will be a smaller set than users in the 30s, and the visits of the former group on a given day will likely exhibit higher volatility.

cookies (Bergemann and Bonatti 2015). This raises questions regarding whether publishers should share information about users' behaviors prior to and/or after the ad exposures, and whether it will benefit publishers to share average CTR or engagement related KPIs for each type of ad inventory.

4.4 Ad Delivery

Ad scheduling is an inherently challenging task for the publisher. On one hand, publishers have guaranteed contracts to fulfill. On another hand, publishers can potentially enjoy high bids from those advertisers who want to cherry pick impressions in RTB. In addition, there is an uncertainty about the supply of impressions which are largely determined by users' browsing behaviors.

For the guaranteed selling channel, researchers looked at allocating arriving impressions among contracted advertisers while maximizing publisher revenue (Feldman et al. 2009; Vee et al. 2010; Devanur et al. 2011; Ciocan and Farias 2012; Turner 2012) and/or ad responses (Korula et al. 2013; Chickering and Heckerman 2003). In recent years, RTB has increased in its importance and size, and researchers have explored dynamic allocation, where publishers first request bids for each impression (often by adopting header bidding) and compare the winning price to the option value of assigning the impression to the best matching guaranteed contract (Ghosh et al. 2009a; Balseiro et al. 2014; Arnosti et al. 2016; Chen 2017; Sayedi 2017). These studies show that dynamic allocation can yield higher profits for the publishers than first sending (random quality) impressions to the guaranteed contracts. Nevertheless, many publishers still prefer selling premium high quality inventories (e.g., front page, leaderboard) via guaranteed contracts, and limit the application of dynamic allocation to lower quality inventories with lower guaranteed contract prices. This is meant to reduce the adverse effect of cherry picking by the advertisers in RTB.

One interesting area for future research is to consider different allocation algorithms for the publishers. For example, publishers may be better off by sending high quality (as opposed to random quality) impressions to guaranteed contracts first, then using dynamic allocation for the rest. Publishers could allow various pacing (Bhalgat et al. 2012; Hojjat et al. 2017) or frequency capping, and assess how these options provided to advertisers affect publisher's optimal ad delivery plan. Another question relates to controlling supply of impressions to facilitate ad delivery or to raise advertiser competition. Publishers could consider buying additional traffic (e.g., via Facebook ad) to intentionally increase supply, or intentionally not sell certain ad spaces at a given point in time. Varying (or smoothing) supply of impressions affects advertiser competition and allocation of impressions among advertisers, making this another potential direction for future research.

5 Intermediaries

As outlined in the display ad market ecosystem discussion, many intermediaries (DSPs, SSPs, ad exchanges, ad networks, data aggregators) serve technological needs in RTB and facilitate the match between advertisers and publishers. These intermediaries affect market outcomes, yet research remains sparse regarding the implications of intermediaries and their decisions.

One characteristic of these intermediaries is that they possess information spanning advertisers and/or publishers and can therefore assess the competitive landscape more completely than a single advertiser or publisher. DSPs have access to ad budgets, bidding strategies, winning bids, and payments across advertisers. SSPs and ad networks have information regarding how impressions are served among advertisers/campaigns and associated ad responses across publishers. Ad exchanges have information about all winning and losing bids from advertisers and inventory characteristics from publishers. Intermediaries may have different data-sharing incentives from advertisers and publishers. Rafieian and Yoganarasimhan 2016 consider the data-sharing arrangements between the ad network and advertisers, finding ad networks have the incentive to withhold information to increase advertiser competition and increase ad networks' revenue. Resolving the information asymmetry between advertisers and publishers by the intermediary will further depend on the competition among intermediaries. For example, header bidding will change the incentives of intermediaries in data sharing arrangements as competition among ad exchanges intensifies and the traditional role of ad networks is undermined.

Second, intermediaries facilitate the match. Balseiro and Candogan 2017 show that intermediation leads to an increase in the overall market efficiency when optimal contracts are offered to budget-constrained advertisers with private information on budgets and targeting criteria. Intermediaries (e.g., SSPs or ad networks) further help publishers to optimize ad delivery by matching impressions to relevant advertisers (Broder et al. 2007; Chakrabarti et al. 2008; Zhang and Katona 2012). On the other hand these intermediaries charge high fees for their services. Although the fees charged by these ad tech intermediaries vary widely depending on the level of services provided in buying and selling ads, their revenues constitute approximately 55% of the ad spend in RTB.[26] The overall welfare implications of intermediaries who enable the RTB are to be examined.[27]

Lastly, economic incentives of principals and intermediaries may not be aligned. Balseiro et al. 2017 explore how the structure of intermediation network affects the profits of participants when advertisers' values are private, and show that intermediaries have incentives to shade bids and not to allocate impressions, even when profitable for their downstream customers. Allouah and Besbes 2017, on the other hand, explore the welfare implications of the collusive behavior of DSPs who represent multiple advertisers, and show that moving from collusion (DSP submitting a single bid to limit competition among the advertisers it represents on a given impression) to non-collusion (DSP bidding for each advertiser independently of other advertisers it represents) leads to a Pareto improvement when taking into account the publisher's reaction such as increasing reserve prices. Future research is called for looking at different objectives of intermediaries, that might differ inherently from those of the principals, and appropriate contract mechanism design.

6 Other Issues: Transparency

There is an increasing demand for transparency in the display ad market, more so with the complexity arising in RTB and with the growth of fraudulent activities (Daswani et al. 2008; Feily et al. 2009; Stone-Gross et al. 2011). In this section, we discuss the industry's effort to adopt better measures (e.g., eligible, viewable, rendered) and research issues related to transparency.

Not all impressions bought are actually viewed by the users, leading to difficulties in predicting ROI and making ad buying decisions at appropriate prices.[28] More advertisers are requesting performance reports with metrics, such as *viewable* impressions, and are demanding pricing schemes to reflect these alternative metrics (e.g., paying only for viewable impressions).[29] In order to provide assurance to advertisers, publishers often use third parties (e.g., eTrust) to verify that they don't serve ads to impressions generated by botnets. Yet, many publishers are wary of meeting the

[26] https://www.iab.com/insights/iab-programmatic-revenue-report-2014-results/

[27] Although in a different industry, Salz 2017 explains how the existence of intermediaries improves overall welfare by reducing incurred search costs and reallocating contracts to lower cost carters in New York City trade-waste market.

[28] The traditional metric, impression, is counted when a creative is sent by the ad server. The counted impression may not be actually viewed by the user, for example, when he/she does not scroll down far enough to see ads loaded at the bottom of a page, or when he/she uses ad blocking tools to purposefully avoid ads. In addition, fake impressions can be generated by fraudulent activities involving botnets or pop-under windows with 0x0 size pixels.

[29] According to the IAB's guideline, an impression is considered viewable when at least 50% of an ad's pixels are in view on the in-focus browser for a minimum of one continuous second (http://mediaratingcouncil.org/081815%20Viewable%20Ad%20Impression%20Guideline_v2.0_Final.pdf).

viewability standards and of the prospect of losing revenues when incorporating the viewability metric into pricing schemes.[30] Understanding the impact of viewability metrics on the profits of advertisers, publishers, and intermediaries is an imminent research question. There exists a stream of research that develops algorithms for detecting fraud impressions and clicks (Stitelman et al. 2013; Crussell et al. 2014; Oentaryo et al. 2014b), and an interesting question is to examine the players' (publishers' or intermediaries') incentives to adopt these algorithms.

Last, there exist concerns that ad targeting/serving algorithms may lead to discriminatory outcomes, intentionally or unintentionally, leading to calls for algorithmic transparency (Datta et al. 2015; Sweeney 2013). For example, Lambrecht and Tucker 2017 find competitive spillover in STEM career ads, where female impressions are highly valued by competing advertisers. This leads to displaying more STEM ads to men by bidding algorithms designed to be cost-effective (and intended to be gender-neutral). Delving further into the underlying mechanisms across various contexts could yield consequential policy insights toward preventing undesirable biases.

7 Conclusion

The display ad market is economically substantial and rapidly growing. The growth is largely driven by RTB, in which many intermediaries exist to serve the technological needs of advertisers and publishers. The guaranteed selling channel and RTB have distinguishing characteristics, especially in terms of players involved, information available, and pricing mechanisms. Despite some industry experts' belief that all ads will be bought and sold via RTB, the dual selling channels will likely co-exist in the long-run as advertisers and publishers have incentives to leverage the distinctive features of both channels (Athey et al. 2017; Sayedi 2017).

In this article, the ecosystem of the display ad market, with dual selling channels, is first outlined. Then for each player involved (advertisers, publishers, and intermediaries), papers are reviewed across disciplines pertaining to each player's decisions. Many interesting inter-disciplinary research questions are open to be explored, especially in joint consideration of guaranteed and non-guaranteed selling channels, and we propose some directions for subsequent studies to stimulate fruitful research in this field.

[30] http://blog.imonomy.com/meeting-the-viewability-threat-to-publishers-revenue/
https://digiday.com/media/publishers-still-bristle-groupms-tougher-viewability-standards/

References

Abhishek, V., P. Fader, and K. Hosanagar. 2012. "Media exposure through the funnel: A model of multi-stage attribution." Working paper.

Abraham, I., S. Athey, M. Babaioff, and M. D. Grubb. 2016. "Peaches, lemons, and cookies: Designing auction markets with dispersed information." Working paper.

Abraham, M. 2008. "The off-line impact of online ads." *Harvard Business Review* 86(4):28.

Acquisti, A., C. Taylor, and L. Wagman. 2016. "The economics of privacy." *Journal of Economic Literature* 54 (2): 442–492.

Agarwal, A., O. Chapelle, M. Dudík, and J. Langford. 2014. "A reliable effective terascale linear learning system." *Journal of Machine Learning Research* 15 (1): 1111–1133.

Agarwal, D., R. Agrawal, R. Khanna, and N. Kota. 2010. "Estimating rates of rare events with multiple hierarchies through scalable log-linear models." In *Proceedings of the 16th international conference on knowledge discovery and data mining,* 213–222. ACM.

Aguirre, E., D. Mahr, D. Grewal, K. de Ruyter, and M. Wetzels. 2015. "Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness." *Journal of Retailing* 91 (1): 34–49.

Ahmed, A., Y. Low, M. Aly, V. Josifovski, and A. J. Smola. 2011. "Scalable distributed inference of dynamic user interests for behavioral targeting." In *Proceedings of the 17th international conference on knowledge discovery and data mining,* 114–122. ACM.

Aksakallı, V. 2012. "Optimizing direct response in internet display advertising." *Electronic Commerce Research and Applications* 11 (3): 229–240.

Allouah, A., and O. Besbes. 2017. "Auctions in the online display advertising chain: A case for independent campaign management." Working paper.

Aly, M., A. Hatch, V. Josifovski, and V. K. Narayanan. 2012. "Web-scale user modeling for targeting." In *Proceedings of the 21st international conference on world wide web,* 3–12. ACM.

Amin, K., A. Rostamizadeh, and U. Syed. 2013. "Learning prices for repeated auctions with strategic buyers." In *Advances in Neural Information Processing Systems,* 1169–1177.

Andrews, M., X. Luo, Z. Fang, and A. Ghose. 2015. "Mobile ad effectiveness: Hyper-contextual targeting with crowdedness." *Marketing Science* 35 (2): 218–233.

Arnosti, N., M. Beck, and P. Milgrom. 2016. "Adverse selection and auction design for internet display advertising." *The American Economic Review* 106 (10): 2852–2866.

Athey, S., E. Calvano, and J. S. Gans. 2014. "The impact of the internet on advertising markets for news media." Working paper.

Athey, S., E. Calvano, and J. S. Gans. 2017. "The impact of consumer multi-homing on advertising markets and media competition." *Management Science, forthcoming.*

Babaioff, M., J. D. Hartline, and R. D. Kleinberg. 2009. "Selling ad campaigns: Online algorithms with cancellations." In *Proceedings of the 10th conference on electronic commerce,* 61–70. ACM.

Bagherjeiran, A., and R. Parekh. 2008. "Combining behavioral and social network data for online advertising." In *International conference on data mining workshops,* 837–846. IEEE.

Bakshy, E., D. Eckles, R. Yan, and I. Rosenn. 2012. "Social influence in social advertising: Evidence from field experiments." In *Proceedings of the 13th conference on electronic eommerce,* 146–161. ACM.

Balseiro, S. R., O. Besbes, and G. Y. Weintraub. 2015. "Repeated auctions with budgets in ad exchanges: Approximations and design." *Management Science* 61 (4): 864–884.

Balseiro, S. R., and O. Candogan. 2017. "Optimal contracts for intermediaries in online advertising." *Operations Research* 65 (4): 878–896.

Balseiro, S. R., J. Feldman, V. Mirrokni, and S. Muthukrishnan. 2014. "Yield optimization of display advertising with ad exchange." *Management Science* 60 (12): 2886–2907.

Balseiro, S., O. Candogan, and H. Gurkan. 2017. "Multi-stage intermediation in display advertising." Working paper.

Balseiro, S., and Y. Gur. 2017. "Learning in repeated auctions with budgets: Regret minimization and equilibrium." Working paper.

Barajas, J., R. Akella, M. Holtan, and A. Flores. 2016. "Experimental designs and estimation for online display advertising attribution in marketplaces." *Marketing Science* 35 (3): 465–483.

Bart, Y., A. T. Stephen, and M. Sarvary. 2014. "Which products are best suited to mobile advertising? A field study of mobile display advertising effects on consumer attitudes and intentions." *Journal of Marketing Research* 51 (3): 270–285.

Baye, M. R., and J. Morgan. 2000. "A simple model of advertising and subscription fees." *Economics Letters* 69 (3): 345–351.

Beales, H. 2010. "The value of behavioral targeting." *Network Advertising Initiative.*
Bergemann, D., and A. Bonatti. 2015. "Selling cookies." *American Economic Journal: Microeconomics* 7 (3): 259–294.
Berman, R. 2015. "Beyond the last touch: Attribution in online advertising." Working paper.
Bhalgat, A., J. Feldman, and V. Mirrokni. 2012. "Online allocation of display ads with smooth delivery." In *Proceedings of the 18th international conference on knowledge discovery and data mining,* 1213–1221. ACM.
Bhawalkar, K., P. Hummel, and S. Vassilvitskii. 2014. "Value of targeting." In *International symposium on algorithmic game theory,* 194–205. Springer.
Bleier, A., and M. Eisenbeiss. 2015. "Personalized online advertising effectiveness: The interplay of what, when, and where." *Marketing Science* 34 (5): 669–688.
Braun, M., and W. W. Moe. 2013. "Online display advertising: Modeling the effects of multiple creatives and individual impression histories." *Marketing science* 32 (5): 753–767.
Broder, A., M. Fontoura, V. Josifovski, and L. Riedel. 2007. "A semantic approach to contextual advertising." In *Proceedings of the 30th international conference on research and development in information retrieval,* 559–566. ACM.
Bruce, N. I., B. Murthi, and R. C. Rao. 2017. "A dynamic model for digital advertising: The effects of creative format, message content, and targeting on engagement." *Journal of Marketing Research* 54 (2): 202–218.
Bucklin, R. E., and P. R. Hoban. 2017. "Marketing models for internet advertising." In *Handbook of Marketing Decision Models,* 431–462. Springer.
Cai, H., K. Ren, W. Zhang, K. Malialis, J. Wang, Y. Yu, and D. Guo. 2017. "Real-time bidding by reinforcement learning in display advertising." In *Proceedings of the 10th international conference on web search and data mining,* 661–670. ACM.
Celis, L. E., G. Lewis, M. Mobius, and H. Nazerzadeh. 2014. "Buy-it-now or take-a-chance: Price discrimination through randomized auctions." *Management Science* 60 (12): 2927–2948.
Cesa-Bianchi, N., C. Gentile, and Y. Mansour. 2015. "Regret minimization for reserve prices in second-price auctions." *IEEE Transactions on Information Theory* 61 (1): 549–564.
Chakrabarti, D., D. Agarwal, and V. Josifovski. 2008. "Contextual advertising by combining relevance with click feedback." In *Proceedings of the 17th international conference on world wide web,* 417–426. ACM.
Chapelle, O. 2015. "Offline evaluation of response prediction in online advertising auctions." In *Proceedings of the 24th international conference on world wide web,* 919–922. ACM.
Chapelle, O., E. Manavoglu, and R. Rosales. 2015. "Simple and scalable response prediction for display advertising." *Transactions on intelligent systems and technology* 5 (4): 61.
Chatterjee, P., D. L. Hoffman, and T. P. Novak. 2003. "Modeling the clickstream: Implications for web-based advertising efforts." *Marketing Science* 22 (4): 520–541.
Chen, B., and J. Wang. 2015. "A lattice framework for pricing display advertisement options with the stochastic volatility underlying model." *Electronic Commerce Research and Applications* 14 (6): 465–479.
Chen, B., S. Yuan, and J. Wang. 2014. "A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising." In *Proceedings of the 8th international workshop on data mining for online advertising,* 1–9. ACM.
Chen, J., and J. Stallaert. 2014. "An economic analysis of online advertising using behavioral targeting." *Mis Quarterly* 38 (2): 429–449.
Chen, Y., D. Pavlov, and J. F. Canny. 2009. "Large-scale behavioral targeting." In *Proceedings of the 15th international conference on knowledge discovery and data mining,* 209–218. ACM.
Chen, Y.-J. 2017. "Optimal dynamic auctions for display advertising." *Operations Research* 65 (4): 897–913.
Chickering, D. M., and D. Heckerman. 2003. "Targeted advertising on the web with inventory management." *Interfaces* 33 (5): 71–77.
Chung, C. Y., J. M. Koran, L.-J. Lin, and H. Yin. 2010. "Model for generating user profiles in a behavioral targeting system." US Patent 7,809,740 (October).
Ciocan, D. F., and V. Farias. 2012. "Model predictive control for dynamic resource allocation." *Mathematics of Operations Research* 37 (3): 501–525.
Constantin, F., J. Feldman, S. Muthukrishnan, and M. Pál. 2009. "An online mechanism for ad slot reservations with cancellations." In *Proceedings of the 20th symposium on discrete algorithms,* 1265–1274. Society for Industrial and Applied Mathematics.

Crussell, J., R. Stevens, and H. Chen. 2014. "Madfraud: Investigating ad fraud in android applications." In *Proceedings of the 12th international conference on mobile systems, applications, and services,* 123–134. ACM.
Cui, Y., R. Zhang, W. Li, and J. Mao. 2011. "Bid landscape forecasting in online ad exchange marketplace." In *Proceedings of the 17th international conference on knowledge discovery and data mining,* 265–273. ACM.
Dalessandro, B., R. Hook, C. Perlich, and F. Provost. 2015. "Evaluating and optimizing online advertising: Forget the click, but there are good proxies." *Big Data* 3 (2): 90–102.
Dalessandro, B., C. Perlich, and T. Raeder. 2014. "Bigger is better, but at what cost? Estimating the economic value of incremental data assets." *Big Data* 2 (2): 87–96.
Dalessandro, B., C. Perlich, O. Stitelman, and F. Provost. 2012. "Causally motivated attribution for online advertising." In *Proceedings of the 6th international workshop on data mining for online advertising and internet economy,* 7. ACM.
Danaher, P. J. 2007. "Modeling page views across multiple websites with an application to internet reach and frequency prediction." *Marketing Science* 26 (3): 422–437.
Danaher, P. J., J. Lee, and L. Kerbache. 2010. "Optimal internet media selection." *Marketing Science* 29 (2): 336–347.
Daswani, N., C. Mysen, V. Rao, S. Weis, K. Gharachorloo, and S. Ghosemajumder. 2008. "Online advertising fraud." *Crimeware: understanding new attacks and defenses* 40 (2): 1–28.
Datta, A., M. C. Tschantz, and A. Datta. 2015. "Automated experiments on ad privacy settings." *Proceedings on privacy enhancing technologies* 2015 (1): 92–112.
De Corniere, A., and R. De Nijs. 2016. "Online advertising and privacy." *The RAND Journal of Economics* 47 (1): 48–72.
Dellarocas, C. 2012. "Double marginalization in performance-based advertising: Implications and solutions." *Management Science* 58 (6): 1178–1195.
Devanur, N. R., K. Jain, B. Sivan, and C. A. Wilkens. 2011. "Near optimal online algorithms and fast approximation algorithms for resource allocation problems." In *Proceedings of the 12th conference on electronic commerce,* 29–38. ACM.
Emek, Y., M. Feldman, I. Gamzu, R. PaesLeme, and M. Tennenholtz. 2014. "Signaling schemes for revenue maximization." *Transactions on Economics and Computation* 2 (2): 5.
Evans, D. S. 2008. "The economics of the online advertising industry." *Review of Network Economics* 7 (3).
Evans, D. S. 2009. "The online advertising industry: Economics, evolution, and privacy." *The Journal of Economic Perspectives* 23 (3): 37–60.
Farahat, A., and M. C. Bailey. 2012. "How effective is targeted advertising?" In *Proceedings of the 21st international conference on world wide web,* 111–120. ACM.
Feige, U., N. Immorlica, V. Mirrokni, and H. Nazerzadeh. 2008. "A combinatorial allocation mechanism with penalties for banner advertising." In *Proceedings of the 17th international conference on world wide web,* 169–178. ACM.
Feily, M., A. Shahrestani, and S. Ramadass. 2009. "A survey of botnet and botnet detection." In *3rd international conference on emerging security information, systems and technologies,* 268–273. IEEE.
Feldman, J., N. Korula, V. Mirrokni, S. Muthukrishnan, and M. Pál. 2009. "Online ad assignment with free disposal." In *International workshop on internet and network economics,* 374–385. Springer.
Fjell, K. 2009. "Online advertising: Pay-per-view versus pay-per-click - A comment." *Journal of Revenue and Pricing Management* 8 (2-3): 200–206.
Fridgeirsdottir, K., and S. Najafi-Asadolahi. 2016. "Cost-per-impression pricing for display advertising." Working paper.
Fu, H., P. Jordan, M. Mahdian, U. Nadav, I. Talgam-Cohen, and S. Vassilvitskii. 2012. "Ad auctions with data." In *Conference on computer communications workshops,* 184–189. IEEE.
Ghose, A., A. Goldfarb, and S. P. Han. 2012. "How is the mobile internet different? Search costs and local activities." *Information Systems Research* 24 (3): 613–631.
Ghose, A., S. Han, and S. Park. 2013. "Analyzing the interdependence between web and mobile advertising: A randomized field experiment." Working paper.
Ghose, A., and V. Todri-Adamopoulos. 2016. "Towards a digital attribution model: Measuring the impact of display advertising on online consumer behavior." *MIS Quarterly* 40 (4): 889–910.

Ghosh, A., M. Mahdian, R. P. McAfee, and S. Vassilvitskii. 2015. "To match or not to match: Economics of cookie matching in online advertising." *Transactions on Economics and Computation* 3 (2): 12.

Ghosh, A., P. McAfee, K. Papineni, and S. Vassilvitskii. 2009a. "Bidding for representative allocations for display advertising." In *International workshop on internet and network economics,* 208–219. Springer.

Ghosh, A., B. Rubinstein, S. Vassilvitskii, and M. Zinkevich. 2009b. "Adaptive bidding for display advertising." In *Proceedings of the 18th international conference on world wide web,* 251–260. ACM.

Goel, S., and D. G. Goldstein. 2013. "Predicting individual behavior with social networks." *Marketing Science* 33 (1): 82–93.

Goldfarb, A. 2014. "What is different about online advertising?" *Review of Industrial Organization* 44 (2): 115–129.

Goldfarb, A., and C. Tucker. 2011a. "Online display advertising: Targeting and obtrusiveness." *Marketing Science* 30 (3): 389–404.

Goldfarb, A., and C. Tucker. 2011b. "Privacy regulation and online advertising." *Management science* 57 (1): 57–71.

Goldfarb, A., and C. Tucker. 2014. "Standardization and the effectiveness of online advertising." *Management Science* 61 (11): 2707–2719.

Goldstein, D. G., R. P. McAfee, and S. Suri. 2015. "Improving the effectiveness of time-based display advertising." *Transactions on Economics and Computation* 3 (2): 7.

Gordon, B. R., F. Zettelmeyer, N. Bhargava, and D. Chapsky. 2016. "A comparison of approaches to advertising measurement: Evidence from big field experiments at Facebook." White paper.

Grewal, D., Y. Bart, M. Spann, and P. P. Zubcsek. 2016. "Mobile advertising: A framework and research agenda." *Journal of Interactive Marketing* 34:3–14.

Ha, L. 2008. "Online advertising research in advertising journals: A review." *Journal of Current Issues & Research in Advertising* 30 (1): 31–48.

He, X., J. Pan, O. Jin, T. Xu, B. Liu, T. Xu, Y. Shi, A. Atallah, R. Herbrich, S. Bowers, et al. 2014. "Practical lessons from predicting clicks on ads at Facebook." In *Proceedings of the 8th international workshop on data mining for online advertising,* 1–9. ACM.

Hoban, P. R., and R. E. Bucklin. 2015. "Effects of internet display advertising in the purchase funnel: Model-based insights from a randomized field experiment." *Journal of Marketing Research* 52 (3): 375–393.

Hojjat, A., J. Turner, S. Cetintas, and J. Yang. 2017. "A unified framework for the scheduling of guaranteed targeted display advertising under reach and frequency requirements." *Operations Research* 65 (2): 289–313.

Hu, Y., J. Shin, and Z. Tang. 2016. "Incentive problems in performance-based online advertising pricing: Cost per click vs. cost per action." *Management Science* 62 (7): 2022–2038.

Hummel, P., and R. P. McAfee. 2015. "Loss functions for predicted click-through rates in auctions for online advertising." *Journal of Applied Econometrics.*

Hummel, P., and R. P. McAfee. 2016. "When does improved targeting increase revenue?" *Transactions on Economics and Computation* 5 (1): 4.

Iyer, K., R. Johari, and M. Sundararajan. 2014. "Mean field equilibria of dynamic auctions with learning." *Management Science* 60 (12): 2949–2970.

Johnson, G. A. 2013. "The impact of privacy policy on the auction market for online display advertising." Working paper.

Johnson, G. A., and R. A. Lewis. 2015. "Cost per incremental action: Efficient pricing of advertising." Working paper.

Johnson, G. A., R. A. Lewis, and E. I. Nubbemeyer. 2017a. "Ghost ads: A revolution in measuring ad effectiveness." *Journal of Marketing Research.*

Johnson, G. A., R. A. Lewis, and D. Reiley. 2016a. "Location, location, location: Repetition and proximity increase advertising effectiveness." Working paper.

Johnson, G. A., R. A. Lewis, and D. H. Reiley. 2016b. "When less is more: Data and power in advertising experiments." *Marketing Science* 36 (1): 43–53.

Johnson, G. A., S. Shriver, and S. Du. 2017b. "Consumer privacy choice in online advertising: Who opts out and at what cost to industry?" Working paper.

Jordan, P., M. Mahdian, S. Vassilvitskii, and E. Vee. 2011. "The multiple attribution problem in pay-per-conversion advertising." In *International symposium on algorithmic game theory,* 31–43. Springer.

Juan, Y., Y. Zhuang, W.-S. Chin, and C.-J. Lin. 2016. "Field-aware factorization machines for CTR prediction." In *Proceedings of the 10th conference on recommender systems,* 43–50. ACM.

Kireyev, P., K. Pauwels, and S. Gupta. 2016. "Do display ads influence search? Attribution and dynamics in online advertising." *International Journal of Research in Marketing* 33 (3): 475–490.

Kohavi, R., R. Longbotham, D. Sommerfield, and R. M. Henne. 2009. "Controlled experiments on the web: Survey and practical guide." *Data mining and knowledge discovery* 18 (1): 140–181.

Korula, N., V. S. Mirrokni, and M. Zadimoghaddam. 2013. "Bicriteria online matching: Maximizing weight and cardinality." In *International conference on web and internet economics,* 305–318. Springer.

Korula, N., V. Mirrokni, and H. Nazerzadeh. 2016. "Optimizing display advertising markets: Challenges and directions." *Internet Computing* 20 (1): 28–35.

Kumar, S., and S. P. Sethi. 2009. "Dynamic pricing and advertising for web content providers." *European Journal of Operational Research* 197 (3): 924–944.

Lambrecht, A., and C. Tucker. 2013. "When does retargeting work? Information specificity in online advertising." *Journal of Marketing Research* 50 (5): 561–576.

Lambrecht, A., and C. E. Tucker. 2017. "Algorithmic bias? An empirical study into apparent gender-based discrimination in the display of STEM career ads." Working paper.

Lang, K. J., B. Moseley, and S. Vassilvitskii. 2012. "Handling forecast errors while bidding for display advertising." In *Proceedings of the 21st international conference on world wide web,* 371–380. ACM.

Lee, D., K. Hosanagar, and H. Nair. 2017. "Advertising content and consumer engagement on social media: Evidence from Facebook." Working paper.

Lee, K.-C., A. Jalali, and A. Dasdan. 2013. "Real time bid optimization with smooth budget delivery in online advertising." In *Proceedings of the 7th international workshop on data mining for online advertising,* 1. ACM.

Levin, J., and P. Milgrom. 2010. "Online advertising: Heterogeneity and conflation in market design." *The American Economic Review* 100 (2): 603–607.

Lewis, R. A., and D. H. Reiley. 2014. "Online ads and offline sales: Measuring the effect of retail advertising via a controlled experiment on Yahoo!" *Quantitative Marketing and Economics* 12 (3): 235–266.

Lewis, R., and D. Nguyen. 2015. "Display advertising's competitive spillovers to consumer search." *Quantitative Marketing and Economics* 13 (2): 93–115.

Lewis, R., J. M. Rao, and D. H. Reiley. 2015. "Measuring the effects of advertising: The digital frontier." In *Economic Analysis of the Digital Economy,* 191–218. University of Chicago Press.

Li, H., and P. Kannan. 2014. "Attributing conversions in a multichannel online marketing environment: An empirical model and a field experiment." *Journal of Marketing Research* 51 (1): 40–56.

Liu, K., and L. Tang. 2011. "Large-scale behavioral targeting with a social twist." In *Proceedings of the 20th international conference on information and knowledge management,* 1815–1824. ACM.

Lu, S., and S. Yang. 2015. "A two-sided market analysis of behaviorally targeted display advertising." Working paper.

Luo, X., M. Andrews, Z. Fang, and C. W. Phang. 2013. "Mobile targeting." *Management Science* 60 (7): 1738–1756.

Manchanda, P., J.-P. Dubé, K. Y. Goh, and P. K. Chintagunta. 2006. "The effect of banner advertising on internet purchasing." *Journal of Marketing Research* 43 (1): 98–108.

Mangani, A. 2004. "Online advertising: Pay-per-view versus pay-per-click." *Journal of Revenue and Pricing Management* 2 (4): 295–302.

McMahan, H. B., G. Holt, D. Sculley, M. Young, D. Ebner, J. Grady, L. Nie, T. Phillips, E. Davydov, D. Golovin, et al. 2013. "Ad click prediction: A view from the trenches." In *Proceedings of the 19th international conference on knowledge discovery and data mining,* 1222–1230. ACM.

Medina, A. M., and M. Mohri. 2014. "Learning theory and algorithms for revenue optimization in second price auctions with reserve." In *Proceedings of the 31st international conference on machine learning,* 262–270.

Menon, A. K., K.-P. Chitrapura, S. Garg, D. Agarwal, and N. Kota. 2011. "Response prediction using collaborative filtering with hierarchies and side-information." In *Proceedings of the 17th international conference on knowledge discovery and data mining,* 141–149. ACM.

Mirrokni, V., and H. Nazerzadeh. 2017. "Deals or no deals: Contract design for online advertising." In *Proceedings of the 26th international conference on world wide web,* 7–14. International World Wide Web Conferences Steering Committee.

Muthukrishnan, S. 2009. "Ad exchanges: Research issues." In *International workshop on internet and network economics,* 1–12. Springer.

Myerson, R. B. 1981. "Optimal auction design." *Mathematics of Operations Research* 6 (1): 58–73.

Najafi-Asadolahi, S., and K. Fridgeirsdottir. 2014. "Cost-per-click pricing for display advertising." *Manufacturing & Service Operations Management* 16 (4): 482–497.

Oentaryo, R. J., E.-P. Lim, J.-W. Low, D. Lo, and M. Finegold. 2014a. "Predicting response in mobile advertising with hierarchical importance-aware factorization machine." In *Proceedings of the 7th international conference on web search and data mining,* 123–132. ACM.

Oentaryo, R. J., E.-P. Lim, M. Finegold, D. Lo, F. Zhu, C. Phua, E.-Y. Cheu, G.-E. Yap, K. Sim, M. N. Nguyen, et al. 2014b. "Detecting click fraud in online advertising: A data mining approach." *Journal of Machine Learning Research* 15 (1): 99–140.

Okazaki, S., A. Katsukura, and M. Nishiyama. 2007. "How mobile advertising works: The role of trust in improving attitudes and recall." *Journal of Advertising Research* 47 (2): 165–178.

Paes Leme, R., M. Pál, and S. Vassilvitskii. 2016. "A field guide to personalized reserve prices." In *Proceedings of the 25th international conference on world wide web,* 1093–1102. International World Wide Web Conferences Steering Committee.

Pandey, S., M. Aly, A. Bagherjeiran, A. Hatch, P. Ciccolo, A. Ratnaparkhi, and M. Zinkevich. 2011. "Learning to target: What works for behavioral targeting." In *Proceedings of the 20th international conference on information and knowledge management,* 1805–1814. ACM.

Papadimitriou, P., H. Garcia-Molina, P. Krishnamurthy, R. A. Lewis, and D. H. Reiley. 2011. "Display advertising impact: Search lift and social influence." In *Proceedings of the 17th international conference on knowledge discovery and data mining,* 1019–1027. ACM.

Perlich, C., B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost. 2012. "Bid optimizing and inventory scoring in targeted online advertising." In *Proceedings of the 18th international conference on knowledge discovery and data mining,* 804–812. ACM.

Perlich, C., B. Dalessandro, T. Raeder, O. Stitelman, and F. Provost. 2014. "Machine learning for targeted display advertising: Transfer learning in action." *Machine learning* 95 (1): 103–127.

Provost, F., B. Dalessandro, R. Hook, X. Zhang, and A. Murray. 2009. "Audience selection for online brand advertising: privacy-friendly social network targeting." In *Proceedings of the 15th international conference on knowledge discovery and data mining,* 707–716. ACM.

Rafieian, O., and H. Yoganarasimhan. 2016. "The value of information in mobile ad targeting." Working paper.

Ren, K., W. Zhang, Y. Rong, H. Zhang, Y. Yu, and J. Wang. 2016. "User response learning for directly optimizing campaign performance in display advertising." In *Proceedings of the 25th international conference on information and knowledge management,* 679–688. ACM.

Roels, G., and K. Fridgeirsdottir. 2009. "Dynamic revenue management for online display advertising." *Journal of Revenue and Pricing Management* 8 (5): 452–466.

Sahni, N. S. 2015. "Effect of temporal spacing between advertising exposures: Evidence from online field experiments." *Quantitative Marketing and Economics* 13 (3): 203–247.

Sahni, N. S., S. Narayanan, and K. Kalyanam. 2017. "An experimental investigation of the effects of retargeted advertising: The role of frequency and timing." Working paper.

Salz, T. 2017. "Intermediation and competition in search markets: An empirical case study." Working paper.

Sayedi, A. 2017. "Real-time bidding in online display advertising." Working paper.

Schwartz, E. M., E. T. Bradlow, and P. S. Fader. 2017. "Customer acquisition via display advertising using multi-armed bandit experiments." *Marketing Science* 36 (4): 500–522.

Scott, S. L. 2010. "A modern Bayesian look at the multi-armed bandit." *Applied Stochastic Models in Business and Industry* 26 (6): 639–658.

Shao, X., and L. Li. 2011. "Data-driven multi-touch attribution models." In *Proceedings of the 17th international conference on knowledge discovery and data mining,* 258–264. ACM.

Stitelman, O., C. Perlich, B. Dalessandro, R. Hook, T. Raeder, and F. Provost. 2013. "Using co-visitation networks for detecting large scale online display advertising exchange fraud." In *Proceedings of the 19th international conference on knowledge discovery and data mining,* 1240–1248. ACM.

Stone-Gross, B., R. Stevens, A. Zarras, R. Kemmerer, C. Kruegel, and G. Vigna. 2011. "Understanding fraudulent activities in online ad exchanges." In *Proceedings of the conference on internet measurement conference,* 279–294. ACM.

Sweeney, L. 2013. “Discrimination in online ad delivery.” *Queue* 11 (3): 10.
Toubiana, V., A. Narayanan, D. Boneh, H. Nissenbaum, and S. Barocas. 2010. “Adnostic: Privacy preserving targeted advertising.” *Proceedings on network and distributed system symposium.*
Trusov, M., L. Ma, and Z. Jamal. 2016. “Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting.” *Marketing Science* 35 (3): 405–426.
Tucker, C. 2012a. “Social advertising.” Working paper.
Tucker, C. 2012b. “The economics of advertising and privacy.” *International Journal of Industrial Organization* 30 (3): 326–329.
Tucker, C. 2014a. “Social networks, personalized advertising, and privacy controls.” *Journal of Marketing Research* 51 (5): 546–562.
Tucker, C. 2014b. “The reach and persuasiveness of viral video ads.” *Marketing Science* 34 (2): 281–296.
Turner, J. 2012. “The planning of guaranteed targeted display advertising.” *Operations Research* 60 (1): 18–33.
Urban, G. L., G. Liberali, E. MacDonald, R. Bordley, and J. R. Hauser. 2013. “Morphing banner advertising.” *Marketing Science* 33 (1): 27–46.
Vee, E., S. Vassilvitskii, and J. Shanmugasundaram. 2010. “Optimal online assignment with forecasts.” In *Proceedings of the 11th conference on electronic commerce,* 109–118. ACM.
Wang, J., and W. Zhang. 2016. “Bid-aware gradient descent for unbiased learning with censored data in display advertising.” In *Proceedings of the 22nd international conference on knowledge discovery and data mining,* 22:665–674. ACM.
Wang, J., W. Zhang, and S. Yuan. 2016. “Display advertising with real-time bidding (RTB) and behavioural targeting.” Working paper.
Wojdynski, B. W., and N. J. Evans. 2016. “Going native: Effects of disclosure position and language on the recognition and evaluation of online native advertising.” *Journal of Advertising* 45 (2): 157–168.
Wu, C. 2015. “Matching value and market design in online advertising networks: An empirical analysis.” *Marketing Science* 34 (6): 906–921.
Wu, W. C.-H., M.-Y. Yeh, and M.-S. Chen. 2015. “Predicting winning price in real time bidding with censored data.” In *Proceedings of the 21th international conference on knowledge discovery and data mining,* 1305–1314. ACM.
Xu, J., K.-c. Lee, W. Li, H. Qi, and Q. Lu. 2015. “Smart pacing for effective online ad campaign optimization.” In *Proceedings of the 21th international conference on knowledge discovery and data mining,* 2217–2226. ACM.
Xu, J., X. Shao, J. Ma, K.-c. Lee, H. Qi, and Q. Lu. 2016. “Lift-based bidding in ad selection.” Working paper.
Xu, L., J. A. Duan, and A. Whinston. 2014. “Path to purchase: A mutually exciting point process model for online advertising and conversion.” *Management Science* 60 (6): 1392–1412.
Yuan, S., J. Wang, B. Chen, P. Mason, and S. Seljan. 2014. “An empirical study of reserve price optimisation in real-time bidding.” In *Proceedings of the 20th international conference on knowledge discovery and data mining,* 1897–1906. ACM.
Yuan, S., J. Wang, and X. Zhao. 2013. “Real-time bidding for online advertising: Measurement and analysis.” In *Proceedings of the 7th international workshop on data mining for online advertising,* 3. ACM.
Zhang, K., and Z. Katona. 2012. “Contextual advertising.” *Marketing Science* 31 (6): 980–994.
Zhang, W., T. Du, and J. Wang. 2016. “Deep learning over multi-field categorical data.” In *European Conference on Information Retrieval,* 45–57. Springer.
Zhang, W., S. Yuan, and J. Wang. 2014. “Optimal real-time bidding for display advertising.” In *Proceedings of the 20th international conference on knowledge discovery and data mining,* 1077–1086. ACM.
Zhu, Y., and K. C. Wilbur. 2011. “Hybrid advertising auctions.” *Marketing Science* 30 (2): 249–273.
Zubcsek, P. P., and M. Sarvary. 2011. “Advertising to a social network.” *Quantitative Marketing and Economics* 9 (1): 71–107.