Revenue-Optimal Dynamic Auctions for Adaptive Ad Sequencing

Omid Rafieian*

Cornell Tech and Cornell University

October 24, 2020

*I am grateful to my advisor Hema Yoganarasimhan, and my committee members Arvind Krishnamurthy, Amin Sayedi, Simha Mummalaneni, and Jacques Lawarrée for their guidance and comments. I am also grateful to an anonymous firm for providing the data and to the UW-Foster High Performance Computing Lab for providing me with computing resources. I thank Caio Waisman for detailed comments that have improved the paper I also thank the participants of the research seminars at University of Wisconsin Madison, University of Colorado Boulder, University of Southern California, University of Texas Dallas at Dallas, Texas AM University, Harvard Business School, Stanford University, Yale University, University of Toronto, Penn State University, University of Rochester, Johns Hopkins University, Rutgers University, Carnegie Mellon University, Cornell Tech, Cornell University, University of California San Diego, and QME 2020 Conference for their feedback. Please address all correspondence to: or83@cornell.edu.

Abstract

Digital publishers often use real-time auctions to allocate their advertising inventory. These auctions are designed with the assumption that advertising exposures within a user's browsing or app-usage session are independent. Rafieian (2019) empirically documents the interdependence in the sequence of ads in mobile in-app advertising, and shows that dynamic sequencing of ads can improve the match between users and ads. In this paper, we examine the revenue gains from adopting a revenue-optimal dynamic auction to sequence ads. We propose a unified framework with two components – (1) a theoretical framework to derive the revenue-optimal dynamic auction that captures both advertisers' strategic bidding and users' ad response and app usage, and (2) an empirical framework that involves the structural estimation of advertisers' click valuations as well as personalized estimation of users' behavior using machine learning techniques. We apply our framework to large-scale data from the leading in-app ad-network of an Asian country. We document significant revenue gains from using the revenue-optimal dynamic auction compared to the revenue-optimal static auction. These gains stem from the improvement in the match between users and ads in the dynamic auction. The revenue-optimal dynamic auction also improves all key market outcomes, such as the total surplus, average advertisers' surplus, and market concentration.

Keywords: online advertising, dynamic mechanism design, ad sequencing, structural models, optimal auctions, reinforcement learning

1 Introduction

Mobile in-app advertising is now a significant source of revenue for publishers and ad networks. In 2018, over 56% of the total digital ad spend came from in-app advertising (eMarketer, 2018). Like other digital advertising environments, mobile publishers use an auction to determine which ad to show inside an app. An auction is a set of rules that characterizes how to allocate each advertising space and how much each advertiser has to pay, given advertisers' bids. Market outcomes under each auction format can accordingly be different since advertisers can strategically vary their bidding behavior. Thus, auction design plays a central role in the success of the digital advertising ecosystem.

Publishers and ad-networks often use auctions that maximize their revenues.[1] The common practice in this industry is to use a first- or second-price auction with an optimally set reserve price. This is in light of the findings from the seminal paper by Myerson (1981) that has shown these auctions are revenue-optimal for a single item, under the regularity assumption. In an advertising environment, these auctions are revenue-optimal if the publisher can treat each advertising space as a single item: i.e., if advertising spaces are independent, and the auction outcome for one advertising space does not create externalities affecting other advertising spaces.

In the context of mobile in-app advertising, Rafieian (2019) provides empirical evidence on the interdependence of exposures within a session where a user is exposed to multiple short-lived ads. He shows that each ad exposure creates externalities that affect future exposures and documents the publisher's gain from dynamic sequencing of ads, i.e., the policy that captures both the immediate and future outcomes in a session and selects the ad that maximizes the expected number of clicks from that point onward. These findings rule out the independence of advertising spaces within a session, which in turn, imply that first- or second-price auction with an optimally set reserve price is not revenue-optimal in the context of mobile in-app advertising.

While the results in Rafieian (2019) elucidate an opportunity to create value in this market by dynamic sequencing of ads by enhancing consumer engagement and match values, the extent to which the publisher can extract this value as revenue is not clear. Notice that advertisers are strategic agents who can change their bids in response to any change in the allocation mechanism, and thereby appropriate most of this created value. The prior literature on advertising dynamics has highlighted the cases where advertisers can strategically make their ad scheduling decisions in competitive environments (Villas-Boas, 1993; Dubé et al., 2005). More specific to online advertising auctions, the prior literature has empirically shown situations where publishers cannot necessarily link the

[1] In this paper, we use the publisher, ad-network, and platform interchangeably, when we refer to the agent who designs the auction and makes the ad allocation decision.

improvement in the match to higher revenues in the presence of advertisers' strategic behavior (Athey and Nekipelov, 2010; Rafieian and Yoganarasimhan, 2020b). Thus, when the revenue is the primary outcome of interest, it is crucial for the publisher to incorporate advertisers' bidding behavior as well as users' ad response when designing the ad sequencing policy.

This brings us to the question of optimal auction design, wherein the publisher designs an auction that maximizes her revenues. Our main goal is to theoretically develop the revenue-optimal dynamic auction and compare its outcomes with the revenue-optimal static auction (second-price auction with optimal reserve price). Overall, we aim to answer the following three research questions in this paper:

1. How can we design a revenue-optimal dynamic auction that captures both inter-temporal trade-offs in ad sequencing and advertisers' strategic bidding behavior?

2. How can we build an empirical framework to evaluate the market outcomes such as publisher's revenues and advertisers' surplus under any auction mechanism?

3. What are the gains from using a revenue-optimal dynamic auction as compared to the static one in mobile in-app advertising? How is the advertisers' surplus distributed across advertisers? Do all advertisers benefit when the publisher uses a dynamic auction?

We need to overcome three major challenges to answer these questions. First, to design a revenue-optimal dynamic auction, we need to specify an allocation rule that incorporates inter-temporal trade-offs in ad interventions and a payment rule that governs advertisers' strategic bidding behavior. Second, to empirically evaluate market outcomes (e.g., publisher's revenues) under counterfactual auctions, we need to obtain accurate estimates of both advertisers' and users' behavior that are valid under any counterfactual auction. For the former, we need to estimate the distribution of advertisers' click valuations as it is the main structural parameter that governs their bidding behavior in any auction. For the latter, we need to estimate users' behavior: their likelihood of clicking on an ad and leaving the session after seeing an ad under any counterfactual allocation policy. Finally, to measure the gains from both dynamic and static revenue-optimal auctions, we need to first solve for the equilibrium outcome under these counterfactual auctions and then use an evaluation method that estimates the corresponding outcomes for each session.

We present an overview of our approach in Figure 1. This figure illustrates the general framework in the top row and the details specific to our problem in the bottom row. In the general framework, we begin with a theoretical framework that informs our empirical approach regarding how to develop and evaluate optimal auctions. More specifically, we start with designing a revenue-optimal dynamic

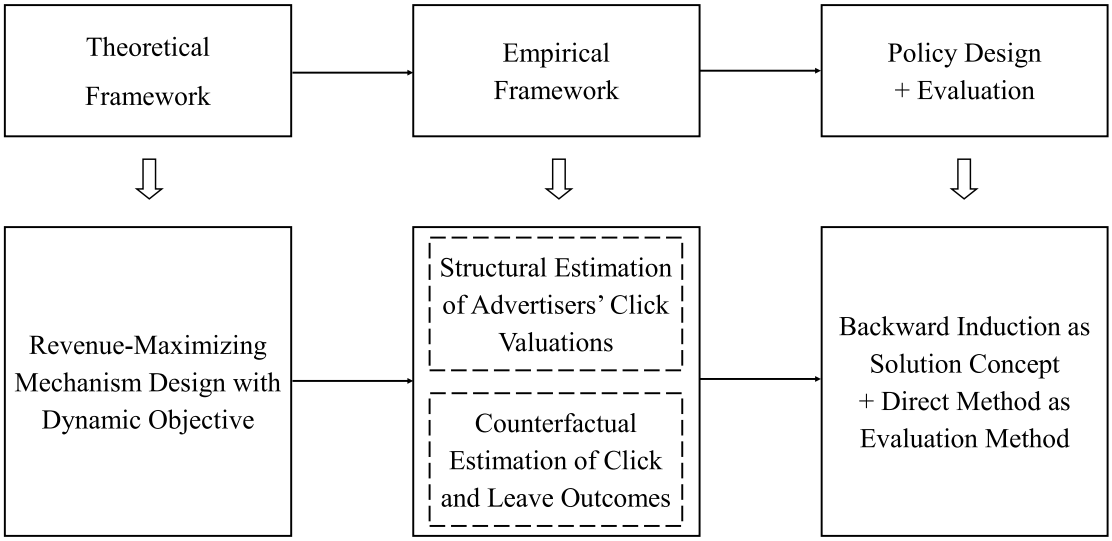

Figure 1: An overview of our approach. The top row presents our general framework and the bottom row shows the specific approach we take in this paper.

auction that gives us a combination of allocation and payment rules. This combination captures both the inter-temporal trade-offs and advertisers' bidding strategies. Since our theoretical framework involves both advertisers' and users' behavior, our empirical framework requires an estimation procedure that cover both these components. We accordingly break our empirical framework into two separate tasks: (1) estimation of the distribution of advertisers' click valuations since click valuation is the key structural parameter governing advertisers' bidding behavior in any auction, and (2) personalized counterfactual estimation of click and leave outcomes, which are the two user-dependent outcomes that affect the expected revenue per session. The second task is the same as the empirical task in Rafieian (2019). We use the same approach in this paper to estimate users' behavior under counterfactual auctions. Finally, we use all our estimates and numerically derive the optimal policy using backward induction and evaluate this policy using the direct method.

Our theoretical framework directly addresses the first challenge and paves the way to address other challenges. We build our theoretical model on the recent literature on dynamic mechanism design that extends the approach in Myerson (1981) to a dynamic setting. The intuitive idea in this literature is to use the Revelation Principle (Myerson, 1981) and exclusively focus on the case where all bidders report their type truthfully (Kakade et al., 2013; Pavan et al., 2014). We use the modeling framework in Kakade et al. (2013) as their separability assumption is particularly suitable in our context: the value an advertiser extracts from an impression is the product of his private click valuation and the expected probability of click on his ad. It allows us to write down the reward function in the Markov Decision Process (MDP henceforth) in terms of virtual valuations and click probabilities (match valuations), thereby deriving the optimal allocation. Using this allocation function, we can then set the payments such that advertisers will participate in the auction and have no incentive to deviate from truthful reporting (IR and IC constraints). We then show that the auction with these allocation and payment rules is revenue-optimal.

To address our second challenge, we propose a structural framework to estimate the distribution of advertisers' click valuation from their observed bids in the data. The key challenge is that the auction format in the data is a quasi-proportional auction, where truthful bidding is not the equilibrium strategy for advertisers. We first characterize the advertisers' utility function in this setting and then derive the equilibrium properties of this auction. Using the first-order condition, we then write the advertisers' click valuation in terms of their cost and the allocation function used by the ad-network. Since both cost and allocation functions can be estimated from the distribution of observed bids and auction configurations, click valuations are identified under the assumption that advertisers are utility-maximizing. This allows us to estimate the distribution of advertisers' click valuations as well as each advertiser's click valuation.

Next, to develop the ad sequencing policy in the revenue-optimal dynamic auction, we need to solve the MDP for the allocation function. While the transition function is the same as the one in Rafieian (2019), the rewards have an additional multiplicative factor – each advertiser's virtual valuation. We can estimate it using the estimated click valuation for each ad as well as the distribution of click valuations. We plug these estimates into the reward function and solve the dynamic allocation policy using a backward induction solution concept. Finally, like Rafieian (2019), we use a direct method approach for evaluation that directly uses our estimates to simulate a session, equilibrium outcomes, and how it evolves. This method allows us to evaluate the revenue outcome for each session.

We first present the results from our auction estimation framework. We theoretically show that advertisers bid roughly half of their click valuations in the quasi-proportional auction. As such, the distribution of bids alone can approximate the distribution of click valuations. Further, it suggests that the current mechanism (quasi-proportional auction) leads to a substantial loss for the platform in terms of both revenue and efficiency. We then focus on the estimated distribution of click valuations from our structural framework, and empirically show that the regularity assumption is satisfied in our context: the virtual valuations are strictly increasing in click valuations. This is an important requirement for our counterfactual analysis, as the solution to the optimal auction is tractable given this assumption.

Next, we conduct our counterfactual analysis to examine the gains from the revenue-optimal dynamic auction. We set the benchmark as the second-price auction with an optimal reserve price, as it is the revenue-optimal static auction. Our results indicate that the expected revenue per session is 1.60% higher under the revenue-optimal dynamic auction compared to that in the revenue-optimal static auction. This is particularly important because most platforms currently use a version of the static revenue-optimal auction. Thus, our results suggest that publishers and ad-networks can significantly benefit from adopting an optimal dynamic auction. Further, we find that the expected number of clicks per session also improve by 1.80% under the dynamic case, suggesting that the gains in revenues can mostly be attributed to the improved match between users and ads as a result of dynamic sequencing, and not to the greater ability of the publisher to extract rent from advertisers.

We then focus on other market outcomes and show that the optimal dynamic auction achieves better outcomes than the optimal static auction in terms of both total surplus and average advertisers' surplus: the total surplus (efficiency) and the average advertisers' surplus increases by 1.77% and 3.00% under the optimal dynamic auction respectively. Hence, we show that the optimal dynamic auction does not achieve revenue optimality at the expense of efficiency. We then explore the surplus gains across advertisers to see whether the market will become more concentrated as a result of using the revenue-optimal dynamic auction. Using a Herfindahl-Hirschman Index (HHI), we find that the optimal dynamic auction has a lower concentration index than the optimal static auction. This suggests that the reason is that the optimal dynamic auction allocates more to the ad with the second largest surplus, thereby closing the gap between the top two advertisers in an auction.

In sum, our paper makes three key contributions to the literature. First, from a methodological point-of-view, we propose a unified dynamic framework that captures both advertisers' and users' behavior to optimize publisher's revenue. A key contribution of our framework is in illustrating how we can use a theoretical framework to break a complex applied problem into a composite of structural estimation and machine learning tasks. To our knowledge, this is the first paper to empirically examine the revenue gains from dynamic sequencing of ads using an optimal dynamic auction. Second, we present a structural estimation framework to recover the distribution of bidders' private valuations from their observed bidding behavior in a quasi-proportional auction. This is the first paper to propose an estimation procedure for quasi-proportional auctions. Our framework can easily be extended to auctions with non-deterministic allocation rules. Third, from a substantive viewpoint, we establish the revenue gains from adopting a dynamic objective in allocating ads, as opposed to a static objective. This is of particular importance, as the current practice in the industry is to use a static objective. We expect our findings to be of relevance to publishers and ad-networks.

2 Related Literature

First, our paper relates to the growing literature on dynamic mechanism design. While early papers in this literature start in 1980s (Baron and Besanko, 1984; Myerson, 1986; Riordan and Sappington, 1987), most of the major developments in this literature appear more recently, with generic characterizations of both efficient mechanisms (Bergemann and Välimäki, 2010; Athey and Segal, 2013) and revenue-maximizing (optimal) mechanisms (Kakade et al., 2013; Pavan et al., 2014). The majority of applied papers on dynamic mechanism design focus on cases where the inter-temporal trade-offs arise through the dynamics of arrival, departure, or population (Vulcano et al., 2002; Parkes and Singh, 2004; Gallien, 2006; Said, 2012). Our paper adds to this literature by empirically evaluating dynamic mechanisms in a digital advertising context. To our knowledge, this is the first paper to provide an empirical framework to examine the performance of dynamic mechanisms.

Second, our paper relates to the literature on the intersection of mechanism design and online advertising. Early papers in this area examines the theoretical properties of different auctions in sponsored search context (Edelman et al., 2007; Varian, 2007; Lahaie et al., 2007). More specific to our context, a series of work takes externalities in search advertising into account and revisits the question of mechanism design in a context where the higher position of an ad may affect the user's decision to even see lower ranked ads (Ghosh and Mahdian, 2008; Kempe and Mahdian, 2008; Ghosh and Sayedi, 2010). In the context of video ads, Kar et al. (2015) adopt a cascade model similar to Kempe and Mahdian (2008), and provide a mechanism for selection and ordering of video ads. While this stream of work proposes simple mechanisms for allocation, they are only applicable to very basic and unrealistic case where the externality is only imposed through the user's leaving decision. We extend this literature by offering a dynamic framework that captures more complex externalities, under a plausible separability assumption.

Lastly, our paper relates to economics and marketing literature on the estimation of auctions. A significant breakthrough in this literature comes from Guerre et al. (2000) who base their identification strategy on the fact that the equilibrium outcome is achieved when all agents maximize their profits given the distribution of others' behavior. While they study the first-price auction with symmetric independent private valuations and without unobserved heterogeneity, other papers in this literature build on this work and extend it to the cases with affiliated private valuations (Li et al., 2002), asymmetric private valuations (Campo et al., 2003), unobserved heterogeneity (Guerre et al., 2009; Krasnokutskaya, 2011), and also different auctions such as scoring auctions (Bajari et al., 2014) and beauty contest auctions (Yoganarasimhan, 2015). Related to our setting, a few papers study online advertising auctions and propose different empirical approaches for the estimation of advertising auctions (Athey and Nekipelov, 2010; Yao and Mela, 2011; Choi and Mela, 2016). Please see Bucklin and Hoban (2017) and Choi et al. (2020) for excellent summaries of models of online advertising in the marketing literature. Our paper adds to this literature by proposing an estimation approach for quasi-proportional auctions that can easily be extended to any randomization-based auction. Further, our counterfactual analysis is the first to consider a dynamic mechanism design in an online advertising context.

3 Optimal Auctions

The main results in Rafieian (2019) establish the gains from the dynamic sequencing of ads in terms of the expected number of clicks generated per session. Intuitively, we expect the increase in the number of clicks to be linked to higher publisher revenues, as clicks are the key revenue-generating source for publishers. However, the fact that advertisers can respond to the change in the allocation by changing their bids in an auction environment makes it unclear how the publisher can extract more revenues from dynamic sequencing of ads. For example, if advertisers decrease their bid as a response to better allocation by dynamic sequencing in the equilibrium, the publisher may end up selling more clicks at a lower price. Thus, it is crucial to take advertisers' bidding behavior into account if the publisher wants to maximize her revenues through adaptive ad sequencing.

To examine revenue gains from adopting a dynamic framework, we incorporate advertisers' utility model into our framework and design revenue-optimal auctions with both dynamic and static objectives. Our framework, in turn, captures both users' behavior and advertisers' strategic bidding. The publisher's problem is then to design an auction with certain allocation and payment rules that maximizes her revenues, i.e., the total payments made by advertisers. To find the revenue-optimal auctions under each objective, we build on the seminal paper by Myerson (1981) and design the allocation and payment rules in the auction to maximize publisher revenues.

This section proceeds as follows: We first define our model environment and assumptions in §3.1. We then introduce the case where the publisher's objective is static and discuss the optimal auction in this case in §3.2. In §3.3, we focus on the optimal auction with a dynamic objective and present the allocation and payments rules in this case.

3.1 Auction Environment and Assumptions

We now describe the auction environment in this problem. For each session i, there are $\mathcal{A}_i$ risk-neutral bidders competing for impressions in this session. Each advertiser a has a private click valuation x_a which is drawn from the distribution F_a with support $[\underline{x}_a, \bar{x}_a]$. This parameter is a private signal that reflects how much the advertiser values a click. We assume that each ad's private valuation is independent of other ads' private valuation and does not vary across impressions.[2] The total value generated from showing ad a at exposure t of session i will then be the product of that ad's click valuation and the probability of click on that ad, which can be written as follows:

$$w_{i,a,t}(x_a; S_{i,t}) = x_a P(Y_{i,t} \mid a, S_{i,t}), \tag{1}$$

where $w_{i,a,t}(x_a; S_{i,t})$ is the total value ad a receives from being shown at exposure t of session i, and $S_{i,t}$ is the state variable at that point which captures the prior history of actions and outcomes within the session, as well as the pre-session information (Please see §4.2.1 in Rafieian (2019) for more details on state variables).

The publisher's problem is to design a mechanism/auction that maximizes her revenues. Each mechanism is characterized by two components – (1) allocation rule, and (2) payment rule. We can define any mechanism as follows:

Definition 1. *A mechanism $M(q, e)$ is defined as a combination of an allocation rule $q(\cdot)$ and a payment rule $e(\cdot)$. Given a profile of reported bids $b = (b_1, b_2, \ldots, b_{\mathcal{A}})$, we can characterize the allocation and payment rules for each ad in the exposure number t in session i as follows:*

- *The allocation rule $q^M_{i,a,t}(b; S_{i,t})$ determines the probability that the item is allocated to a.*
- *The payment rule $e^M_{i,a,t}(b; S_{i,t})$ determines what ad a pays in expectation.*

We assume that the publisher has full commitment power, i.e., the players believe that the publisher follows the rules. Now, under any mechanism $M(q, e)$, we can characterize advertiser's utility function. The only decision variable for any advertiser is to submit a bid that reflects their willingness to pay for a click on their ad. Given advertiser a's bid, we can characterize their utility in exposure t of session i as follows:

$$u^M_{i,a,t}(b_a; x_a, S_{i,t}) = \mathbb{E}_{b_{-a}}\left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b; S_{i,t}) - e^M_{i,a,t}(b; S_{i,t})\right], \tag{2}$$

where b_{-a} is a bid profile of all ads except a, and b is the profile of all bids. We assume that advertisers maximize their utility.

3.2 Warm-Up Case: Optimal Static Auction

We begin by describing the case where the publisher's objective is static. In this case, the publisher only considers the current period rewards. As such, at any point, the goal is to sell the slot to the ad that maximizes the publisher's revenues. The analysis of this case is almost identical to that of the

[2] We can extend our theoretical analysis to the cases where click valuations vary across sessions and exposures. However, in our empirical analysis, we can only identify one value for each ad. This is why we restrict our attention to this case.

seminal paper by Myerson (1981) on optimal auctions. However, we present this case as a warm-up example for our main goal – deriving the optimal dynamic auction.

In general, the choice of the optimal auction may seem impossible as there is no bound on the set of feasible auctions. However, in light of the Revelation Principle and without loss of generality, we can only focus on direct-revelation mechanisms wherein advertisers truthfully bid their click valuations (Myerson, 1981). A direct revelation mechanism is feasible if it satisfies:

1. *Plausibility:* For any profile of reported click valuations x, we have $\sum_{a\in\mathcal{A}_i} q_{i,a,t}(x; S_{i,t}) \leq 1$ and $q_{i,a,t}(x; S_{i,t}) \geq 1$ for all $a \in \mathcal{A}_i$. This condition guarantees that each impression is allocated to at most one ad.

2. *Individual Rationality (IR):* Given reported click valuations, all advertisers receive non-negative utility, i.e., $u^M_{i,a,t}(x_a; x_a, S_{i,t}) \geq 0$ for all $a \in \mathcal{A}_i$. This condition guarantees that all competing ads for a session will participate in the auction.

3. *Incentive Compatibility (IC):* No advertiser has incentive to deviate from bidding truthfully, given that everyone else bids truthfully. Therefore, we have:

$$u^M_{i,a,t}(x_a; x_a, S_{i,t}) \geq u^M_{i,a,t}(b_a; x_a, S_{i,t}) \tag{3}$$

   for any $b_a \in [\underline{x}_a, \bar{x}_a]$. This condition guarantees that reporting truthfully is a Bayesian Nash Equilibrium for all advertisers.

The Revelation Principle helps us reduce the set of all auction to the set of feasible direct revelation mechanism denoted by $\mathcal{M}_{\text{direct}}$. As such, our search is over a more structured set with clear constraints. We can write the publisher's optimization problem as follows:

$$\max_{M\in\mathcal{M}_{\text{direct}}} \mathbb{E}_x \left[\sum_{a\in\mathcal{A}_i} e^M_{i,a,t}(x; S_{i,t}) \right] \tag{4}$$

In this optimization, $M \in \mathcal{M}_{\text{direct}}$ implies that the mechanism must satisfy all three constraints presented above. While focusing only on feasible direct revelation mechanisms helps constrain the problem, we still need further transformations to find the optimal solution. One key transformation is to use envelope condition instead of the IC constraint. The following lemma shows the link between these two:

Lemma 1. *If the mechanism $M(q, e)$ is IC, we have:*

$$
\begin{aligned}
u^M_{i,a,t}(x_a; x_a, S_{i,t}) = \; & u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) \\
& + P(Y_{i,t} \mid a, S_{i,t}) \int_{x'_a}^{x_a} \mathbb{E}_{x_{-a}} \left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) \right] db_a
\end{aligned}
\tag{5}
$$

Now, we can use this lemma to derive the publisher's expected revenue under any IC mechanism as follows:

Lemma 2. *If the mechanism $M(q, e)$ is IC, the publisher's expected revenue can be written as follows:*

$$
\mathbb{E}_x \left[\sum_{a \in \mathcal{A}_i} \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) \right] - \sum_{a \in \mathcal{A}_i} u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) \tag{6}
$$

The result of Lemma 2 is transformative in finding the optimal auction, as for any given environment, the expected revenue of an auction only depends on the allocation at the equilibrium and advertisers' expected utility of their lowest click valuation. We can now use Equation (6) as the objective function and maximize it subject to the *Plausibility* and *Individual Rationality (IR)* constraints and obtain the optimal auction.

Another important feature of writing the objective function in Equation (6) is that it additively separates allocation and payment functions: the first component is independent of the payment function. Thus, one candidate for the optimal auction is to find a plausible allocation function that maximizes the first component and a payment function that minimizes the second component. This brings us to the following proposition:

Proposition 1. *The mechanism $M(q, e)$ is optimal if q maximizes*

$$
\mathbb{E}_x \left[\sum_{a \in \mathcal{A}_i} \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) \right] \tag{7}
$$

subject to q being plausible and $\mathbb{E}_{x_{-a}} \left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t}) \right]$ increasing in x_a, and the payment function e is

$$
e^M_{i,a,t}(x; S_{i,t}) = w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) - P(Y_{i,t} \mid a, S_{i,t}) \int_{\underline{x}_a}^{x_a} q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) db_a \tag{8}
$$

As shown in Proposition 1, we can solve for the optimal allocation and payment functions. In next sections, we discuss the details of each component.

3.2.1 Allocation Rule

Proposition 1 offers a constrained optimization to find the allocation function under the static case. Without the constraint on $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ being increasing in x_a, we can simply maximize the objective function by allocating to the ad with the highest $\left(x_a - \frac{1-F_a(x_a)}{f_a(x_a)}\right) P(Y_{i,t} \mid a, S_{i,t})$. However, there is no guarantee whether the constraint is satisfied unless we impose the following assumption on the distribution F_a for all ads:

Assumption 1. *Distribution F_a is regular, i.e., the function $c_a(x_a) = x_a - \frac{1-F_a(x_a)}{f_a(x_a)}$ is strictly increasing in x_a.*

It is worth noting that this is not an unrealistic assumption, since most familiar distributions satisfy this condition. Under this assumption, the resulting allocation rule will allocate the item to ad with the highest non-negative $c_a(x_a)P(Y_{i,t} \mid a, S_{i,t})$ for any exposure t in session i. If all values are negative, the publisher does not sell the item. It is easy to show that under this allocation, the constraint on $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ being increasing in x_a is satisfied: an increase in x_a will increase the expected probability of winning the item for ad a, since $c_a(x_a)$ is increasing in x_a.

3.2.2 Payments

Following Equation (8) in the second part of Proposition 1, we can now determine the optimal payment functions, consistent with the allocation function presented in §3.2.1. We can show that in this case, losing ads will not pay anything: the first component in Equation (8) is zero, and it is easy to show that the integral is also zero. The payment for the winning ad, however, can be calculated as follows:

$$
\begin{aligned}
e^M_{i,a,t}(x; S_{i,t}) &= w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) - P(Y_{i,t} \mid a, S_{i,t}) \int_{\underline{x}_a}^{x_a} q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) db_a \\
&= w_{i,a,t}(x_a; S_{i,t}) - P(Y_{i,t} \mid a, S_{i,t}) \int_{x_a}^{x_a} db_a \\
&= x_a P(Y_{i,t} \mid a, S_{i,t}),
\end{aligned}
\tag{9}
$$

where x_a is the minimum bid that still wins the impression for the winning ad. The simple allocation rule in the static case helps us find the analytical solution to the integral in Equation (8). With this payment rule, it is easy to check the incentive compatibility of the proposed optimal mechanism.

Finally, if we consider the case of symmetric click valuations, we can simplify the optimal auction to a greater extent. In this case, instead of having ad-specific distributions F_a for each ad, we have one distribution F from which all click valuations are drawn independently. We can show the following corollary for this specific case:

Corollary 1. *With symmetric independent click valuations, the optimal auction is a second-price auction with a reserve price of* $c^{-1}(0)$*. The auction allocates the item to the ad with the highest* $w_{i,a,t}(x_a; S_{i,t})$ *and that ad pays the second-highest* $w_{i,a,t}(x_a; S_{i,t})$ *in expectation.*

Thus, the optimal auction either allocates the impression to the highest valuation advertiser or does not allocate it at all, meaning that it never allocates an impression to an advertiser who does not have the highest valuation for that impression.

3.3 Optimal Dynamic Auction

We now focus on the optimal auction with the dynamic objective. In this case, the publisher wants to incorporate the expected future revenues as well as expected revenues from the current period. This is in contrast with optimal static auction where the publisher only cares about her expected revenues in the current period. The publisher's goal in this dynamic environment is to design a mechanism $M(q, e)$ that maximizes her expected revenues from the session, i.e., $\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(\sum_{a \in \mathcal{A}_i} e^M_{i,a,t}(b; S_{i,t})\right)\right]$.

This change in the publisher's objective clearly changes critical aspects of our environment. The most important change is that the publisher no more sells each exposure t in session i, but rather auctions off the entire session i. As such, the optimal auction in the static case is no more optimal under the current objective. An important factor that helps us simplify the dynamic problem is that only one piece is private to each advertiser in the entire session: their click valuation x_a. That is, advertisers' click valuation will not change within the session, and anything that changes their overall valuation of each impression (e.g., the probability of click in different states) is common knowledge. Thus, we can treat this problem as a case of static auction where advertisers just submit only one bid and the publisher decides how to allocate the exposures within the session, using a dynamic objective.

In line with the change in the publisher's objective, we must re-define advertisers' utility function for the session as follows:

$$U^M_{i,a}(b_a; x_a, S_i) = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b; S_{i,t}) - e^M_{i,a,t}(b; S_{i,t})\right)\right], \tag{10}$$

where the expectation is taken over bidding strategies and the mechanism, and S_i denotes the pre-session information.

Similar to the static case, without loss of generality, we only focus on dynamic direct revelation mechanisms to find the optimal auction (Myerson, 1986). As such, for the dynamic case, we re-write the requirements for a feasible direct revelation mechanism as follows:

1. *Plausibility:* This condition states that at each time period, the exposure is allocated to at most one ad, given that advertisers report their click valuations truthfully. That is, we have $\sum_{a \in \mathcal{A}} q_{i,a,t}(x; S_{i,t}) \leq 1$ and $q_{i,a,t}(x; S_{i,t}) \geq 1$ for all $a \in \mathcal{A}$ and for any t in session i.

2. *Individual Rationality (IR):* Given truthful reporting of click valuations, the advertiser's expected utility over the session is non-negative, i.e., $U^M_{i,a}(x_a; x_a, S_i) \geq 0$ for all $a \in \mathcal{A}_i$.

3. *Incentive Compatibility (IC):* No advertiser has incentive to deviate from bidding truthfully, given that everyone else reports their bid truthfully. Hence, we can write:

$$U^M_{i,a}(x_a; x_a, S_i) \geq U^M_{i,a}(x'_a; x_a, S_i) \tag{11}$$

   where the expectation is taken over other advertisers' click valuations and the mechanism. This constraint guarantees that truth-telling is a Bayesian Nash Equilibrium for all advertisers.

While this restriction to the direct revelation mechanisms reduces the set of auctions the publisher considers, we still need some transformations in these constraints to be able to solve for the optimal mechanism. Like the static case, we show the resulting envelope condition from the IC constraint in the dynamic case as follows:

Lemma 3. *If the mechanism $M(q, e)$ is IC, we have:*

$$\begin{aligned} U^M_{i,a}(x_a; x_a, S_i) = \; & U^M_{i,a}(x'_a; x'_a, S_i) \\ & + \int_{x'_a}^{x_a} \mathbb{E}_{x_{-a}} \left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) \right] db_a \end{aligned} \tag{12}$$

Now, with the dynamic version of the envelope condition, we can show that the publisher's expected revenues under any IC mechanism can be written as follows:

Lemma 4. *If the mechanism $M(q, e)$ is IC, the publisher's expected revenue can be written as follows:*

$$\begin{aligned} \mathbb{E}_x \Bigg[& \sum_{t=1}^{\infty} \beta^{t-1} \left(\sum_{a \in \mathcal{A}_i} \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)} \right) P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x_a; S_{i,t}) \right) \\ & - \sum_{a \in \mathcal{A}_i} U^M_{i,a}(\underline{x}_a; \underline{x}_a, S_i) \Bigg] \end{aligned} \tag{13}$$

Lemma 4 is the equivalent of Lemma 2 for the dynamic case. Now, if we optimize this new objective subject to both *Plausibility* and *Individual Rationality*, we can find the optimal auction.

Further, an important finding of this lemma is that the publisher's revenues cannot exceed the first component in Equation (13). Thus, roughly speaking, if we find a dynamic allocation policy that maximizes the first component and payments are such that the second component is zero, the corresponding mechanism is optimal. More precisely, we can write the following proposition on the optimal auction with dynamic objective as follows:

Proposition 2. *The mechanism $M(q, e)$ is optimal if q maximizes*

$$\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a;S_{i,t})\right)\right] \tag{14}$$

subject to q being plausible and $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ increasing in x_a, and the payment function e is

$$\begin{aligned} e^M_{i,a}(x;S_i) =&\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right)\right] \\ &-\int_{\underline{x}_a}^{x_a}\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(b_a,x_{-a};S_{i,t})\right]db_a, \end{aligned} \tag{15}$$

where the expectation is taken over the stochasticity induced by the dynamic process and state transitions, and not over other advertisers' click valuations.

This proposition states that if an allocation mechanism that maximizes the first part of Equation (13) and sets the payment to make the second part zero, this mechanism is optimal if the allocation mechanism satisfies both plausibility and monotonicity conditions. We explain the details of both allocation and payment components in next sections.

3.3.1 Allocation Rule

We start by finding the optimal allocation rule. Again, we impose the assumption that the distribution F_a is regular for each ad a, i.e., $c_a(x_a) = x_a - \frac{1-F_a(x_a)}{f_a(x_a)}$ is increasing in x_a. In the static case, we show that under this assumption, the optimal ad at any time period is simply the one that maximizes $c_a(x_a)P(Y_{i,t}\mid a,S_{i,t})$. In the dynamic case, however, we want to design a dynamic allocation policy that maximizes the expected revenues for the entire session.

An important result of Lemma 4 is that we can re-write the reward function that is independent of payments. We present a generic definition of reward function with revenue-maximizing objective as follows:

$$R^r_t(a;S_{i,t}) = \left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)P(Y_{i,t}\mid a,S_{i,t}) \tag{16}$$

Now, we can use the reward function in Equation (16) to write down the publisher's optimization problem and find the optimal allocation. The following lemma characterizes the optimal allocation function:

Lemma 5. *If the distribution F_a is regular for each ad a (i.e., $c_a(x_a) = x_a - \frac{1-F_a(x_a)}{f_a(x_a)}$ is increasing in x_a), then the optimal allocation is the solution to the following Markov Decision Process:*

$$\underset{a}{\operatorname{argmax}}\, R_t^r(a; S_{i,t}) + \beta \mathbb{E}_{G_{i,t+1}|S_{i,t},a}\left[V_{t+1}^r(G_{i,t+1})\right], \tag{17}$$

where the value function is defined for each exposure number as follows:

$$V_{t+1}^r(S_{i,t}) = \max_a R_t^r(a; S_{i,t}) + \beta \mathbb{E}_{G_{i,t+1}|S_{i,t},a}\left[V_{t+1}^r(G_{i,t+1})\right] \tag{18}$$

Depending on the context of our problem, we can use various approaches to design the optimal dynamic allocation policy. Based on this dynamic policy, we can then easily define the allocation function $q_{i,a,t}(x, S_{i,t})$. This lemma guarantees that under the assumption that click valuations come from a regular distribution, the chosen q based on Equation (17) and Equation (18) satisfies both plausibility and monotonicity constraints.

3.3.2 Payments

The intuition behind the payment function in the dynamic case is the same as that in the static case. Each advertiser pays the expected valuation they received from the session, minus an informational rent which is determined by integration of their allocation over the set of lower possible values. This payment function guarantees that the IR constraint is satisfied and the second component in Equation (13) will be zero.

While the informational rent has an analytical solution in the static case as shown in Equation (9), it is harder to derive it analytically in the dynamic case, since the allocation function contains more elaborate rules. The first component in Equation (15) is the expected valuation advertiser a receives from the session, and the second component is the informational rent. To calculate the amount of this rent, we basically need to move down from the true click valuation and see how the number of impressions allocated to advertiser a shrinks. Integrating over this function over the possible values will then give us the amount of rent advertiser a is able to extract. Intuitively, it is the total leeway advertiser a has in this optimal auction.

4 Empirical Strategy

In light of the theoretical results in Proposition 1 and 2, we can characterize the reward function at any point for both static and dynamic cases as follows:

$$R_t^r(a; S_{i,t}) = \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)}\right) P(Y_{i,t} \mid a, S_{i,t}), \tag{19}$$

where $S_{i,t}$ is the set of state variables in session i at exposure t. Since the reward specification is the same at any point, we can present a unifying specification of the publisher's optimization in both static and dynamic cases as follows:

$$\underset{a}{\operatorname{argmax}}\, R_t^r(a; S_{i,t}) + \beta \mathbb{E}_{S_{i,t+1}|S_{i,t},a} \left[V_{t+1}^r(S_{i,t+1})\right], \tag{20}$$

where $\beta = 0$ when the objective is static and the value function is defined as the maximum future rewards at a given state.

Given the unified optimization problem in Equation (20), we need to know three key elements to determine the outcomes: (1) distribution of click valuations, (2) match values or click probabilities, and (3) distribution of transitions. The first two elements are required for both static and dynamic objectives, whereas the distribution of transitions is only required for the dynamic case. To conduct an empirical analysis of both static and dynamic optimal auctions presented in §3, we need to obtain estimates for all three unknown elements in Equation (20). Thus, we can identify three empirical tasks as follows:

Task 1: For any ad a in the data, we want to estimate the click valuations x_a based on their observed bidding behavior in the data, under the quasi-proportional auction.

Task 2: For any set of state variables observed in the data, we want to accurately estimate the click probability for all ads if shown in that impression. That is:

$$\hat{y}_{i,t}(a; S_{i,t}) = P(Y_{i,t} \mid a, S_{i,t}), \forall a \tag{21}$$

Task 3: For any set of state variables observed in the data, we want to accurately estimate the leave probability for all ads if shown in that impression. That is:

$$\hat{l}_{i,t}(a; S_{i,t}) = P(L_{i,t} \mid a, S_{i,t}), \forall a \tag{22}$$

The empirical strategy for Tasks 2 and 3 is presented in Rafieian (2019) in great details. We use the same approach for these tasks. We present our empirical strategy to estimate the click valuations in this section.

4.1 Setting

The setting of this problem is the same as that of Rafieian (2019). However, since we are interested in estimation of click valuations, we present more details on the auction environment and advertisers' decisions in this section.

4.1.1 Auction Mechanism

For any exposure that is recognized, the ad-network runs an auction to serve an ad. Unlike the common practice in this industry, the ad-network runs a quasi-proportional auction to select the ad for each exposure. The most notable feature of this auction is in the probabilistic allocation rule that is in contrast with the commonly used mechanisms such as first- or second-price auctions.

1. **Reserve Price:** There is a reserve price b_0 that advertisers' bid must exceed to participate in the auction.

2. **Allocation Function:** For any exposure i and any set of participating ads $\mathcal{A}_i$ with bidding profile $(b_1, b_2, \ldots, b_{|\mathcal{A}_i|})$, ad a has the following probability to win each exposure t within session i:
$$q^p_{i,a,t}(b; z) = \frac{b_a z_a}{\sum_{j \in \mathcal{A}_i} b_j z_j}, \tag{23}$$
where z_a is ad a's quality score which is a measure reflecting the profitability of ad a. The rationale behind using such quality score adjustments is to run the auction based on the expected revenue the ad-network can extract from the ad, rather than their willingness to pay per click. For example, there may be an ad with a very high bid but no chance of getting a click. So despite its high willingness to pay per click, the ad-network cannot actually extract much from it as it will not get many clicks.

   While quality scores can technically vary across auctions, the ad-network does not update them regularly. They only take value zero when the ad is not available for the auction due to budget exhaustion or targeting decisions. Likewise, bids do not change across t, since advertisers cannot choose their bids per auction[3]. Further, changing bids is unlikely as bids reflect advertisers' structural parameters that are stable across auctions.

3. **Payment Scheme:** The ad-network employs a cost-per-click (CPC henceforth) payment scheme wherein advertisers are only charged when a user clicks on their ad. The amount an

[3]There are over 500 auctions run in a second and the information shared with advertisers is minimal about them, disabling them from distinguishing auctions.

ad is charged per click is determined by a next-price rule similar to that of Google's sponsored search auctions. That is, ads are first ranked based on their product of bid and quality score, and each ad pays the minimum amount that guarantees their rank, if a click happens on their ad:

$$e^p_{i,a,t}(b;z) = \inf\left\{ b' \mid \sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b'z_a \leq b_j z_j) = \sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b_a z_a \leq b_j z_j) \right\}, \qquad (24)$$

where $\sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b_a q_a \leq b_j q_j)$ indicates the rank of advertiser a, and the payment b' is the minimum amount of bid that guarantees the same rank for ad a. For example, if there are three bidders with bids 1, 2, and 3, and quality scores 0.1, 0.2, and 0.3, the scores will be 0.1, 0.4, and 0.9. Now, if the second-ranked bidder gets a click, she will pay the price that would have guaranteed her second rank. That is, she only needs to pay $\frac{1\times 0.1}{0.2} = 0.5$, as it guarantees her score to be higher than the third-ranked bidder. Since the item being sold is a click, advertisers' bids reflect their willingness to pay per click. Our goal is to use advertisers' observed bid to estimate their click valuations.

4.1.2 Advertisers' Decisions

Advertisers can make the following decisions:

- **Bid:** Advertisers can set their bid indicating how much they are willing to pay per click.
- **Targeting:** Advertisers can specify their targeting decisions on the following variables: (1) app category, (2) province, (3) hour of the day, (4) smartphone brand, (5) connectivity type, and (6) mobile service provider (MSP). As such, they can exclude some categories from the variables listed above. It will guarantee that their ad will not be shown in the excluded categories.
- **Budget:** Advertisers become unavailable if they do not have enough budget in their account. They can refill their budget and make their campaigns available again.
- **Design of the Banner:** They can design a small banner for their ad.

While we mainly focus on their bidding behavior, it is important to take into account what other decisions they can make. Further, it is important to notice what information they receive from the auction. Each advertiser has a profile in which she can track some key performance metrics on an hourly basis such as the number of impressions, number of clicks, and the total cost of clicks. As such, they do not have granular access to exposure-level information and can only incorporate this information at an aggregate level.

4.2 Data

We use data from a leading in-app ad-network in a large Asian country for over a one week time period from October 22 to 30 in 2015. The analysis sample is the same as the one in Rafieian

(2019). However, for estimation of auction, we use all the impressions as it adds to the precision of our estimates. The original data are at the impression-level, indicating the characteristics of each impression, the ad shown, and the final clicking decision by the user. Please see § 3.2.1 in Rafieian (2019) for the description of impression-level data.

The data required for the estimation of auction are not readily available, but we can construct that from the impression-level data. Auction information in the impression-level data is limited to the winning ad and the bid and potential CPC for that particular ad. However, for estimation of auction, we need more information on each auction (impression) regarding the losing ads as well as their bids. We use the filtering strategy presented in § 5.1.1 in Rafieian (2019).

The idea is simple: ad a is available for session i if these two conditions hold: 1) ad a has enough budget at the time of session i, and 2) ad a has not excluded any targeting categories in auction i. The former can easily be verified by looking at time-stamp for impressions ad a has won and inferring unavailability by identifying discontinuities. The latter can also be verified by inferring ad a's targeting decisions from impressions ad a has won around the time of session i: if ad a's targeting decisions match with targeting variables in session i and this ad is available around the time of this session, it satisfies both requirements and participates in this auction. A few noteworthy factors in the data help in drawing accurate inferences. First, in a data-rich environment like this, identification of discontinuities and targeting decisions is not challenging. Second, advertisers do not frequently change their targeting decisions, making the inference easier on them.

Overall, it provides us with 547,626,756 impressions and 398 distinct ads competing to win the impressions.

4.2.1 Summary Statistics

Here we present some summary statistics of the data. Overall, there are 398 ads participating in the timeline of the study. Table 1 shows mean, standard deviation, minimum and maximum of key ad-level variables. All these variables are defined by ads and reflect an important aspect of advertiser's decision, including bidding, targeting, and budget.

Variable	**Mean**	**Std. Dev**	**Min**	**Max**
Bid	411.13	208.25	300.00	2976.00
Avg. CPC	363.30	103.98	300.00	1375.84
Total Hours of Availability	48	70	1	217
No. of Impressions	1,379,411	5,402,346	7	66,228,977
No. of Clicks	12,619	55,208	0	656,680
Click-Through Rate	0.0109	0.0112	0.0000	0.1429
Total Expenditure	4,880,013	21,032,815	0	216,434,237

Table 1: Summary statistics of key ad-level variables

As shown in this table, there is substantial variation in ads participating in this auction, ranging from smaller ads with a very short lifetime and few clicks, to larger ads with permanent availability and significant expenditure. Further, Table 1 presents number of categories targeted out of all 86 targeting categories, indicating that most ads are not targeting at all.

Figure 2 shows the the empirical CDF of ads' bids and their average CPC. The figure on the left (Figure 2b) shows the distribution for all values of bids. As shown in this figure, there is a reserve price 300 that censors the left side of the graph. With no reserve price, there could have been bids lower than 300. This raises an important identification challenge that we address in the estimation approach, especially because a substantial portion of ads are reserve bidders.

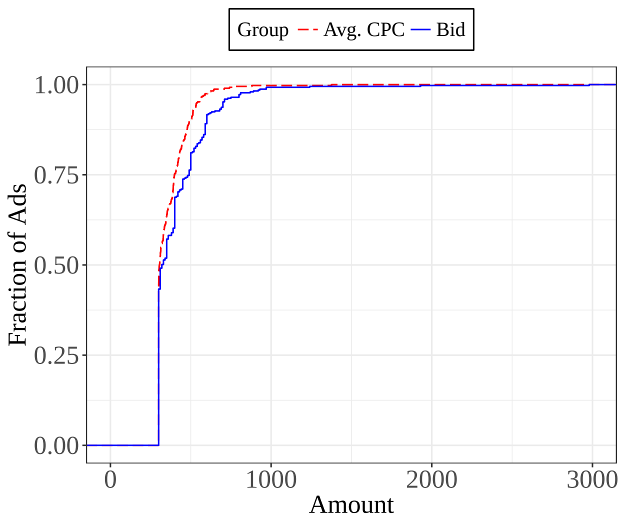

(a) Zoomed Out (b) Zoomed In

Figure 2: Empirical CDF of ads' bids and average CPCs.

Figure 2b zooms into the bids for a shorter interval of amounts that capture over 90% of all ads. While there are discontinuities in bids especially at round numbers, average CPCs show a smoother pattern. Overall, both figures demonstrate a first-order stochastic dominance relationship between the empirical CDF of average CPC and bids which emphasizes the fact an ad's CPC cannot exceed her bid.

4.3 Estimation of Auction

We now present our approach to estimate the distribution of click valuations from the observed auction data. We first develop advertisers' utility model and provide and equilibrium analysis in §4.3.1. We then state the set of assumptions required for the identification of the distribution of click valuations in §4.3.2. Finally, in §4.3.3, we propose our estimation approach.

4.3.1 Advertisers' Model

We begin by characterizing advertisers' utility model in the context of our data. Given the allocation and payment functions defined in §4.1.1, we know that advertiser a receives the following utility from exposure t in session i:

$$\mathbb{1}(A_{i,t} = a, Y_{i,t} = 1)\left(x_a - e^p_{i,a,t}(b)\right),$$

where the indicator function takes value one if ad a is selected through the allocation mechanism and got clicked in that impression. Here we explicitly assume that advertisers only receive utility from clicks and there is no utility from an impression alone.[4]

Mirrokni et al. (2010) study this auction with the case of identity cost function, i.e., $e^p_{i,a,t}(b) = b_a$. They show that advertisers' optimal bidding strategy depends on their expected probability of winning. While this expected probability varies across auctions in our case, we do not observe bid-changing behavior. This is possibly because the effect of the expected probability of winning on their bid is very small in a competitive market and it does not exceed the bid-changing cost.[5]

The observation that advertisers do not change their bids informs us in modeling advertisers' utility function. We can characterize the expected utility of advertiser a from quasi-proportional auction p as follows:

$$u^p_a(b_a; x_a) = m_a\left(x_a \tilde{q}^p_a(b_a) - \tilde{e}^p_a(b_a)\right), \tag{25}$$

where m_a is the expected probability of click on ad a conditional on winning. Since we treat this probability as independent to other parts of the expected utility, we separate that in the specification. Further, functions $\tilde{q}^p_a$ and $\tilde{e}^p_a$ are advertiser a's expected allocation and payment functions. Before defining these two functions, we define an important distribution that we use in our estimation. Let $\mathcal{C}_a$ denote the joint distribution of auctions that advertiser a participates in. Each draw C_a from this distribution contains the information about the distribution of bids and quality scores, the impression characteristics (session i and exposure t), and the number of ads competing in that auction. Using this distribution, we can now characterize the allocation function as follows:

$$\tilde{q}^p_a(b_a) = \mathbb{E}_{\mathcal{C}_a}\left[q^p_{i,a,t}(b_a, b^{C_a}_{-a})\right], \tag{26}$$

[4]We believe this assumption is valid in our context, because almost all ads are mobile apps whose objective is more app installs. This assumption may be violated in the presence of brand ads whose objective is more reach.

[5]Please notice that this is just a behavioral cost and the platform does not charge them for bid-changing. As shown in Rafieian and Yoganarasimhan (2020b), the marginal effect of the expected winning probability α is $\frac{1}{1-\alpha}$. Hence, in a market with many homogeneous competitors, α has a very small effect on advertisers' equilibrium bidding strategy.

where the expectation is taken over all the configurations C_a. This function essentially returns the expected probability of winning by advertiser a in a random impression. Similarly, we can define the payment function as follows:

$$\tilde{e}_a^p(b_a) = \mathbb{E}_{\mathcal{C}_a}\left[\mathbb{1}\left(A_{i,t} = a; C_a\right) e_{i,a,t}^p\left(b_a, b_{-a}^{C_a}; z_a, z_{-a}^{C_a}\right)\right] \tag{27}$$

The term $\mathbb{1}\left(A_{i,t} = a; C_a\right)$ in Equation (27) shows whether ad a wins the impressions and the second term basically computes the cost-per-click using Equation (24). Now, using the utility specification in Equation (25), we can write the first-order condition as follows:

$$x_a = \frac{\partial \tilde{e}_a^p(b_a)}{\partial b_a}\left(\frac{\partial \tilde{q}_a^p(b_a)}{\partial b_a}\right)^{-1} \tag{28}$$

This equation lays out our estimation approach, as we need to empirically estimate both $\tilde{q}$ and $\tilde{e}$ using the data at hand. In the next section, we discuss the assumptions required for identification and overall identification strategy.

4.3.2 Assumptions and Identification

We make a series of assumptions required for our estimation task. Some of these are commonly used in the context of auction estimation, whereas some other assumptions are more specific to the context of quasi-proportional auction in our data. While the former is necessary for identification, we impose the latter mostly for the ease of estimation. For robustness check, we relax those specific assumptions and show the results will not change.

Our first assumption characterizes how advertisers make decisions regarding their own bidding strategy:

Assumption 2. ***[Profit-Maximizing Advertisers]*** *Advertisers are profit-maximizing, i.e., they choose the bid that maximizes their profit.*

In light of Assumption 2, we can treat observed bids in the data as equilibrium bids and use Lemma 6 to estimate click valuations. As shown in Equation (25), advertisers choose their optimal bids given their own private click valuations and their belief about other advertisers. The following assumption characterizes other advertisers' private click valuations:

Assumption 3. ***[Independent Private Values (IPV)]*** *Advertisers' private click valuations are drawn independently from a distribution $F(\cdot)$ with a continuous density.*

The assumption of Independent Private Values (IPV) is an assumption used in most auction settings (Guerre et al., 2000; Athey and Haile, 2007). This assumption hints a straightforward approach to estimate both $\tilde{q}$ and $\tilde{e}$ in Equation (25), by simulating the distribution of $\mathcal{C}_a$ for each ad.

Now, we make assumptions more specific to the case of quasi-proportional auctions. Some of these assumptions are necessary to use the first-order condition in Lemma 6. However, we impose some of these assumptions to provide an analytically simpler solution. We start with the following assumption:

Assumption 4. ***[Zero Impression Value]*** *Advertisers' valuation from an unclicked impression is zero.*

This assumption is widely used in both theoretical and empirical literature on cost-per-click auctions (Edelman et al., 2007; Varian, 2007; Athey and Nekipelov, 2010). In our case, we believe the value from impressions is negligible as most ads are mobile apps whose objective is to get more app installs and can be considered as performance ads. In the next assumption, we make an assumption about the role of budget in advertisers bidding behavior:

Assumption 5. ***[Budget Independent Bidding]*** *Advertisers' bidding behavior is independent of their budget.*

In principle the equilibrium bidding behavior may change for budget-constrained advertisers (Borgs et al., 2007; Asadpour et al., 2014). We make this assumption in our case for two reasons. First, we observe quite a few advertisers who have run out of budget during our study and refilled it later. However, there is no difference in their bidding behavior. The second reason is more empirical, as we do not observe the exact budget in our data. As such, we need to approximate this budget which can be quite noisy.

The assumptions so far are enough to establish the identification of the distribution of click valuations for bidders who bid above the reserve price. However, we make two additional assumption that helps us derive a simpler analytical solution for this auction.

Assumption 6. ***[Separability of Allocation]*** *We can separate the allocation function from the payment function as follows:*

$$\tilde{e}^p_a(b_a) = \tilde{q}^p_a(b_a)\mathbb{E}_{\mathcal{C}_a}\left[e^p_{i,a,t}\left(b_a, b^{C_a}_{-a}; z_a, z^{C_a}_{-a}\right)\right]$$

Assumption 7. ***[Proportional Functional Form of Allocation]*** *We can write the functional form for the function $\tilde{q}$ as follows:*

$$\tilde{q}^p_a(b_a) = \frac{b_a z_a}{b_a z_a + \zeta_a}, \tag{29}$$

where ζ_a captures the overall competition ad a faces.[6]

These assumptions allow us to simplify the first-order condition presented in Equation (28). However, it is worth noting that we can actually relax both assumptions and neither of them is crucial for identification.

For brevity, let $\epsilon_a^p(b_a) = \mathbb{E}_{\mathcal{C}_a}\left[e_{i,a,t}^p\left(b_a, b_{-a}^{C_a}; z_a, z_{-a}^{C_a}\right)\right]$. It is easy to show that the function ϵ_a^p is monotonic for each advertiser a. However, we need to make the following assumption about this function:

Assumption 8. ***[Twice Differentiability of Payment]*** *The expenditure function ϵ_a^p is twice differentiable in b_a.*

Now, we can use all these assumptions and show the following lemma which is the key idea behind our identification:

Lemma 6. *If advertiser a's equilibrium bid is greater than the reserve price ($b_a > b_0$) and the function $b^2 \frac{\partial \epsilon_a^p(b)}{\partial b}$ is increasing in the local neighbourhood around b_a, her click valuation x_a can be written in terms of equilibrium bids as follows:*

$$x_a = \epsilon_a^p(b_a) + \frac{b_a \frac{\partial \epsilon_a^p(b)}{\partial b}\,|_{b=b_a}}{1 - \tilde{q}_a(b_a)} \tag{30}$$

Lemma 6 shows the first-order condition for the case where the advertiser's equilibrium bid is greater than the reserve price. We need $b^2 \frac{\partial \epsilon_a^p(b)}{\partial b}$ to be increasing at b_a to satisfy the second-order condition. Intuitively, if the function e is too concave, this assumption may not hold. However, it is a testable assumption as we can empirically test it for all the bidders.

Now, we focus on the advertisers whose bid is equal to the reserve price. For these advertisers, we cannot use the inverse bidding equation in Equation (30), as their optimal bid could have been lower than the reserve price, if they had been allowed to bid lower. Instead of the first-order condition, we have other conditions for reserve price bidders. First, we know that their participation constraint is satisfied, i.e., $x_a \geq b_0$. Second, we know that their first-derivative at b_0 is lower than or equal to zero, as the utility must be decreasing at the truncation point. Together, we can write the following lemma to characterize the link between the click valuation and reserve bidding behavior:

[6]It is worth mentioning that the functional form of the expected probability of winning is not necessarily quasi-proportional, but it is well-approximated by a quasi-proportional functional form.

Lemma 7. *If advertiser a's equilibrium bid is equal the reserve price ($b_a = b_0$), we can obtain lower and upper bounds for her click valuation x_a as follows:*

$$b_0 \leq x_a \leq \epsilon_a^p(b_0) + \frac{b_0 \frac{\partial \epsilon_a^p(b)}{\partial b} \mid_{b=b_0}}{1 - \tilde{q}_a(b_0)} \tag{31}$$

In light of Lemma 31, we cannot point-identify click valuations for the reserve bidders. However, we can use lower and upper bounds in Equation (31). To complete our identification for the distribution of valuations for all participating advertisers, we need one more assumption on the click valuations of the reserve bidders:

Assumption 9. ***[Uniformity of Reserve Bidders' Valuations]*** *The click valuation x_a for any reserve price bidder is drawn from a uniform distribution with the following bounds:*

$$x_a \sim \mathcal{U}\left(b_0, \epsilon_a^p(b_0) + \frac{b_0 \frac{\partial \epsilon_a^p(b)}{\partial b} \mid_{b=b_0}}{1 - \tilde{q}_a(b_0)}\right) \tag{32}$$

With this assumption, we can now state our identification proposition:

Proposition 3. *If all Assumptions 2 to 9 hold, then the distribution of advertisers' private click valuations are non-parametrically identified.*

The proof for this part is similar to the identification proof for most auction models with independent private values (Guerre et al., 2000; Athey and Haile, 2007). We can directly observe all the elements in Lemma 6 and 7. Thus, we can estimate advertisers' click valuations and form the distribution F.

4.3.3 Estimation Method

Our estimation approach relies on the findings in §4.3.2. We first estimate the joint distribution of configurations $\mathcal{C}_a$ for all ads. This involves the estimation of quality scores, distribution of observed bids, and observed impressions. We then use this distribution to form both allocation and payment functions, which in turn, allows us to estimate the distribution of click valuations.

Before describing the estimation procedure, we need to define a time period η at which advertisers update their bids. While advertisers do not change their bid in our data, we set this time period for two main reasons. First, advertisers can technically change their bids if they want. Therefore, it is more reasonable to assume that they revise their decision every once in a while. Second, from an empirical point-of-view, it allows us to capture the variance in the set of advertisers competing. Thus, following the arguments in Rafieian and Yoganarasimhan (2020a), we set an hourly time period such that each η distinguish a different hour-day combination. The denote the last time period by L.

We provide a step-by-step procedure for our estimation as follows:

- **Step 1:** We estimate all the quality scores $z_{a,\eta}$ for all ads and time periods. We use the proportional nature of the allocation to identify quality scores. If ads a and a' both participate in the auction for exposure t in session i, the denominator of their winning probability is the same. Thus, the odds ratio of these two ads can be written as follows:

$$\frac{\Pr(A_{i,t} = a)}{\Pr(A_{i,t} = a')} = \frac{b_a z_a}{b_{a'} z_{a'}}$$

  Since we observe their bids, we can easily estimate the ratio $\frac{z_a}{z_{a'}}$ by calculating the number of impressions awarded to each ad over the set of auctions where they both have participated. Let $\mathcal{I}^{\eta}_{a,a'}$ denote the set of impressions wherein both ads a and a' participate in time period η. We can write:

$$\frac{\hat{z}_{a,\eta}}{\hat{z}_{a',\eta}} = \frac{b_{a'} \sum_{(i,t)\in\mathcal{I}^{\eta}_{a,a'}} \mathbb{1}(A_{i,t} = a)}{b_a \sum_{(i,t)\in\mathcal{I}^{\eta}_{a,a'}} \mathbb{1}(A_{i,t} = a')} \tag{33}$$

  As such, we can estimate all the ratios. For ad a that has participated in all the impressions, we set $\hat{z}_{a,\eta} = 1$ for all time periods. This allows us to estimate all quality scores $\hat{z}_{a,\eta}$.
- **Step 2:** We empirically estimate the joint distribution $\mathcal{C}_{a,\eta}$ for all ads in all time periods. This distribution contains the information about all the impressions that ad a has participated in, number of bidders in corresponding impressions, and the joint distribution of bids and quality scores. We call the estimated distribution $\hat{\mathcal{C}}_{a,\eta}$.
- **Step 3:** Given all estimated configurations $\hat{\mathcal{C}}_{a,\eta}$, we can estimate the allocation probability for all ads over all time periods as follows:

$$\hat{q}_{a,\eta}(b_a) = \frac{1}{N_a} \sum_{(i,t)\sim\hat{\mathcal{C}}_{a,\eta}}^{N_a} \mathbb{1}(A_{i,t} = a), \tag{34}$$

  where N_a is the number of draws we get from the distribution $\hat{\mathcal{C}}_{a,\eta}$. It is worth noting that we do not need to fully estimate the allocation function as the relationship in Equation (28) only depends on the final allocation probability given the observed bid.
- **Step 4:** Again, we use the distribution estimate configurations $\hat{\mathcal{C}}_{a,\eta}$ to estimate our cost-per-click

function $\hat{\epsilon}_{a,\eta}(b)$. For any value b, we can estimate the cost-per-click function as follows:

$$\hat{\epsilon}_{a,\eta}(b) = \frac{1}{N_a} \sum_{(i,t)\sim\hat{\mathcal{C}}_{a,\eta}}^{N_a} \inf \left\{ b' \mid \sum_{j\in\mathcal{A}_i, j\neq a} \mathbb{1}(b'z_a \leq b_j z_j) - \mathbb{1}(bz_a \leq b_j z_j) = 0 \right\} \tag{35}$$

This gives us the estimate for the cost-per-click function. We can then take the numerical derivatives of $\hat{\epsilon}_{a,\eta}(b)$ and estimate $\hat{\epsilon'}_{a,\eta}(b)$ for all values of b.

- **Step 5:** Using the estimates for the allocation and cost-per-click functions, we can now identify click valuations $x_{a,\eta}$ for any combination of a and η as follows:

$$\hat{x}_{a,\eta} = \begin{cases} \hat{\epsilon}_{a,\eta}(b_a) + \frac{b_a \hat{\epsilon'}_{a,\eta}(b_a)}{1-\hat{q}_{a,\eta}(b_a)} & b_a > b_0 \\ x_0 \sim \mathcal{U}\left(b_0, \hat{\epsilon}_{a,\eta}(b_a) + \frac{b_a \hat{\epsilon'}_{a,\eta}(b_a)}{1-\hat{q}_{a,\eta}(b_a)}\right) & b_a = b_0 \end{cases} \tag{36}$$

  We also obtain single estimates for each x_a without using the time periods. The procedure is the same as what described above. We only need to use $\hat{\mathcal{C}}_a$ instead of this distribution for all time periods. We then use these estimates throughout for click valuations in our analysis.

- **Step 6:** Finally, we use the estimates in Equation (36) to form the distribution of click valuations. We can write:

$$\hat{F}(x) = \frac{1}{L} \sum_{\eta=1}^{L} \frac{1}{|\mathcal{A}_\eta|} \sum_{a\in\mathcal{A}_\eta} \mathbb{1}(x_{a,\eta} \leq x) \tag{37}$$

$$\hat{f}(x) = \frac{1}{L} \sum_{\eta=1}^{L} \frac{1}{|\mathcal{A}_\eta|} \sum_{a\in\mathcal{A}_\eta} \frac{1}{h} \mathcal{K}\left(\frac{x - \hat{x}_{a,\eta}}{h}\right), \tag{38}$$

  where $\mathcal{K}$ is the kernel function. In this case, we use Epanechnikov kernel, i.e., $\mathcal{K}(u) = \frac{3}{4}(1-u^2)\mathbb{1}(|u| \leq 1)$.

4.4 Results

We now present our results on the estimation of click valuations. We first show the main findings on the distribution of click valuations in §4.4.1. Next, in §4.4.2, we discuss the validity of our assumption that the distribution of click valuations is regular.

4.4.1 Estimated Distribution of Click Valuations

We first show the estimated distribution of click valuations for top ads. We only focus on the top ads, as it is the set of ads that we use for the counterfactual evaluation of the optimal auctions. However, we can technically recover the distribution for any set of ads. Figure 3 show the empirical CDF

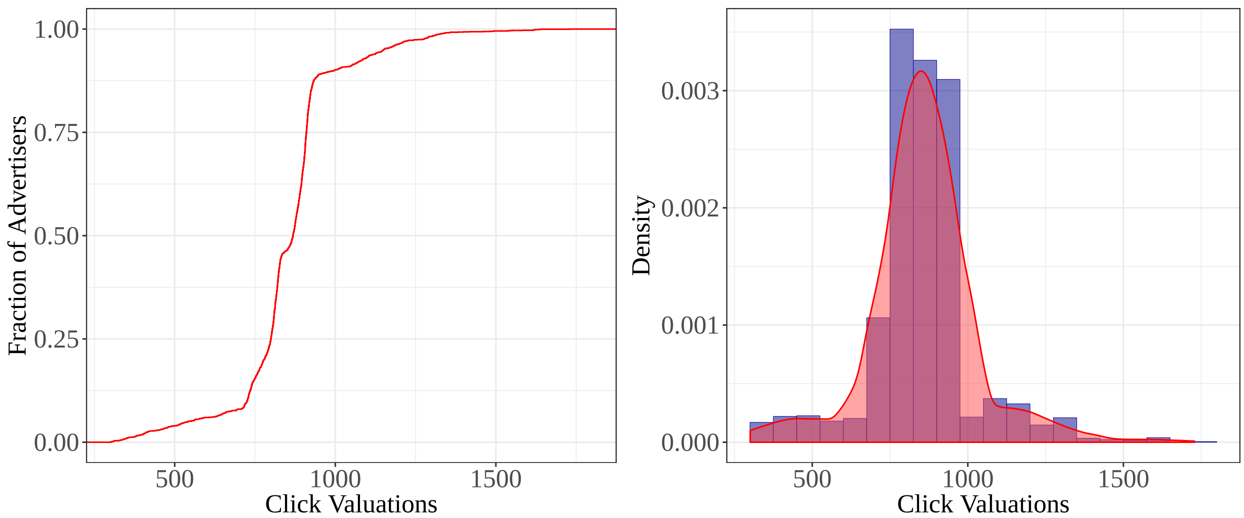

(a) CDF

(b) Density

Figure 3: Estimated distribution of click valuations for top 15 ads.

and estimated density of the distribution. As shown in Figure 3a, the range between 700 to 1000 constitutes the vast majority of click valuations. Over 80% of click valuations lie in this range.

Figure 3b shows the estimated density for the distribution of click valuations. We use Epanechnikov kernel with bandwidth size of 75. The shape of density is similar to a bell curve with a low variance. The mean and standard deviation for this distribution are 854.91 and 164.73 respectively.

4.4.2 On the Regularity of the Distribution of Click Valuations

Now, we use the results shown in Figure 3 and check whether the estimated distribution is regular. As described in Assumption 1, the estimated distribution $\hat{F}$ is regular if $x - \frac{1-\hat{F}(x)}{\hat{f}(x)}$ is strictly increasing in x. Using our estimates for $\hat{F}$ and $\hat{f}$, we plot the virtual valuations in Figure 4. As demonstrated in this figure, virtual values are monotonic in click valuations.

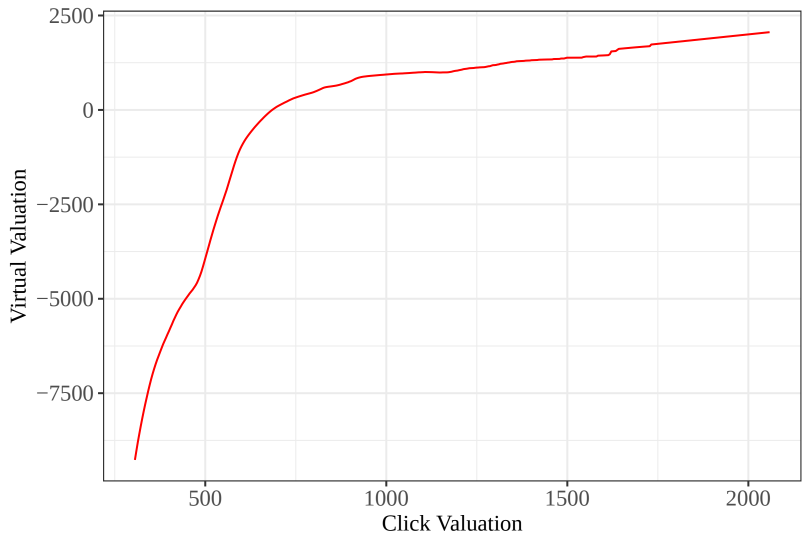

Figure 4: Virtual valuations against click valuations

5 Counterfactual Evaluation of Optimal Auctions

In this section, we take our theoretical models in §3.2 and §3.3 to the data and present our approach to evaluate the counterfactual outcomes under optimal auctions. In both cases, we need to use the reward function that is presented Equation (16) in a generic way. More specific to our context, we can write it as follows:

$$R_t^r(a; S_{i,t}) = \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)}\right) P(Y_{i,t} \mid a, S_{i,t}), \tag{39}$$

Now, we want to estimate the reward function in this case, given our data. We use estimates for both click probability at any state (i.e., $P(Y_{i,t} \mid a, S_{i,t})$) as well as click valuations x_a for all ads from Rafieian (2019). We also impose symmetry assumption on the distribution of click valuations and estimate $\hat{F}$ and $\hat{f}$, based on the estimated click valuations. Given these estimates, we can now estimate the reward function in Equation (40) and re-write it as follows:

$$\hat{R}_t^r(a; S_{i,t}) = \left(\hat{x}_a - \frac{1 - \hat{F}(\hat{x}_a)}{\hat{f}(\hat{x}_a)}\right) \hat{y}_{i,t}(a; S_{i,t}) \tag{40}$$

We use our estimates for the reward function in this case to obtain optimal auctions in both static and dynamic cases and then present how we evaluate the resulting auctions.

5.1 Solution Concept for the Optimal Static Auction

We start with the simpler case where the publisher's objective is static. In §3.2.1 and §3.2.2, we show how we can derive the optimal mechanism $M^m(q^m, e^m)$. More specifically, we show that when $\hat{F}$ is symmetric, the optimal auction will be a second-price auction with an optimal reserve price. We can estimate the reserve price $\hat{e}_0$ as follows:

$$\hat{e}_0 = \hat{c}^{-1}(0),$$

where $\hat{c}(x) = \frac{1-\hat{F}(x)}{\hat{f}(x)}$ and we can find $\hat{e}_0$ by solving for $\hat{e}_0 = \frac{1-\hat{F}(\hat{e}_0)}{\hat{f}(\hat{e}_0)}$. Using this reserve price, we can then derive the optimal allocation q^m as follows:

$$q^m_{i,a,t}(\hat{x}; S_{i,t}) = \begin{cases} 1 & \hat{R}^r_t(a; S_{i,t}) > \max_{a' \in \mathcal{A}_i \setminus a} \left(\hat{R}^r_t(a'; S_{i,t}), \hat{e}_0\right) \\ 0 & \text{otherwise} \end{cases} \tag{41}$$

This allocation rule indicates that the highest reward ad (which is the ad with the highest expected valuation) will win the impression.[7] As specified in §3.2.2, the payments are determined as follows:

$$e^m_{i,a,t}(\hat{x}; S_{i,t}) = \begin{cases} \max_{a' \in \mathcal{A}_i \setminus a} \left(\hat{R}^r_t(a'; S_{i,t}), \hat{e}_0\right) & q^m_{i,a,t}(x \mid S_{i,t}) = 1 \\ 0 & \text{otherwise} \end{cases} \tag{42}$$

We can now use both Equation (41) and Equation (42) to find the optimal mechanism and apply it to each session.[8]

5.2 Solution Concept for the Optimal Dynamic Auction

Now, we focus on the optimal auction with dynamic objective and show how we can use our estimates and determine the optimal mechanism $M^d(q^d, e^d)$. As shown in §3.3.1, we can find the optimal allocation by solving a Markov Decision Process that incorporates expected future rewards as well as the current period reward. Like Rafieian (2019), we consider a finite case with and solve for the optimal allocation using backward induction. For notational convenience, we first define the

[7]It is worth noting that in a case where there are multiple highest reward ads, we randomly allocate the item to one of the highest reward ads. However, since this a very rare event in an empirical setting, for brevity, we exclude that from Equation (41).

[8]It is important to notice that we can recover the distribution of click valuations specific to each session. For simplicity, we focus on the case one global distribution for all sessions. As a robustness check, we also consider the case where distributions are session-specific.

function $\tilde{V}_t^r(a, S_{i,t})$ for a pair of action and state as follows:

$$\begin{aligned}\tilde{V}_t^r(a, S_{i,t}) = &\left(\hat{x}_a - \frac{1-\hat{F}(\hat{x}_a)}{\hat{f}(\hat{x}_a)}\right)\hat{y}_{i,t}(a; S_{i,t}) \\ &+ \left(1 - \hat{l}_{i,t}(a; S_{i,t})\right)\hat{y}_{i,t}(a; S_{i,t})V_{t+1}^r\left(\langle S_{i,t}, a, Y_{i,t} = 1\rangle\right) \\ &+ \left(1 - \hat{l}_{i,t}(a; S_{i,t})\right)\left(1 - \hat{y}_{i,t}(a; S_{i,t})\right)V_{t+1}^r\left(\langle S_{i,t}, a, Y_{i,t} = 0\rangle\right),\end{aligned} \tag{43}$$

where $\hat{y}_{i,t}(a; S_{i,t})$ and $\hat{l}_{i,t}(a; S_{i,t})$ are estimated leave and click probabilities respectively. We can use Equation (43) and describe our backward induction solution concept as follows:

1. We begin from the last period T. Since there is no expected future at that point, the value function can be estimated as follows:

$$\hat{V}_T^c(S_{i,T}) = \max_{a \in \mathcal{A}_i} \hat{R}_t^r(a, S_{i,t}) \tag{44}$$

2. For any $t < T$, we can determine the value function as follows:

$$\hat{V}_t^c(S_{i,t}) = \max_{a \in \mathcal{A}_i} \tilde{V}_t^r(a, S_{i,T}) \tag{45}$$

   We can easily estimate the value function for any $t < T$ if we have all the value functions for the next periods. We can satisfy that by going backward and find the value function for all states at any t and continue this process until exposure number 1.

Once we have identified the value function for all the states, we can find the optimal allocation with dynamic objective as follows:

$$q_{i,a,t}^d(\hat{x}; S_{i,t}) = \begin{cases} 1 & a = \text{argmax}_{a \in \mathcal{A}_i} \tilde{V}_t^r(a, S_{i,t}), t < T \\ 1 & a = \text{argmax}_{a \in \mathcal{A}_i} \hat{R}_t^r(a, S_{i,t}), t = T \\ 0 & \text{otherwise} \end{cases} \tag{46}$$

Now, given the allocation, we can determine the payments using Equation (15). It is important to notice that the payment is determined in expectation over the entire session for each ad. For any ad a, the first component in Equation (15) is the average value ad a would get given the allocation, and the second component is the integral of the expected value ad a would get if she reduces her bid, taken over all possible bids that she could submit. The first component is easier to calculate as it is an expected sum of the total value each ad receives given each sequence. The second part, however, is more computationally intensive as it involves a numerical integration. To do that, we first need to estimate the function inside the integral. Let $\hat{Q}_{i,a}(b_a; S_{i,1})$ denote the expectation inside the integral in Equation (15). This will be the expected number of clicks ad a would get for any bid, given other players report their click valuations. We interpolate this function by finding its values for a set of points on the interval $[\underline{x}_a, x_a]$. Since the maximum number of expected number of clicks an ad could get is 6 when $T = 6$, we split this interval into 6 equally length intervals and find the function values for the points splitting the interval. We operationalize that with $h_a = \frac{\hat{x}_a - \hat{\underline{x}}_a}{T}$, that indicates the length of each interval for ad a. As such, if $b_a \in (\hat{\underline{x}}_a + (i-1)h_a, \hat{\underline{x}}_a + ih_a]$, we can estimate this function as follows:

$$\hat{Q}_{i,a}(b_a, \hat{x}_{-a}; S_{i,1}) = \sum_{t=1}^{T} \sum_{g_t \in \mathcal{G}_T} q^d_{i,a,t}(\hat{x}_a + ih_a, \hat{x}_{-a}; S_{i,t}) P(g_t \mid \tau, \pi), \tag{47}$$

where g_t contains the states and ads prior to period t. In Equation (47), we approximate the function $\mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right]$ with a step function, which makes the integral computation significantly easier. Now, we can determine the payments as follows:

$$\begin{aligned} e^d_{i,a}(x; S_{i,1}) = & \sum_{t=1}^{T} \sum_{g_t \in \mathcal{G}_T} \hat{x}_a \hat{y}_{i,t}(a; S_{i,t}) q^d_{i,a,t}(x; S_{i,t}) P(g_t \mid \tau, \pi) \\ & - h_a \sum_{t=1}^{T} \hat{Q}_{i,a}(\hat{x}_a + th_a, \hat{x}_{-a}; S_{i,1}), \end{aligned} \tag{48}$$

where the first component is the total value ad a receives from the session given the allocation, and the second component is the estimate of advertiser's surplus.

5.3 Evaluation

To evaluate the performance of each auction, we implement the direct method with the reward function with the revenue-maximizing objective. As such, we can define the expected revenue from session i with the horizon length T as $\rho^r_T(M; S_{i,1})$ for any mechanism M as follows:

$$\rho_T(M; S_{i,1}) = \mathbb{E}_{g_t \sim (\tau, M)} \left[\sum_{t=1}^{T} \beta^{t-1} r^r_t \right], \tag{49}$$

where the sequence g_t is determined by the joint distribution of transitions τ, and the allocation rule in mechanism M. Further, r^r_t denotes the reward function with the revenue-maximizing objective for the pair of state and action shown in the corresponding g_t. We can use our estimates for the distribution of transitions and evaluate $\rho_T(M; S_{i,1})$ for any mechanism M as follows:

$$\hat{\rho}_T(M; S_{i,1}) = \sum_{t=1}^{T} \sum_{g_t \in \mathcal{G}_T} \sum_{a \in \mathcal{A}_i} q^M_{i,a,t}(\hat{x}; S_{i,t}) \hat{R}^r_t(a; S_{i,t}) P(g_t \mid \tau, M) \tag{50}$$

This gives us the expected revenue that the publisher can extract from session i in the first T periods, when using the mechanism M.

6 Results

6.1 Gains from the Optimal Dynamic Auction

We start by presenting different session-level outcomes under both dynamic and static optimal auctions as well as the actual outcomes under the current auction – quasi-proportional auction. We evaluate the counterfactual outcomes under optimal auctions using our approach in §5. The sample of sessions that we use is the same as the one in Rafieian (2019). In our evaluation, we only focus on the first six exposures in any session, as we set $T = 6$ as the length of the horizon. We present our results in Table 2.

As expected, both optimal auctions generate substantial gains over the current auction in terms of all session-level outcomes, except the expected advertisers' surplus. Three key components of the current auction explain its performance relative to the optimal auctions. First, as discussed before, advertisers bid roughly half of their click valuation in this auction, lowering the revenues the publisher is able to extract. This, in turn, explains higher advertisers' surplus in the current auction as advertisers can extract huge informational rent by bidding half of their valuations. Second, the allocation mechanism is probabilistic in the current auction, which enables advertisers with low valuations to win impressions and clicks. This significantly reduces the efficiency or total surplus generated in the current auction compared to optimal auctions. Finally, as mentioned before, the ad-network does not incorporate sophisticated customization tools to show more relevant ads to users. This also greatly contributes to the gap between the current auction and other optimal auctions.

Next, we focus on our main goal in this paper, and compare the session-level outcomes under dynamic and static optimal auctions. We find that using an optimal dynamic auction leads to 1.60% increase in the expected revenue per session, compared to the optimal static auction. This is of particular importance to the publishers and ad-networks as they usually use optimal static auctions (e.g., second-price auction).

While revenue is the main outcome most publishers and ad-networks care about, they do not usually want to achieve better revenue outcomes at the expense of efficiency and advertisers' surplus.

	Auctions		
	Dynamic	*Static*	*Current*
Expected Publisher's Revenue Per Session	126.49	124.49	31.63
Expected Total Surplus Per Session	143.45	140.96	66.60
Expected Advertisers' Surplus Per Session	16.96	16.47	34.97
Expected No. of Clicks Per Session	0.1577	0.1549	0.0823
Expected Session Length	3.24	3.25	3.12
Expected CTR	4.86%	4.76%	2.64%
No. of Users	1000	1000	1000
No. of Sessions	12,136	12,136	12,136

Table 2: Market outcomes under different auctions for a sequence size of 6

Our results show that the optimal dynamic auction yields higher total and advertisers' surplus than the optimal static auction: the expected total surplus and advertisers' surplus grow by 1.77% and 3.00% respectively. Thus, advertisers also benefit from the revenue-optimal dynamic auction, as compared to the revenue-optimal static auction. Finally, we focus on the user-level outcomes – expected number of clicks per session and the session length. While the difference in the session length is very small, we find 1.83% increase in the expected number of clicks. This finding is in line with the findings in Rafieian (2019). More importantly, it suggest that the revenue gains from the optimal dynamic auction likely come from the improvement in the match between ads and users, and not from the greater ability to extract informational rent from advertisers.

6.2 Distribution of Advertisers' Surplus

Our results in Table 2 indicate that the average advertisers' surplus increases by 3.00% in the optimal dynamic auction, as compared to the optimal static auction. We now explore how this surplus is distributed across advertisers. We use two main approaches to examine the distribution of advertisers' surplus: (1) number of advertisers who benefit from the optimal dynamic auction compared to the optimal static auction, and (2) the Herfindahl-Hirschman Index (HHI) which is a well-known measure to study market concentration.

We first focus on the distribution of advertisers' surplus under both auctions. We compute the log of each advertiser's surplus over all 12,136 sessions and present it in Figure 5. Interestingly, we find that only 3 out of 15 advertisers benefit from the optimal dynamic auction compared to the one with the optimal static auction. In contrast, 9 out of 15 advertisers prefer the optimal static auction. However, it is worth noting that these advertisers have a very small surplus in both cases and their gains are small in magnitude. Therefore, the average advertisers' surplus is higher under

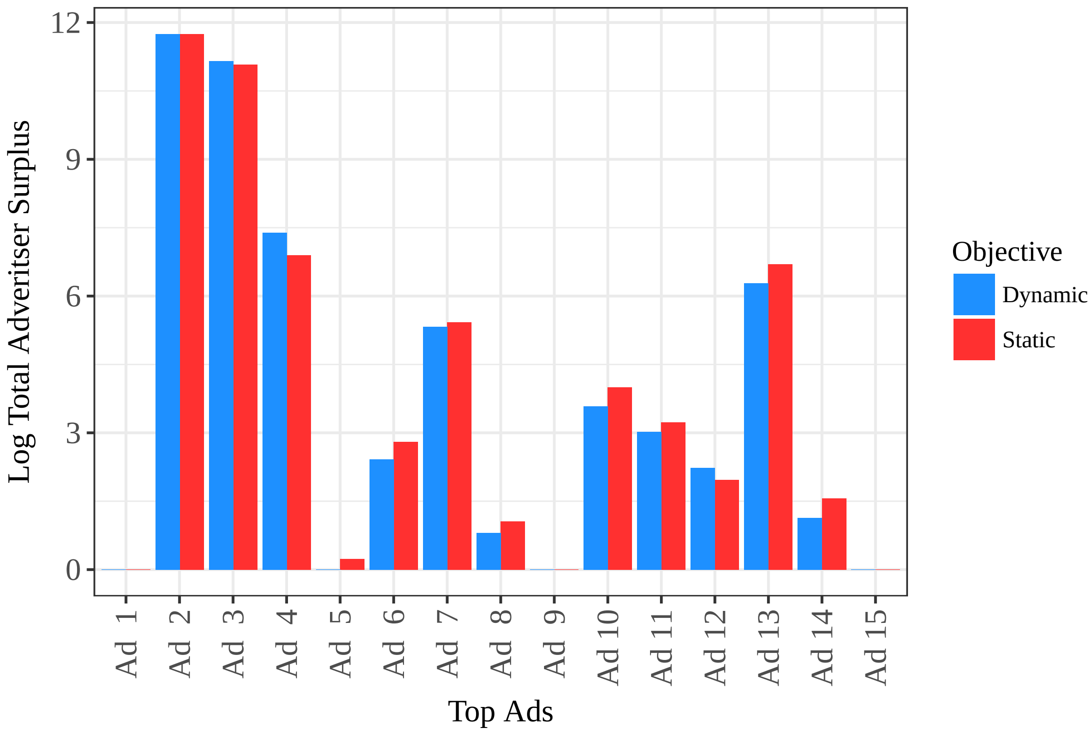

Figure 5: Log advertiser surplus under auctions with both dynamic and static objectives. The values are computed over all 12,136 session.

the optimal dynamic auction.

Next, we focus on the concentration of advertiser surplus in the market and calculate the Herfindahl-Hirschman Index (HHI) under both auctions. While more advertisers prefer the optimal static auction, the concentration in the optimal dynamic auction is lower: the Herfindahl-Hirschman Index (HHI) for the optimal dynamic auction is 0.5285, whereas it is 0.5398 for the optimal static auction. Although both auctions seem quite concentrated as they both allocate most impressions to one ad, optimal dynamic auction achieves a lower concentration by allocating more to the second-largest ad, thereby closing the gap between the two largest advertisers.

7 Conclusion

Mobile in-app advertising has grown exponentially over the last years. One important reason contributing to this growth is the publishers' ability to make adaptive interventions – using the time-varying information about the users within the session to personalize ad interventions for these users. Rafieian (2019) focuses on the match between ads and users as the main outcome of interest and shows that dynamic sequencing of ads improves the match outcome per session, compared static sequencing of ads. While dynamic sequencing of ads leads to improvements in the match outcome, it is not clear whether it can be linked to higher revenues, as advertisers can change their bids in response to the change in the allocation. In this paper, we investigate the revenue gains from adopting a dynamic sequencing framework in a competitive environment, as opposed to a static sequencing framework. We propose a unified framework that contains two key components: (1) a theoretical framework that solves for the revenue-optimal auction with the dynamic objective, and (2) an empirical framework that estimates the counterfactual market outcomes under this auction. Our empirical framework comprises of structural estimation of advertisers' private valuations as well as personalized estimation of the click outcome given any pair of user-ad using machine learning methods. We apply our framework to large-scale data from the leading in-app ad-network of an Asian country. We demonstrate that adopting a dynamic sequencing framework increases the expected revenue by 1.60% compared to the static sequencing framework. We then show that the improved match outcome is the key factor in achieving these gains. Further, we explore other outcomes such as the total surplus (efficiency) and advertisers' surplus and document gains from the dynamic framework over the static framework. Thus, our optimal dynamic auction leads to improvement in all primary market outcomes.

Our paper makes three key contributions to the literature. First, from a methodological standpoint, we present a unified dynamic framework that incorporates both advertisers' and users' behavior and examines the market outcomes under the optimal dynamic auction. To our knowledge, this is the first paper to empirically study the revenue gains from adopting an optimal dynamic auction. Second, we propose a methodological framework for structural estimation of quasi-proportional auction. Our framework adds to the literature on the non-parametric estimation of auctions and can easily be extended to any auction that employs a randomized allocation rule. Finally, from a substantive viewpoint, we show that dynamic sequencing of ads can lead to considerable gains in the publisher's revenue, over the existing auctions. This is particularly important, because the current practice in marketing ignores the gains from using a dynamic framework.

Our findings in this paper have several implications for marketing practitioners. First, we propose a framework that helps publishers evaluate market outcomes under counterfactual auctions. Second, our substantive finding shows that adopting an optimal dynamic auction results in considerable gains in terms of three key market outcomes – revenue, total surplus, and advertisers' surplus. Thus, we expect this finding to inform the publishers' decision as to whether use an optimal dynamic auction. Further, we highlight that adaptive interventions require more careful consideration of players' incentives, when there are strategic players whose actions can affect market outcomes. Thus, our framework has implications for the practitioners and policy makers that increasingly use adaptive interventions in the presence of strategic players.

Nevertheless, there remains some limitations in our study that serve as excellent avenues for future research. First, our optimal dynamic auction involves a bit complex allocation and payment rules. As a result, it may be possible that advertisers will not behave in the expected way. Finding auctions with easier rules that achieve approximately the same outcomes would be an interesting avenue for future study. Second, our structural estimation framework assumes that advertisers are risk-neutral and aware of how their bid affects their probability of winning. However, advertisers' behavior would be different if they are risk-averse. Future research can study risk aversion in probabilistic auctions and empirically identify the risk elements in advertisers' utility function. Finally, while our paper provides counterfactual estimates optimal auctions, we do not run these auctions in the field. An interesting future research is to test these auction in a field setting and examine market outcomes.

References

A. Asadpour, M. H. Bateni, K. Bhawalkar, and V. Mirrokni. Concise bid optimization strategies with multiple budget constraints. In *International Conference on Web and Internet Economics*, pages 263–276. Springer, 2014.

S. Athey and P. A. Haile. Nonparametric approaches to auctions. *Handbook of econometrics*, 6:3847–3965, 2007.

S. Athey and D. Nekipelov. A structural model of sponsored search advertising auctions. In *Sixth Ad Auctions Workshop*, volume 15, 2010.

S. Athey and I. Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

P. Bajari, S. Houghton, and S. Tadelis. Bidding for incomplete contracts: An empirical analysis of adaptation costs. *American Economic Review*, 104(4):1288–1319, 2014.

D. P. Baron and D. Besanko. Regulation and information in a continuing relationship. *Information Economics and policy*, 1(3):267–302, 1984.

D. Bergemann and J. Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.

C. Borgs, J. Chayes, N. Immorlica, K. Jain, O. Etesami, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*, pages 531–540. ACM, 2007.

R. E. Bucklin and P. R. Hoban. Marketing models for internet advertising. In *Handbook of marketing decision models*, pages 431–462. Springer, 2017.

S. Campo, I. Perrigne, and Q. Vuong. Asymmetry in first-price auctions with affiliated private values. *Journal of Applied Econometrics*, 18(2):179–207, 2003.

H. Choi and C. F. Mela. Online marketplace advertising. *Available at SSRN*, 2016.

H. Choi, C. F. Mela, S. R. Balseiro, and A. Leary. Online display advertising markets: A literature review and future directions. *Information Systems Research*, 2020.

J.-P. Dubé, G. J. Hitsch, and P. Manchanda. An empirical model of advertising dynamics. *Quantitative marketing and economics*, 3(2):107–144, 2005.

B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.

eMarketer. Mobile In-App Ad Spending , 2018. URL `https://forecasts-na1.emarketer.com/584b26021403070290f93a5c/5851918a0626310a2c186a5e`.

J. Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.

A. Ghosh and M. Mahdian. Externalities in online advertising. In *Proceedings of the 17th international conference on World Wide Web*, pages 161–168. ACM, 2008.

A. Ghosh and A. Sayedi. Expressive auctions for externalities in online advertising. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pages 371–380, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-799-8. doi: 10.1145/1772690.1772729. URL `http://doi.acm.org/10.1145/1772690.1772729`.

E. Guerre, I. Perrigne, and Q. Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.

E. Guerre, I. Perrigne, and Q. Vuong. Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. *Econometrica*, 77(4):1193–1227, 2009.

S. M. Kakade, I. Lobel, and H. Nazerzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013.

W. Kar, V. Swaminathan, and P. Albuquerque. Selection and ordering of linear online video ads. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, pages 203–210, New

York, NY, USA, 2015. ACM. ISBN 978-1-4503-3692-5. doi: 10.1145/2792838.2800194. URL http://doi.acm.org/10.1145/2792838.2800194.

D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. In *International Workshop on Internet and Network Economics*, pages 585–596. Springer, 2008.

E. Krasnokutskaya. Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies*, 78(1):293–327, 2011.

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

H. Li, S. M. Edwards, and J.-H. Lee. Measuring the intrusiveness of advertisements: Scale development and validation. *Journal of advertising*, 31(2):37–47, 2002.

V. Mirrokni, S. Muthukrishnan, and U. Nadav. Quasi-proportional mechanisms: Prior-free revenue maximization. In *Latin American Symposium on Theoretical Informatics*, pages 565–576. Springer, 2010.

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

R. B. Myerson. Multistage games with communication. *Econometrica: Journal of the Econometric Society*, pages 323–358, 1986.

D. C. Parkes and S. P. Singh. An mdp-based approach to online mechanism design. In *Advances in neural information processing systems*, pages 791–798, 2004.

A. Pavan, I. Segal, and J. Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82 (2):601–653, 2014.

O. Rafieian. Optimizing user engagement through adaptive ad sequencing. Technical report, Working paper, 2019.

O. Rafieian and H. Yoganarasimhan. How does variety of previous ads influence consumer's ad response? *Available at SSRN 3647373*, 2020a.

O. Rafieian and H. Yoganarasimhan. Targeting and privacy in mobile advertising, 2020b.

M. H. Riordan and D. E. Sappington. Information, incentives, and organizational mode. *The Quarterly Journal of Economics*, 102(2):243–263, 1987.

M. Said. Auctions with dynamic populations: Efficiency and revenue maximization. *Journal of Economic Theory*, 147(6):2419–2438, 2012.

H. R. Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.

J. M. Villas-Boas. Predicting advertising pulsing policies in an oligopoly: A model and empirical test. *Marketing Science*, 12(1):88–102, 1993.

G. Vulcano, G. Van Ryzin, and C. Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 48(11):1388–1407, 2002.

S. Yao and C. F. Mela. A dynamic model of sponsored search advertising. *Marketing Science*, 30(3):447–468, 2011.

H. Yoganarasimhan. Estimation of beauty contest auctions. *Marketing Science*, 35(1):27–54, 2015.

Appendices

A Proofs

Proof of Lemma 1. In a mechanism M, we can write the maximum utility advertiser a can receive as follows:

$$\max_{b_a} u^M_{i,a,t}(b_a; x_a, S_{i,t}) = \max_{b_a} \mathbb{E}_{x_{-a}} \left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) - e^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) \right], \quad (51)$$

where the expectation is over other advertisers' click valuations, as we have IC. The first derivative of advertiser a with respect to her click valuation x_a as follows:

$$\frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} = \mathbb{E}_{x_{-a}} \left[P(Y_{i,a} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) \right] \quad (52)$$

Given IC constraint, we know that $u^M_{i,a,t}(x_a; x_a, S_{i,t}) = \max_{b_a} u^M_{i,a,t}(b_a; x_a, S_{i,t})$. Therefore, based on envelope theorem, we have:

$$\frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} \Big|_{b_a = x_a} = \mathbb{E}_{x_{-a}} \left[P(Y_{i,a} \mid a, S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) \right] \quad (53)$$

Now, since u is differentiable, we can write:

$$\begin{aligned} u^M_{i,a,t}(x_a; x_a, S_{i,t}) - u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) &= \int_{x'_a}^{x_a} \frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} db_a \\ &= P(Y_{i,a} \mid a, S_{i,t}) \int_{x'_a}^{x_a} \mathbb{E}_{x_{-a}} \left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) \right] db_a \end{aligned} \quad (54)$$

This directly implies Equation (5) and completes the proof for Lemma 1. □

Proof of Lemma 2. We can write the publisher's revenues as follows:

$$\begin{aligned} \mathbb{E}_x \left[\sum_{a \in \mathcal{A}_i} e^M_{i,a,t}(x; S_{i,t}) \right] &= \sum_{a \in \mathcal{A}_i} \mathbb{E}_x \left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) \right] \\ &- \sum_{a \in \mathcal{A}_i} \mathbb{E}_x \left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) - e^M_{i,a,t}(x; S_{i,t}) \right], \end{aligned} \quad (55)$$

where each component of the second term is advertiser a's surplus and allows us to write it as follows:

$$\mathbb{E}_x \left[\sum_{a \in \mathcal{A}_i} e^M_{i,a,t}(x; S_{i,t}) \right] = \sum_{a \in \mathcal{A}_i} \mathbb{E}_x \left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) \right] - \sum_{a \in \mathcal{A}_i} \mathbb{E}_{x_a} \left[u^M_{i,a,t}(x_a; x_a, S_{i,t}) \right] \quad (56)$$

Using calculus theorems and Lemma 1, we can transform each element of the second term. For brevity and with slight abuse of notation, we drop $S_{i,t}$ from the function specification, and define $p_{i,a} = P(Y_{i,a} \mid a, S_{i,t})$.

We can write:

$$
\begin{aligned}
\mathbb{E}_{x_a}\left[u^M_{i,a,t}(x_a;x_a)\right] &= \int_{\underline{x}_a}^{\bar{x}_a} u^M_{i,a,t}(x_a) f_a(x_a) dx_a \\
&= \int_{\underline{x}_a}^{\bar{x}_a} \left(u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{\underline{x}_a}^{x_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a})\right] db_a \right) f_a(x_a) dx_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{\underline{x}_a}^{\bar{x}_a} \int_{\underline{x}_a}^{x_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a})\right] f_a(x_a) db_a dx_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{\underline{x}_a}^{\bar{x}_a} \int_{b_a}^{\bar{x}_a} f_a(x_a) \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a})\right] dx_a db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{\underline{x}_a}^{\bar{x}_a} \left(1 - F_a(b_a)\right) \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a})\right] db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{\underline{x}_a}^{\bar{x}_a} \left(1 - F_a(b_a)\right) \int_{x_{-a}} q^M_{i,a,t}(b_a, x_{-a}) f_{-a}(x_{-a}) dx_{-a} db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{\underline{x}_a}^{\bar{x}_a} \frac{1 - F_a(b_a)}{f_a(b_a)} \int_{x_{-a}} q^M_{i,a,t}(b_a, x_{-a}) f_{-a}(x_{-a}) dx_{-a} f_a(b_a) db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{x} \frac{1 - F_a(b_a)}{f_a(b_a)} q^M_{i,a,t}(b_a, x_{-a}) f(b_a, x_{-a}) dx_{-a} db_a \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \int_{x} \frac{1 - F_a(x_a)}{f_a(x_a)} q^M_{i,a,t}(x) f(x) dx \\
&= u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) + p_{i,a} \mathbb{E}_x\left[\frac{1 - F_a(x_a)}{f_a(x_a)} q^M_{i,a,t}(x)\right]
\end{aligned} \tag{57}
$$

Now we can transform the publisher's revenues in Equation (56) using Equation (57) to complete the proof:

$$
\begin{aligned}
\mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i} e^M_{i,a,t}(x;S_{i,t})\right] &= \sum_{a\in\mathcal{A}_i} \mathbb{E}_x\left[w_{i,a,t}(x_a;S_{i,t}) q^M_{i,a,t}(x;S_{i,t})\right] \\
&\quad - \sum_{a\in\mathcal{A}_i} u^M_{i,a,t}(\underline{x}_a;\underline{x}_a) - \sum_{a\in\mathcal{A}_i} p_{i,a} \mathbb{E}_x\left[\frac{1 - F_a(x_a)}{f_a(x_a)} q^M_{i,a,t}(x)\right] \\
&= \mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i} \left(x_a - \frac{1 - F_a(x_a)}{f_a(x_a)}\right) p_{i,a} q^M_{i,a,t}(x;S_{i,t})\right] - \sum_{a\in\mathcal{A}_i} u^M_{i,a,t}(\underline{x}_a;\underline{x}_a, S_{i,t})
\end{aligned}
$$

□

Proof of Proposition 1. Given the payments in Equation (8), it is easy to check that $u^M_{i,a,t}(\underline{x}_a;\underline{x}_a, S_{i,t}) = 0$.

We can write:

$$
\begin{aligned}
u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) &= \mathbb{E}_{x_{-a}}\left[w_{i,a,t}(\underline{x}_a; S_{i,t}) q^M_{i,a,t}(\underline{x}_a, x_{-a}; S_{i,t}) - e^M_{i,a,t}(\underline{x}_a, x_{-a}; S_{i,t})\right] \\
&= \mathbb{E}_{x_{-a}}\left[P(Y_{i,t} \mid a, S_{i,t}) \int_{\underline{x}_a}^{\underline{x}_a} q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) db_a\right] \\
&= 0
\end{aligned}
\tag{58}
$$

This implies that the second part of the publisher's revenues in Equation (6) is zero, i.e., $\sum_{a \in \mathcal{A}_i} u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) = 0$. Thus, since the q is chosen to maximize the first part of Equation (6) given some constraints, we know that mechanism M is optimal given those constraint. It is now sufficient to show the following two statements: 1) mechanism M is a direct revelation mechanism, and 2) any direct mechanism satisfies the constraints: q is plausible and $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a.

We start by showing that M is a direct revelation mechanism. The plausibility is satisfied by definition. We only need to show both IR and IC. Given the payment function, we can write the utility function for advertiser a as follows:

$$
\begin{aligned}
u^M_{i,a,t}(x_a; x_a, S_{i,t}) &= \mathbb{E}_{x_{-a}}\left[w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x; S_{i,t}) - e^M_{i,a,t}(x; S_{i,t})\right] \\
&= \mathbb{E}_{x_{-a}}\left[P(Y_{i,t} \mid a, S_{i,t}) \int_{\underline{x}_a}^{x_a} q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) db_a\right] \\
&= P(Y_{i,t} \mid a, S_{i,t}) \int_{\underline{x}_a}^{x_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a
\end{aligned}
\tag{59}
$$

Given $u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) = 0$, this is equivalent to the envelope condition. Since the integral on the RHS is always non-negative, $u_{i,a,t}$ is increasing, which combined with $u^M_{i,a,t}(\underline{x}_a; \underline{x}_a, S_{i,t}) = 0$ imply IR constraint. We now need to show IC constraint is also satisfied. We prove that by contradiction. Suppose that there is an x'_a that gives a higher utility to advertiser a than truthful reporting, given everyone else bidding their true click valuations. Let γ denote the gains advertiser a receives by reporting x'_a instead of x_a. We can write:

$$
\begin{aligned}
\gamma &= u^M_{i,a,t}(x'_a; x_a, S_{i,t}) - u^M_{i,a,t}(x_a; x_a, S_{i,t}) \\
&= u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) - \left(u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) - u^M_{i,a,t}(x'_a; x_a, S_{i,t})\right) - u^M_{i,a,t}(x_a; x_a, S_{i,t}) \\
&= u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) - u^M_{i,a,t}(x_a; x_a, S_{i,t}) - (x'_a - x_a) P(Y_{i,t} \mid a, S_{i,t}) \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] \\
&= P(Y_{i,t} \mid a, S_{i,t}) \left(\int_{x_a}^{x'_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a - (x'_a - x_a) \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right]\right) \\
&= P(Y_{i,t} \mid a, S_{i,t}) \int_{x_a}^{x'_a} \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) - q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] db_a
\end{aligned}
\tag{60}
$$

Now, given that $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a, it is easy to check $\gamma \leq 0$ regardless of whether $x'_a > x_a$ or not. This completes the proof of part 1: mechanism M is a direct revelation mechanism.

Now, we show the second part: any direct revelation mechanism satisfies the constraints: q is plausible and $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a. A direct revelation mechanism M satisfies IC constraints.

Hence, for $x'_a > x_a$ we can write:

$$u^M_{i,a,t}(x_a; x_a, S_{i,t}) \geq u^M_{i,a,t}(x'_a; x_a, S_{i,t}) \tag{61}$$
$$u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) \geq u^M_{i,a,t}(x_a; x'_a, S_{i,t}) \tag{62}$$

Subtracting these two equations, we have:

$$u^M_{i,a,t}(x_a; x_a, S_{i,t}) - u^M_{i,a,t}(x_a; x'_a, S_{i,t}) \geq u^M_{i,a,t}(x'_a; x_a, S_{i,t}) - u^M_{i,a,t}(x'_a; x'_a, S_{i,t}) \tag{63}$$

Simplifying Equation (73) gives us:

$$(x'_a - x_a)\left(\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] - \mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]\right) \geq 0 \tag{64}$$

Since $x'_a > x_a$, we can show that $\mathbb{E}_{x_{-a}}\left[q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a. This completes the proof for the second part. □

Proof of Lemma 3. The proof for this lemma is almost identical to the proof for Lemma 1. We start by writing down the maximization problem advertiser a faces in mechanism M:

$$\max_{b_a} U^M_{i,a}(b_a; x_a, S_i) = \max_{b_a} \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}\left(w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t}) - e^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right)\right], \tag{65}$$

where the expectation is over other advertisers' click valuations as well as the stochasticity induce by the dynamic process. The reason we take the expectation over other advertisers' click valuation is the main condition in the lemma: IC constraint. We can write the first derivative of advertiser a with respect to her click valuation x_a as follows:

$$\frac{\partial U^M_{i,a}(b_a; x_a, S_i)}{\partial x_a} = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right], \tag{66}$$

where, again, the expectation is over other advertisers' click valuations as well as the stochasticity induce by the dynamic process. IC constraint implies that $U^M_{i,a}(x_a; x_a, S_i) = \max_{b_a} U^M_{i,a}(b_a; x_a, S_i)$. We can now apply the envelope theorem as follows:

$$\frac{\partial u^M_{i,a,t}(b_a; x_a, S_{i,t})}{\partial x_a} \Big|_{b_a = x_a} = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] \tag{67}$$

Now, due to the first-differentiability of U, we have:

$$\begin{aligned} U^M_{i,a}(x_a; x_a, S_i) - U^M_{i,a}(x'_a; x'_a, S_{i,t}) &= \int_{x'_a}^{x_a} \frac{\partial U^M_{i,a}(b_a; x_a, S_i)}{\partial x_a} db_a \\ &= \int_{x'_a}^{x_a} \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a, \end{aligned} \tag{68}$$

which is equivalent to Equation (12) and completes the proof for Lemma 3. □

Proof of Lemma 4. The steps for this proof is almost identical to the proof for Lemma 2. We start by writing the publisher's objective function:

$$
\begin{aligned}
\mathbb{E}\left[\sum_{a\in\mathcal{A}_i} e^M_{i,a}(x;S_i)\right] &= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right] \\
&\sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t}) - e^M_{i,a}(x;S_i)\right] \\
&= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right] - \sum_{a\in\mathcal{A}_i}\mathbb{E}_{x_a}\left[U^M_{i,a}(x_a;x_a,S_i)\right],
\end{aligned} \tag{69}
$$

where all the expectations are over the specified click valuations and the stochasticity induced by the dynamic process. Now, we can transform each element of the second term. For brevity and with slight abuse of notation, we drop S_i and $S_{i,t}$ from the function specification, and define $p_{i,a,t} = P(Y_{i,a} \mid a, S_{i,t})$. We can write:

$$
\begin{aligned}
\mathbb{E}_{x_a}\left[U^M_{i,a}(x_a;x_a)\right] &= \int_{\underline{x}_a}^{\bar{x}_a} U^M_{i,a}(x_a)f_a(x_a)dx_a \\
&= \int_{\underline{x}_a}^{\bar{x}_a}\left(U^M_{i,a}(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{x_a}\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(b_a,x_{-a})\right]db_a\right)f_a(x_a)dx_a \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a}\int_{\underline{x}_a}^{x_a}\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(b_a,x_{-a})\right]f_a(x_a)db_a dx_a \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a}\int_{b_a}^{\bar{x}_a} f_a(x_a)\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(b_a,x_{-a})\right]dx_a db_a \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a}(1-F_a(b_a))\,\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(b_a,x_{-a})\right]db_a \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + p_{i,a}\int_{\underline{x}_a}^{\bar{x}_a}(1-F_a(b_a))\int_{x_{-a}}\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(b_a,x_{-a})f_{-a}(x_{-a})dx_{-a}db_a \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + \int_{\underline{x}_a}^{\bar{x}_a}\frac{1-F_a(b_a)}{f_a(b_a)}\int_{x_{-a}}\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(b_a,x_{-a})f_{-a}(x_{-a})dx_{-a}f_a(b_a)db_a \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + \int_{x}\frac{1-F_a(b_a)}{f_a(b_a)}\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(b_a,x_{-a})f(b_a,x_{-a})dx_{-a}db_a \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + \int_{x}\frac{1-F_a(x_a)}{f_a(x_a)}\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(x)f(x)dx \\
&= U^M_{i,a}(\underline{x}_a;\underline{x}_a) + p_{i,a}\mathbb{E}_x\left[\frac{1-F_a(x_a)}{f_a(x_a)}\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(x)\right]
\end{aligned} \tag{70}
$$

Now we can re-write Equation (69) using Equation (70) to complete the proof:

$$
\begin{aligned}
\mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i} e^M_{i,a}(x;S_i)\right] &= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}w_{i,a,t}(x_a;S_{i,t})q^M_{i,a,t}(x;S_{i,t})\right] \\
&\quad -\sum_{a\in\mathcal{A}_i}U^M_{i,a}(\underline{x}_a;\underline{x}_a)-\sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\frac{1-F_a(x_a)}{f_a(x_a)}\sum_{t=1}^{\infty}\beta^{t-1}p_{i,a,t}q^M_{i,a,t}(x)\right] \\
&= \mathbb{E}_x\left[\sum_{a\in\mathcal{A}_i}\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)p_{i,a,t}q^M_{i,a,t}(x;S_{i,t})\right] \\
&\quad -\sum_{a\in\mathcal{A}_i}U^M_{i,a}(\underline{x}_a;\underline{x}_a,S_i) \\
&= \sum_{a\in\mathcal{A}_i}\mathbb{E}_x\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)p_{i,a,t}q^M_{i,a,t}(x;S_{i,t})\right] \\
&\quad -\sum_{a\in\mathcal{A}_i}U^M_{i,a}(\underline{x}_a;\underline{x}_a,S_i)
\end{aligned}
$$

□

Proof of Proposition 2. The proof for this proposition is very similar to the one for Proposition 1. We begin by providing the necessary and sufficient conditions for the IC constraint. We first show that a mechanism M is IC if and only if $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ is increasing in x_a and we have the envelope condition as presented in Equation (12).

We start our proof by the only if part. We want to show that if the mechanism M is IC, then $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ is increasing in x_a and we have the envelope condition as presented in Equation (12). The latter is the result of Lemma 3. So we only need to show the that IC implies monotonicity. Given the IC, for any $x'_a > x_a$ we can write:

$$U^M_{i,a}(x_a;x_a,S_i)\geq U^M_{i,a}(x'_a;x_a,S_i) \tag{71}$$
$$U^M_{i,a}(x'_a;x'_a,S_i)\geq U^M_{i,a}(x_a;x'_a,S_i) \tag{72}$$

Subtracting these two equations, we have:

$$U^M_{i,a}(x_a;x_a,S_i)-U^M_{i,a}(x_a;x'_a,S_i)\geq U^M_{i,a}(x'_a;x_a,S_i)-U^M_{i,a}(x'_a;x'_a,S_i) \tag{73}$$

Simplifying Equation (73) gives us:

$$(x'_a-x_a)\left(\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]\right)\geq 0 \tag{74}$$

Since $x'_a > x_a$, we can show that $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ is increasing in x_a. This completes the *only if* part of the statement.

Now, we show the *if* part of the statement. If we have plausibility and monotonicity conditions as above, then we IC constraint is satisfied. We assume that the IC is not satisfied and then show contradiction. If IC is not satisfied, there exists an x'_a that gives a higher utility to advertiser a than truthful reporting, given IC for other advertisers. We denote the gains from deviating by γ. We can write:

$$\begin{aligned}
\gamma &= U^M_{i,a}(x'_a; x_a, S_i) - U^M_{i,a}(x_a; x_a, S_i) \\
&= U^M_{i,a}(x'_a; x'_a, S_i) - \left(U^M_{i,a}(x'_a; x'_a, S_i) - U^M_{i,a}(x'_a; x_a, S_i)\right) - U^M_{i,a}(x_a; x_a, S_i) \\
&= U^M_{i,a}(x'_a; x'_a, S_i) - U^M_{i,a}(x_a; x_a, S_i) \\
&\quad - (x'_a - x_a)\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] \\
&= \int_{x_a}^{x'_a} \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a \\
&\quad - (x'_a - x_a)\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] \\
&= \int_{x_a}^{x'_a} \mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right. \\
&\quad \left. - \sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x'_a, x_{-a}; S_{i,t})\right] db_a
\end{aligned} \tag{75}$$

Now, given that $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ is increasing in x_a, it is easy to check $\gamma \leq 0$ independent of the relationship between x'_a and x_a. This contradicts the assumption that $\gamma > 0$ and shows that mechanism M is IC.

Now we show that the proposed mechanism is optimal. This mechanism maximizes the first component in Equation (13) subject to the plausibility and monotonicity conditions, while setting the payment such that the second component is zero. Given the payments, we can write:

$$\begin{aligned}
U^M_{i,a}(x_a; x_a, S_i) &= \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(w_{i,a,t}(x_a; S_{i,t}) q^M_{i,a,t}(x_a, x_{-a}; S_{i,t}) - e^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right)\right] \\
&= \int_{\underline{x}_a}^{x_a} \mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(b_a, x_{-a}; S_{i,t})\right] db_a
\end{aligned} \tag{76}$$

Equation Equation (76) shows that $U^M_{i,a}(\underline{x}_a; \underline{x}_a, S_i) = 0$ for all a, which implies that we have the envelope condition as in Equation (12). Together, these imply IR constraint. Further, the envelope condition and the monotonicity constraint are equivalent to the IC constraint. Therefore, IC is also satisfied in our case.

Now we only need to show the optimality of the mechanism M. Since $\sum_{a \in \mathcal{A}_i} U^M_{i,a}(\underline{x}_a; \underline{x}_a, S_i) = 0$, the choice of q maximizes the publisher's objective given the plausibility and monotonicity constraint. Since these two constraints are necessary for any direct revelation mechanism, the mechanism M is optimal.

This equivalence proves two statements: 1) IC is satisfied, and 2) any direct revelation mechanism must satisfy monotonicity and plausibility. $\square$

Proof of Lemma 5. We only need to show that $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty} \beta^{t-1} P(Y_{i,t} \mid a, S_{i,t}) q^M_{i,a,t}(x_a, x_{-a}; S_{i,t})\right]$ increasing in x_a. Let q^M and $q^{M'}$ denote the optimal allocation functions derived by Equation (17) and Equation (18) for click valuation profiles x and x' respectively. Since these mechanisms are optimal for corresponding cases, we can write the following two inequalities:

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a;S_{i,t})\right)\right] \\ &\geq \mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right)\right], \end{aligned} \tag{77}$$

and

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x'_a-\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right)\right] \\ &\geq \mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x'_a-\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M}_{i,a,t}(x_a;S_{i,t})\right)\right] \end{aligned} \tag{78}$$

Subtracting these two inequalities will give us the following inequality:

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a;S_{i,t})\right)\right] \\ &\geq \mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(\sum_{a\in\mathcal{A}_i}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right)\right] \end{aligned} \tag{79}$$

Now, suppose that x and x' are the same at each element except the a-th element, i.e., $x_a \neq x'_a$ and $x_j = x'_j$ for all $j \neq a$. Further, suppose that $x_a > x'_a$. We can then write:

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a;S_{i,t})\right] \\ &\geq \mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}\left(x_a-\frac{1-F_a(x_a)}{f_a(x_a)}-x'_a+\frac{1-F_a(x'_a)}{f_a(x'_a)}\right)P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right] \end{aligned} \tag{80}$$

Now, since the distribution F_a is regular, we have:

$$\mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a;S_{i,t})\right] \geq \mathbb{E}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^{M'}_{i,a,t}(x'_a;S_{i,t})\right] \tag{81}$$

The last inequality directly implies $\mathbb{E}_{x_{-a}}\left[\sum_{t=1}^{\infty}\beta^{t-1}P(Y_{i,t}\mid a,S_{i,t})q^M_{i,a,t}(x_a,x_{-a};S_{i,t})\right]$ increasing in x_a and completes the proof. $\square$

Proof of Lemma 6. To write the FOC, we first need to take the first derivative of the expected utility function for ad a. We can write:

$$\begin{aligned}\frac{\partial u_a^p(b_a;x_a)}{\partial b_a} &= \frac{\partial}{\partial b_a}\big(x_a\tilde{q}_a^p(b_a) - \tilde{e}_a^p(b_a)\big)m_a \\ &= \frac{\partial}{\partial b_a}\big(x_a - \epsilon_a^p(b_a)\big)\tilde{q}_a^p(b_a)m_a \\ &= \left(x_a\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a} - \epsilon_a^p(b_a)\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a} - \tilde{q}_a^p(b_a)\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\right)m_a,\end{aligned} \tag{82}$$

where the second line comes from Assumption 6. We now use Assumption 7 and derive the following relationship for the first derivative of the allocation function:

$$\frac{\tilde{q}_a^p(b_a)}{\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}} = \frac{\frac{b_a z_a}{b_a z_a+\zeta_a}}{\frac{z_a\zeta_a}{(b_a z_a+\zeta_a)^2}} = \frac{b_a(b_a z_a+\zeta_a)}{\zeta_a} = \frac{b_a}{1-\tilde{q}_a^p(b_a)} \tag{83}$$

Now, we need to have $\frac{\partial u_a^p(b_a;x_a)}{\partial b_a} = 0$ for the equilibrium bid b_a. Using Equation (82), we can write the FOC as follows:

$$\begin{aligned}x_a &= \epsilon_a^p(b_a) + \frac{\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\tilde{q}_a^p(b_a)}{\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}} \\ &= \epsilon_a^p(b_a) + \frac{\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}b_a}{1-\tilde{q}_a^p(b_a)},\end{aligned} \tag{84}$$

where the second line is resulted directly from Equation (83). Now, to complete the proof, we need to show that the second-order condition (SOC) is also satisfied. We start by writing a useful property in the relationship between the first and second derivative of the allocation function:

$$\frac{\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}}{\frac{\partial^2\tilde{q}_a^p(b_a)}{\partial b_a^2}} = \frac{\frac{z_a\zeta_a}{(b_a z_a+\zeta_a)^2}}{\frac{-2z_a^2\zeta_a}{(b_a z_a+\zeta_a)^3}} = \frac{(b_a z_a+\zeta_a)}{2z_a} = -\frac{b_a}{2\tilde{q}_a^p(b_a)} \tag{85}$$

We can now write the second derivative of the expected utility function as follows:

$$\begin{aligned}
\frac{\partial^2 u_a^p(b_a;x_a)}{\partial b_a^2} &= m_a\left((x_a-\epsilon_a^p(b_a))\frac{\partial^2\tilde{q}_a^p(b_a)}{\partial b_a^2}-\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}-\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}-\frac{\partial^2\epsilon_a^p(b_a)}{\partial b_a^2}\tilde{q}_a^p(b_a)\right)\\
&= m_a\left(\frac{\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\tilde{q}_a^p(b_a)}{\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}}\frac{\partial^2\tilde{q}_a^p(b_a)}{\partial b_a^2}-2\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}-\frac{\partial^2\epsilon_a^p(b_a)}{\partial b_a^2}\tilde{q}_a^p(b_a)\right)\\
&= m_a\left(\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\tilde{q}_a^p(b_a)\left(\frac{\frac{\partial^2\tilde{q}_a^p(b_a)}{\partial b_a^2}}{\frac{\partial\tilde{q}_a^p(b_a)}{\partial b_a}}-2\frac{\pi'_a(b_a)}{\pi_a(b_a)}\right)-\frac{\partial^2\epsilon_a^p(b_a)}{\partial b_a^2}\pi_a(b_a)\right)\\
&= m_a\left(\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\tilde{q}_a^p(b_a)\left(\frac{2\pi_a(b_a)}{b_a}-2\frac{1-\pi_a(b_a)}{b_a}\right)-\frac{\partial^2\epsilon_a^p(b_a)}{\partial b_a^2}\pi_a(b_a)\right)\\
&= m_a\left(-\frac{2}{b_a}\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}\tilde{q}_a^p(b_a)-\frac{\partial^2\epsilon_a^p(b_a)}{\partial b_a^2}\pi_a(b_a)\right)\\
&= m_a\tilde{q}_a^p(b_a)\left(-\frac{2}{b_a}\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}-\frac{\partial^2\epsilon_a^p(b_a)}{\partial b_a^2}\right)
\end{aligned} \tag{86}$$

In the equation above, we know that both m_a and $\tilde{q}_a^p(b_a)$ are positive, so it is sufficient to show that $-\frac{2}{b_a}\frac{\partial\epsilon_a^p(b_a)}{\partial b_a}-\frac{\partial^2\epsilon_a^p(b_a)}{\partial b_a^2}\leq 0$, which is direct result of the condition in Lemma 6, i.e., $b^2\frac{\partial\epsilon_a^p(b)}{\partial b}$ is increasing in the local neighbourhood around b_a. Thus, the SOC is satisfied and this completes the proof. □

Proof of Lemma 7. The proof is straighforward for this lemma. First, we know that the advertiser has participated in the auction. Therefore, we know her click valuation was above the reserve price. That is:

$$b_0 \leq x_a \tag{87}$$

On the other hand, we know that her equilibrium bid is an upper bound for what would satisfy her FOC. Thus, we can use the FOC from Lemma 6 and write:

$$x_a \leq \epsilon_a^p(b_0)+\frac{b_0\frac{\partial\epsilon_a^p(b)}{\partial b}\,|_{b=b_0}}{1-\tilde{q}_a(b_0)} \tag{88}$$

□