Multi-sided Advertising Markets: Dynamic Mechanisms and Incremental User Compensations

Moran Feldman
Math. and Computer Science Dept.
The Open University of Israel
moranfe@openu.ac.il

Gonen Frim
Math. and Computer Science Dept.
The Open University of Israel
gonen987@gmail.com

Rica Gonen
Management and Economics Dept.
The Open University of Israel
gonenr@openu.ac.il

Abstract

Online advertising has motivated companies to collect vast amounts of information about users, which increasingly creates privacy concerns. One way to answer these concerns is by enabling end users to choose which aspects of their private information can be collected. Based on principles suggested by Feldman and Gonen (2016), we introduce a new online advertising market model which uses information brokers to give users such control. Unlike (Feldman and Gonen 2016), our model is dynamic and involves multi-sided markets where all participating sides are strategic. We describe a mechanism for this model which is theoretically guaranteed to (approximately) maximize the gain from trade, avoid a budget deficit and incentivize truthfulness and voluntary participation. As far as we know, this is the first known *dynamic* mechanism for a multi-sided market having these properties.

We experimentally examine and compare our theoretical results using real world advertising bid data. The experiments suggest that our mechanism performs well in practice even in input regimes for which our theoretical guarantee is weak or not relevant.

1 Introduction

Online advertising currently supports some of the most important Internet services, including: search, social media and user generated content sites. For online advertising to be effective, companies collect vast amounts of information about users, which increasingly creates privacy concerns (Conitzer, Taylor, and Wagman 2012). As these concerns are especially pronounced in the European society, EU regulators have actively been looking for ways to improve users' privacy. One way suggested by the EU regulators toward this goal is development of tools that enable end users to choose which parts of their private information online advertising platforms are allowed to collect.

Based on this motivation, and extending principles suggested by Feldman and Gonen (2016), we introduce a new model capturing a foreseeable future form of online advertising. The market in this model includes advertisers as buyers, users as sellers (each willing to sell her own information portfolio through a broker) and information brokers as mediators representing the users. The objective of a mechanism for this setting is to end up with a match between users and advertisers maximizing the gain from trade. Towards that goal, the mechanism has to collect information from the mediators and advertisers; and thus, needs to incentivize the mediators and advertisers to report truthfully, which it can do by charging the advertisers and paying the mediators. Additionally, unlike in (Feldman and Gonen 2016), we assume here that the users are strategic as well, which requires the mechanism to incentivize them also by recommending for each mediator to forward some of the payment he received to his users.

As the online advertising ecosystem is dynamic, the market in our model is dynamic as well. We assume the mediators and advertisers arrive at a uniformly random order, and refer to the arriving advertisers and mediators as arriving entities. Every time that a new entity arrives, the mechanism has an opportunity to assign users to advertisers. More specifically, users enroll with a mediator offline, i.e., when the mediator arrives at the market it has a list of users that are his customers and the mechanism is allowed to assign users of the newly arriving mediator to advertisers that have already arrived. Similarly, when an advertiser arrives the mechanism is allowed to assign users of mediators that have already arrived to the newly arriving advertiser. The result of an assignment is that when a user of a mediator?s deal views interstitial advertising it will be from the assigned advertiser. Notice that this means that the mechanism is not allowed to cancel assignments that have already been made, or assign a user of a mediator that has already arrived to an advertiser that has already arrived. These restrictions, together with the random arrival order, represent the dynamicity of the setting. We note that our choice to model a dynamic market using a random arrival order is a well established practice—for a few examples, see (Babaioff et al. 2009; Vaze and Coupechoux 2016). Intuitively, this modeling choice reflects the assumption that real arrival orders are arbitrary rather than adversarial.

A natural expectation from a dynamic exchange mechanism is to (approximately) maximize the gain from trade, while maintaining desirable economic properties such as incentivizing truthfulness, voluntary participation and avoiding budget deficit. Unfortunately, as far as we know, no previous work has managed to achieve these goals simultaneously. Wurman, Walsh, and Wellman (1998) presented a mechanism incentivizing truthful reporting from either the buyers or the sellers, but not simultaneously from both. A different mechanism given by Blum, Sandholm, and Zinkevich (2002) maximizes the social welfare of buyers and non-selling sellers (as opposed to maximizing the gain from trade). Finally, Bredin, Parkes, and Duong (2007) present a truthful dynamic double-sided auction that is constructed from a truthful offline double-sided auction rule, however its competitiveness with respect to the optimal trade was only studied empirically.

The failure of the above works to maximize the gain from trade while maintaining truthfulness, individual rationality (voluntary participation) and budget balance (avoiding budget deficit) can be partially attributed to an impossibility result of (Myerson and Satterthwaite 1983). This impossibility result states that, even in an offline setting involving a single buyer and a single seller, maximizing the gain from trade while maintaining truthfulness and individual rationality perforce runs a deficit (i.e., is not budget balanced). An additional reason for the above failure is that the matching problem faced by the market maker (exchange mechanism) in multi-sided dynamic markets combines elements of dynamic algorithms and sequential decision making with considerations from mechanism design. More specifically, unlike in a traditional dynamic algorithm, a mechanism for such a setting must provide incentives for agents to report truthful information to the mechanism. On the other hand, unlike in traditional mechanism design, this is a dynamic setting with agents that arrive over time, and the mechanism must deal with uncertainty and make irrevocable decisions before the arrival of all the agents.

1.1 Our Result

In this work we present the first (to the best of our knowledge) dynamic mechanism for a multi-sided market setting which theoretically guarantees the economic properties of truthfulness, individual rationality, and budget balance while (approximately) maximizing the gain from trade. As our setting involves multi-dimensional agents, our result shows that dynamic multi-sided markets can be handled even in the presence of multi-dimensional agents. Moreover, we study the practical performance of our mechanism using simulations based on real-world advertisers' bids. The data for these experiments was gathered from Facebook advertising campaigns. These experiments suggest that our mechanism performs well in practice even in input regimes for which our theoretical guarantee is weak.

The dynamic nature of our setting raises the question of what it means for a mechanism to be individually rational. As usual, individual rationality should imply that an agent never losses by participating. However, in a dynamic setting it is natural to require also that an agent never losses by not leaving prematurely. We introduce a new concept called "continuous individual rationality" which captures the above intuitive requirement. Formally, a mechanism is *continuously individually rational* for an agent (a user, a mediator or an advertiser) if the agent's utility can only increase over time when the agent is truthful[1]. Note that this newly presented concept of continuous individual rationality is a stronger concept than individual rationality in its classic form as it implies ex post individual rationality.

Satisfying the requirements of continuous individual rationality, together with the other economic properties our mechanism guarantees, requires our mechanism to use a novel pricing scheme where users may be paid ongoing increments during the mechanism's execution. The maximum total payment that a user may end up with is pre-known (when the user arrives), however, the actual increments are not pre-known and depend on the market's dynamically changing demands and supplies. As users rarely ever get paid in reality, this pricing scheme is new to mechanism design and might look odd at first glance. Nevertheless, the principle it is based on can be observed in many common real life scenarios such as executive compensation payments and company acquisition deals. For example, the eBay acquisition of Skype in 2005 involved both an upfront payment and an additional payment whose amount depended on the future performance of the bought company (see `https://investors.ebayinc.com/releasedetail.cfm?releaseid=176402`).

Like in (Babaioff, Feldman, and Tennenholtz 2016), we say that a mechanism is *user-side incentive compatible* if truthfulness is a dominant strategy[2] for each user given that her mediator is truthful. Similarly, the mechanism is *user-side continuously individually rational* if it is continuously individually rational for each user given that her mediator is truthful. A mechanism is *mediator-side incentive compatible* if truthfulness is a dominant strategy for each mediator whose users are all truthful, and it is *mediator-side continuously individually rational* if it is continuously individually rational for every such mediator. Finally, a mechanism is *advertiser-side incentive compatible* if truthfulness is a dominant strategy for every advertiser, and it is *advertiser-side continuously individually rational* if it is continuously individually rational for every advertiser. We construct a mechanism which is *three-sided incentive compatible* (i.e., it is simultaneously user-side incentive compatible, mediator-side incentive compatible and advertiser-side incentive compatible) and also *three-sided continuously individually rational* (i.e., it is simultaneously user-side continuously individually rational, mediator-side continuously individually rational and advertiser-side continuously individually rational).

Our mechanism is termed "Observe and Price Mechanism" (`OPM`). The following theorem analyzes the economic properties guaranteed by `OPM` and its competitive ratio. The parameter α is an upper bound, known to the mechanism, on the market importance of any single agent. Formally, α bounds the ratio between the size of the optimal trade and the maximum capacity of an advertiser or the maximum number of users that a mediator can represent.

[1] Informally, an agent is truthful if he/she reports the information as it is known to him/her. A formal definition of what it means for a user, mediator or advertiser to be truthful is given in Section 2.

[2] Here and throughout the paper, a reference to domination of strategies should be understood as a reference to weak domination.

Theorem 1.1. *`OPM` is budget balanced*[3]*, three-sided continuously individually rational, three-sided incentive compatible and* $(1 - 9.5\sqrt[6]{\alpha} - 10e^{-2/\sqrt[3]{\alpha}})$*-competitive.*

From a theoretical perspective, the most important feature of this ratio is that it approaches 1 when no agent has too much market power. Though this is a desirable aspect of the algorithm, we are aware that the competitive ratio has an unintuitive form and is often non-positive for markets of a moderate size. The later can be significantly alleviated by making the proof tighter (and less readable). Instead we chose to address intuitiveness and readability by including experimental results in the paper. The experimental results demonstrate that our mechanism performs well on inputs derived by real world data even for moderate size markets despite what current theoretical analysis shows. We note that for large markets, such as the market we study in this work, α is expected to be much smaller than 1. Nevertheless, our simulation results suggest that in practice `OPM` performs well even for markets having a more moderate size and a larger value of α. We also note that the three-sided incentive compatibility of our mechanism implies that it is universally truthful, i.e., truthful for all possible random coin flips.

1.2 Additional Related Work

From a motivational point of view our model is closely related to models involving mediators and online advertising markets, such as the models studied by (Ashlagi, Monderer, and Tennenholtz 2009; Stavrogiannis, Gerding, and Polukarov 2014). However, despite their network exchange motivation, these models are actually auctions (i.e., one-sided mechanisms). Moreover, they focus on offline revenue maximization mechanisms, which is very different from our focus. Other works with a different motivation, such as (McAfee 1992; Gonen, Gonen, and Pavlov 2007; Segal-Halevi, Hassidim, and Aumann 2016), have studied mechanisms for two-sided *non-dynamic* settings. However, with the exception of the very recent last reference, they all considered single-dimensional agents. We are not aware of any previous mechanism for a two-sided dynamic setting.

There is also a significant body of works studying dyanmic matching problems with an adversarial arrival order. This body of work was originated by the work of Karp, Vazirani, and Vazirani (1990) who described an optimal dynamic algorithm for unweighted bipartite online matching. Later works considered more general settings allowing various kinds of weights—see, for example, (Charikar, Henzinger, and Nguyen 2014). We note that none of these works refers to strategic considerations.

2 Model and Definitions

Let us now present the exact details of the model we consider. The model consists of a set P of users, a set M of mediators and a set A of advertisers. Each user $p \in P$ has a non-negative cost $c(p)$ which she suffers if she is assigned to an advertiser; thus, the utility of p is 0 if she is not assigned and $t - c(p)$ if she is assigned and paid t. The users are partitioned among the mediators, and we denote by $P(m) \subseteq P$ the set of users associated with mediator $m \in M$ (i.e., the sets $\{P(m) \mid m \in M\}$ form a disjoint partition of P). The utility of a mediator $m \in M$ is the amount he is paid minus the total cost his users suffer; hence, if $x(p) \in \{0, 1\}$ is an indicator for the event that user $p \in P(m)$ is assigned and t is the payment received by m (part of which might have been forwarded by the mediator to his users), then the utility of m is $t - \sum_{p \in P(m)} x(p) \cdot c(p)$[4]. Finally, each advertiser $a \in A$ has a positive capacity $u(a)$, and she gains a non-negative value $v(a)$ from every one of the first $u(a)$ users assigned to her; thus, if advertiser a is assigned $n \leq u(a)$ users and has to pay t then her utility is $n \cdot v(a) - t$.

As explained in Section 1, we assume the entities (i.e., the mediators and advertisers) arrive at a uniformly random order. A mechanism for this model knows the total number of entities[5], and views the entities as they arrive; however, it has no prior knowledge about the parameters of the entities or about the users. To compensate for this lack of knowledge, each arriving entity reports information to the mechanism. Each advertiser reports her capacity and value. The reports of the mediators are formed in a slightly more involved way. Each user reports her cost to her mediator, and based on these reports each mediator reports the number of his users and their costs to the mechanism. The users, mediators and advertisers are all strategic, and thus, free to produce incorrect reports. In other words, an advertiser may report incorrect capacity and value, a user may report an incorrect cost and a mediator may report a smaller number of users and associate with each one of them an arbitrary cost.

Every time that a new entity arrives, the mechanism has an opportunity to assign users to advertisers. More specifically, when a mediator arrives the mechanism is allowed to assign users of the newly arriving mediator to advertisers that have already arrived. Similarly, when an advertiser arrives the mechanism is allowed to assign users of mediators that have already arrived to the newly arriving advertiser. The objective of the mechanism is to end up with an assignment of users to advertisers maximizing the *gain from trade*. In order to incentivize the mediators and advertisers to report truthfully, the mechanism may charge the advertisers and pay the mediators. Additionally, the mechanism is also allowed to recommend for each mediator how much

[3] A mechanism is *budget balanced* if the amount it charges (from the advertisers) is at least as large as the amount it pays.

[4] The mediators' utility functions are independent of the amount of money transferred from the mediators to the users. This choice was made with the aim of balancing two of the mediators? conflicting objectives: on the one hand, mediators want to make as much money as possible, and on the other hand, they want to acquire users and have them use their services rather than switch to another mediator who is known for paying more money to his users.

[5] In some cases the assumption that the mechanism has a prior knowledge about the number of entities might be considered unnatural. The mechanism we present can be modified using standard techniques to work with an alternative assumption stating that each entity arrives at a uniformly random time from some range (for example, $[0, 1]$). We refer the reader to (Feldman, Naor, and Schwartz 2011) for more details.

of the payment he received to forward to each one of his user. It is important to observe that the utility function of the mediators is not affected by the forwarding of payments to the users, and thus, it is reasonable to believe that mediators follow the forwarding recommendations.

We say that a user is *truthful* if she reports her true cost. Similarly, an advertiser is *truthful* if she reports her true capacity and value. Finally, a mediator is considered *truthful* if he reports to the mechanism his true number of users and the costs of the users as reported to him; and, in addition, he also pays the users according to the recommendation of the mechanism (i.e., he lets them know about their true balance).

We associate a set $B(a)$ of $u(a)$ slots with each advertiser $a \in A$. This allows us to think of the users as assigned to slots instead of directly to advertisers. Formally, let B be the set of all slots (i.e., $B = \bigcup_{a \in A} B(a)$), then an assignment is a set $S \subseteq B \times P$ in which no user or slot appears in more than one ordered pair. We say that an assignment S assigns a user p to slot b if $(p, b) \in S$. Similarly, we say that an assignment S assigns user p to advertiser a if there exists a slot $b \in B(a)$ such that $(p, b) \in S$. It is also useful to define values for the slots. For every slot b of advertiser a, we define its value $v(b)$ as equal to the value $v(a)$ of a. Using this notation, the gain from trade of an assignment S can be stated as: $\texttt{GfT}(S) = \sum_{(p,b) \in S}[v(b) - c(p)]$.

Finally, let us define two additional useful shorthands. Given a set $A' \subseteq A$ of advertisers, we denote by $B(A') = \bigcup_{a \in A'} B(a)$ the set of slots belonging to the advertisers of A'. Similarly, given a set $M' \subseteq M$ of mediators, $P(M') = \bigcup_{m \in M'} P(m)$ is the set of users associated with the mediators of M'.

Comparison of Costs and Values. The presentation of our mechanism is simpler when the values of slots and the costs of users are all unique. Clearly, this is extremely unrealistic since all the slots of a given advertiser have the exact same value in our model. Thus, we simulate uniqueness using a tie-breaking rule (which must be independent of the reports of the agents). In the rest of this paper, whenever costs/values are compared, the comparison is assumed to use such a tie breaking rule.

Canonical Assignment. Given a set $B' \subseteq B$ of users and a set $P' \subseteq P$ of slots, the canonical assignment $S_c(P', B')$ is the assignment constructed as follows. First, we order the slots of B' in a decreasing value order $b_1, b_2, \ldots, b_{|B'|}$ and the users of P' in an increasing cost order $p_1, p_2, \ldots, p_{|P'|}$. Then, for every $1 \leq i \leq \min\{|B'|, |P'|\}$ the canonical assignment $S_c(B', P')$ assigns user p_i to slot b_i if and only if $v(b_i) > c(p_i)$.

The canonical assignment is an important tool we use often in this paper, and it was proved by (Feldman and Gonen 2016) that $S_c(P', B')$ is always an assignment of users from P' to slots of B' maximizing the gain from trade (among all such assignments). Occasionally, we refer to the user or slot at location i of a canonical assignment $S_c(P', B')$, by which we mean user p_i or slot b_i, respectively.

3 Our Mechanism

In this section we describe and analyze our dynamic mechanism "Observe and Price Mechanism" (`OPM`). `OPM` assumes $|S_c(P, B)| > 0$, and that there exists a value $\alpha \in [|S_c(P, B)|^{-1}, 1]$, known to the mechanism, such that we are guaranteed that, for every advertiser $a \in A$ and mediator $m \in M$:

$$\frac{u(a)}{|S_c(P, B)|} \leq \alpha \quad \text{and} \quad \frac{|P(m)|}{|S_c(P, B)|} \leq \alpha \ .$$

In other words, α is an upper bound on how large can the capacity of an advertiser or the number of users of a mediator be compared to the size of the optimal assignment $S_c(P, B)$. We remind the reader that α can be informally understood as a bound on the market importance of any single entity.

A description of `OPM` is given as Mechanism 1. Notice that Mechanism 1 accepts a parameter $r \in (0, 1/2]$ whose value is specified later. Additionally, Mechanism 1 often refers to parameters of the model that are not known to the mechanism, such as the value of an advertiser or the number of users of a mediator. Whenever this happens, this should be understood as referring to the reported values of these parameters.

Mechanism 1: Observe and Price Mechanism (`OPM`)

1. Draw a random value t from the binomial distribution $\mathcal{B}(|A| + |M|, r)$, and observe the first t entities that arrive without assigning any users. Let A_T and M_T be the set of the observed advertisers and mediators, respectively. We later refer to this step of the mechanism as the "observation phase".

2. Let $\hat{p}$ and $\hat{b}$ be the user and slot, respectively, at location $\lceil (1 - 2r^{-1} \cdot \sqrt[3]{\alpha}) \cdot |S_c(P(M_T), B(A_T))| \rceil$ of the canonical assignment $S_c(P(M_T), B(A_T))$. If $(1 - 2r^{-1} \cdot \sqrt[3]{\alpha}) \cdot |S_c(P(M_T), B(A_T))| \leq 0$, then the previous definition of $\hat{p}$ and $\hat{b}$ cannot be used. Instead, define $\hat{p}$ as a dummy user of cost $-\infty$ and $\hat{b}$ as a dummy slot of value ∞. We say that a slot b or a user p corresponding to an entity that arrived *after* the observation phase is *assignable* if $v(b) > v(\hat{b})$ or $c(p) < c(\hat{p})$, respectively.

3. Let σ_E be the sequence of the entities that arrived so far after the observation phase. Initially σ_E is empty, and entities are added to it as they arrive.

4. For every arriving entity:
   a. Add the new entity to the end of σ_E.
   b. If the arriving entity is a mediator m (advertiser a), then, as long as m (a) has unassigned assignable users (slots) and there is an advertiser (mediator) in σ_E having unassigned assignable slots (users), do:
   - Let a (m) be the earliest advertiser (mediator) in σ_E having unassigned assignable slots (users).
   - Assign the unassigned assignable user of mediator m with the lowest cost to an arbitrary unassigned assignable slot of a, charge an amount of $v(\hat{b})$ from advertiser a and pay $c(\hat{p})$ to mediator m.

   c. For every mediator $m \in \sigma_E$, recommend m to transfer his assigned users an additional amount that guarantees the following:

- If all the assignable users of m are assigned, the additional amount should increase the total payment received so far by each assigned user of m to $c(\hat{p})$.
- Otherwise, let p be the unassigned assignable user of m with the minimum cost. In this case the additional amount should increase the total payment received so far by each assigned user of m to $c(p)$.[6]

We would like to note that OPM is based on a mechanism of (Feldman and Gonen 2016) named "Threshold by Partition Mechanism", and the analyses of both mechanisms go along similar lines. However, OPM introduces additional ideas that allow it to work in a dynamic setting. In particular, OPM uses an involved recommended payments updating rule that keeps it three-sided continuously individually rational. Moreover, OPM is able to use an observation phase whose size is a small fraction of the entire input (for $\alpha \ll 1$), whereas the analysis of the mechanism of (Feldman and Gonen 2016) relies on the symmetry properties induced by an even partition of the input (which is inappropriate in a dynamic setting).

Let us start the analysis of OPM with the following simple observation, showing that OPM obeys the restriction of our model on the way a mechanism may update its assignment.

Observation 3.1. *Each time* OPM *assigns a user to a slot, either the user belongs to the newly arrived mediator or the slot belongs to the newly arrived advertiser.*

In the rest of this section we prove the following restatement of Theorem 1.1, which implies the original statement of the theorem from Section 1.1. Due to space constraints, the part of the proof proving the competitive ratio guaranteed by Theorem 1.1 and the proofs of some of the following observations are deferred to the full version of this paper.

Theorem 1.1. OPM *is budget balanced, three-sided continuously individually rational, three-sided incentive compatible and* $(1 - r - 22r^{-1} \cdot \sqrt[3]{\alpha} - 10e^{-2/\sqrt[3]{\alpha}})$*-competitive. Hence, for* $r = \min\{1/2, 4\sqrt[6]{\alpha}\}$ *the competitive ratio of* OPM *is at least:* $1 - 9.5\sqrt[6]{\alpha} - 10e^{-2/\sqrt[3]{\alpha}}$.

One part of Theorem 1.1 is proved by the following observation.

Observation 3.2. OPM *is budget balanced.*

We prove the observation by showing that whenever OPM assigns a user p to a slot b, it charges the advertiser of b more than it pays the mediator of p. We omit the details of the proof due to space constants.

Following is a useful observation about OPM that we occasionally use in the next proofs.

Observation 3.3. OPM *preserves the invariant that one of the following is always true immediately after* OPM *processes the arrival of an entity:*

1. OPM *assigned all the assignable users of mediators that have already arrived.*
2. OPM *assigned users to all the assignable slots of advertisers that have already arrived.*

We are now ready to prove the incentive parts of Theorem 1.1. Specifically, we prove three lemmata showing that OPM is three-sided continuously individually rational and three-sided incentive compatible. The first lemma analyzes the incentive properties of OPM for users.

Lemma 3.4. *For every user* p *whose mediator* m *is truthful,* OPM *is continuously individually rational for* p*, and truthfulness is a dominant strategy for her.*

Proof. If m arrives during the observation phase (i.e., $m \in M_T$), then no user of m is ever assigned to a slot or paid; which makes the lemma trivial. Thus, we assume in the rest of the proof that m arrives after the observation phase.

Note that OPM calculates the threshold $c(\hat{p})$ based on the reports of advertisers and mediators in A_T and M_T, respectively. Thus, p, who is associated with a mediator not belonging to M_T, cannot affect this threshold. Next, let us denote by k the number of users of m that are assigned to slots when p reports a cost smaller than $c(\hat{p})$. We claim that k is independent of the exact cost reported by p, as long as this cost is smaller than $c(\hat{p})$. The reason for that is that most of the time OPM accesses the reported cost of p only by checking whether p is assignable, and the answer for that check does not change as long as the reported cost of p is smaller than $c(\hat{p})$. The exact value of $c(p)$ is only used by OPM after OPM decides to assign *some* user of m to a slot, and then this exact value is used to decide which user of m will be assigned to the slot—which does not affect k.

Let p' be the user of m with the k^{th} smallest cost among his users that are not p. If m does not have k users other than p, then p' is a dummy user of cost ∞. In the next two paragraphs we show that p is assigned to a slot if and only if she reports a cost smaller than $\min\{c(\hat{p}), c(p')\}$. Moreover, when p is assigned to a slot the total payment she gets is this minimum (which is her critical value). Clearly, the incentive compatibility of OPM for p follows from this claim.

Let us begin proving the above claim by showing that p is left unassigned when she reports a cost larger than $\min\{c(\hat{p}), c(p')\}$. There are two cases to consider. If p reports a cost larger than $c(\hat{p})$ then she is not assignable, and thus, she is left unassigned. On the other hand, consider the case that p reports a value smaller than $c(\hat{p})$, but larger than $c(p')$. In this case p is not one of the k users of m with the smallest reported costs, and thus, is again left unassigned.

Next, we prove the other side of the above claim, i.e., that when p reports a cost smaller than $\min\{c(\hat{p}), c(p')\}$ she is assigned and the total payment she gets is this minimum. The fact that p reports a value smaller than $c(\hat{p})$ implies that p is assignable, and the fact that she reports a value smaller than $c(p')$ guarantees that p is one of the k users of m with the smallest reported costs. This already guarantees that p is assigned to some slot, and that p' is the unassigned user of m with the smallest cost when OPM updates for the last time the recommended total payment from m to p (unless

[6]Note that at every point in time m is budget balanced since he receives a payment of $c(\hat{p})$ for each one of his assigned users, and the total amount recommended for him to pay to each one of these users is either $c(\hat{p})$ or equal to the cost of some assignable user (and thus, is upper bounded by $c(\hat{p})$).

p' is a dummy user). Hence, the recommended total payment for p is determined as follows. If p' is assignable (i.e., $c(p') < c(\hat{p})$), then the recommended total payment for p is set to $c(p')$. Otherwise, m has no unassigned assignable users left, and thus, the recommended total payment to p is set to $c(\hat{p}) < c(p')$.

It remains to prove that OPM is continuously individually rational for p, i.e., that the utility of p can only increase over time when p is truthful. The first time that the utility of p might change is when p is assigned. When this happens p is immediately payed an amount equal either to $c(\hat{p})$ or to the cost of an unassigned assignable user of m. Since OPM chooses the user to assign as the unassigned assignable user of m with the lowest cost, both possible payments are larger than $c(p)$, and thus, the utility of p does not become negative following its assignment. Next, we prove that the recommended total payment for p can only increase over time, which proves that p's utility can only increase from the moment p is assigned. To see why that is true, recall that, at every time point in which OPM updates the recommended total payment to p, this total payment is updated to be either the cost of the unassigned assignable user of m with the lowest cost, or $c(\hat{p})$ if m has no unassigned assignable users left. As long as m has unassigned assignable users this update rule yields a recommended total payment which can only increase over time since users occasionally get removed from the set of unassigned assignable users of m (when they get assigned), but no user is ever added to this set. Moreover, the recommended total payment to p also increases when the last unassigned assignable user of m gets assigned since the recommended total payment before this point was equal to the cost of some assignable user of m, which is smaller, by definition, than the new recommended total payment $c(\hat{p})$. □

The next lemma analyzes the incentive properties of OPM for mediators.

Lemma 3.5. *For every mediator m whose users are truthful, OPM is continuously individually rational for m, and truthfulness is a dominant strategy for him.*

Proof. If m arrives during the observation phase (i.e., $m \in M_T$), then no user of m is ever assigned to a slot and m receives no payment. Hence, the lemma is trivial in this case. Thus, we assume in the rest of the proof that m arrives after the observation phase.

Note that OPM calculates the threshold $c(\hat{p})$ based on the reports of advertisers and mediators in A_T and M_T, respectively. Thus, m, who does not belong to M_T, cannot affect this threshold. Whenever a user $p \in P(m)$ is assigned to a slot the utility of m (and the user) decreases by $c(p)$ and increases by the additional payment m gets, which is $c(\hat{p})$. In other words, the utility of m changes by $c(\hat{p}) - c(p)$ (independently of the amount m forwards to p). When m is truthful this change is always non-negative since the assignment of p implies that she is assignable, i.e., her reported cost is smaller than $c(\hat{p})$. This already proves that each assignment of a user of m increases his utility by a non-negative amount when he is truthful (assuming his users are also truthful), thus, OPM is continuously individually rational for m.

Let s be the number of assignable users of m, according to his report. We claim that there exists a value k which is independent of the report of m such that for any report of m the mechanism assigns the $\min\{k, s\}$ users of m with the lowest reported costs. Before proving this claim, let us explain why the lemma follows from this claim. The above description shows that the utility of m changes by a $c(\hat{p}) - c(p)$ for every assigned user $p \in P(m)$, thus, m wishes to assign as many as possible users having cost less than $c(\hat{p})$, and if he cannot assign all of them then he prefers to assign the users with the lowest costs. By being truthful m guarantees that only users of cost less than $c(\hat{p})$ are considered assignable, and thus, have a chance to be assigned. Moreover, by the above claim OPM assigns the k assignable users of m with the lowest costs (or all of them if $s < k$), which is the best result m can hope for given that at most k of his users can be assigned. Hence, truthfulness is a dominant strategy for m.

We are only left to prove the above claim. Note that Observation 3.3 implies that OPM assigns no users of m as long as there are mediators appearing earlier in σ_E which still have unassigned assignable users. Once there are no more such mediators, OPM assigns users of m, in an increasing costs order, to unassigned assignable slots till one of two things happens: either m runs out of unassigned assignable users, or the input for OPM ends. This means that when the input for OPM ends before all the assignable users of mediators appearing before m in σ_E are assigned, then no users of m are assigned and the claim holds with $k = 0$. Otherwise, we choose k to be the number of unassigned assignable slots immediately before OPM assigns the first user of m (we count in k both unassigned assignable slots of advertisers that have already arrived at this moment and unassigned assignable slots of advertisers that arrive later). Notice that the report of m does not affect the behavior of OPM up to the moment it starts assigning users of m, thus, k is independent of the report of m. If $s > k$, then the k users of m with the lowest costs are assigned before OPM runs out of input and stops. Otherwise, if $s \leq k$ then OPM stops assigning users of m only after assigning all the s assignable users of m. □

Finally, the next lemma considers the incentive properties of OPM for advertisers. The proof of this lemma is analogous to the proof of the previous lemma (with slots exchanging roles with users, $v(\hat{b})$ exchanging roles with $c(\hat{p})$, etc.), and thus, we omit it.

Lemma 3.6. *For every advertiser a, OPM is continuously individually rational for a, and truthfulness is a dominant strategy for her.*

4 Simulations

We have used simulations to study the empirical performance of our mechanism OPM. Our simulations involved two methods for generating the input. The more interesting of these methods, which we call *real-data based input* was as follows. The creation of the advertisers was based on data collected as part of a Horizon 2020 project from Facebook campaigns targeting Europeans between the ages 18 and 22 who are interested in entertainment. Every bid collected consisted of a budget for the relevant campaign, the maximal CPC (cost-per-click) value, the minimal CPC value and the median CPC value that the advertiser was willing to pay. Based on these bids we constructed three advertisers for our generated input, one advertiser for each one of the CPC values. More specifically, let β be the budget specified by the bid, and let δ be one of the three CPC values specified by this bid, then the advertiser created for this CPC value has a value of δ and a capacity of $\frac{\beta}{100\delta}$.[7] For every advertiser we also created a single mediator having the same number of users as the capacity of the advertiser. Every user of these mediators was assigned an independent cost chosen uniformly at random from the range between the smallest and largest CPC values encountered in the real world bids. To verify that the data fed into the mechanism was unbiased, we designed a secondary input generation method which we call the *random bids input*. The way input is generated by this method is very similar to the way real-data based input was generated, except that the advertisers' values were selected as uniformly random independent values between the smallest and largest CPC values encountered in the real world bids (rather then being taken directly from the input bids, as in the real-data based input).

Our first simulation was designed to study the effect of market size on the performance of the mechanism. In this experiment we used the above methods to generate markets of various sizes and then we sent the entities of each generated market into `OPM` in a uniformly random order. The observable performance of the assignments produced by `OPM` (as a percent of the optimal canonical assignment) are depicted in Figure 1. In order to reduce variance and error margins, every value given by this figure (and the next one) was produced by averaging 3000 independent executions. As expected, the performance of the algorithm improves with the size of the market (as the size of the market is roughly inversely proportional to α). Moreover, one can observe that the performance of the mechanism is quite good even for moderate size markets, which is better than what can be predicted based on our theoretical result alone.

In the previous experiment, we used the value of the parameter r of `OPM` which was specified by the version of Theorem 1.1 given in Section 3. Our second simulation was designed to study the possibility of improving the performance of the mechanism by varying the value of r. Specifically, we repeated the previous experiment with a market of 11961 advertisers (which is close to the size of the largest market we considered before), but varied the value of the parameter r. The results of this experiment are depicted in Figure 2. As in the previous experiment, we see again that using the value of r specified by Theorem 1.1 ($1/2$ in this case) leads to good performance for both input generation methods. For the real-data based input, varying r does not improve the performance of `OPM`, but for the random input bids one can significantly improve the outcome by decreasing r. This is likely a consequence of a higher variance in the random bids input, which allows `OPM` to calculate good thresholds based on a shorter observation phase (which are induced by decreasing r). Thus, for the random bids input, decreasing r leads to improved performance as it allows `OPM` to harvest value from a larger fraction of the market while inducing only a weak adversarial effect on the selected thresholds.

[7] Our experiments are based on only a fraction of the entire data set, which significantly increased the market strength of the entities in the input. To compensate for this increase, and keep the market strength of each advertiser in the simulation similar to the market strength of the corresponding real world advertiser, we introduced a division by 100 into the capacity formula.

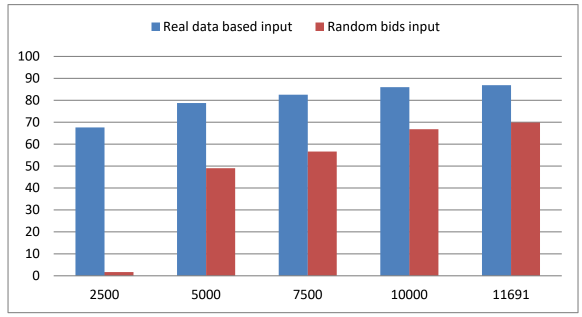

Figure 1: The performance of `OPM` as a function of the market size. The number below each column specifies the number of advertisers in the market.

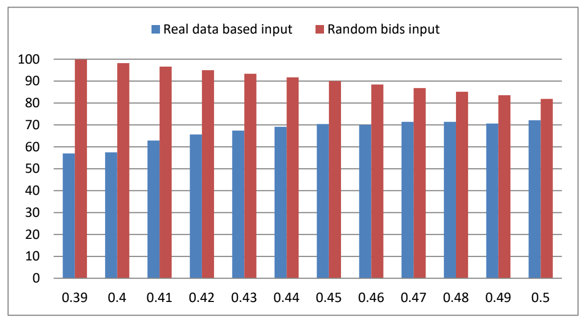

Figure 2: The performance of `OPM` as a function of the value of the parameter r on markets with 11961 advertisers.

5 Conclusion

In this paper we have presented a dynamic model for a foreseeable form of the online advertising market based on principals suggested by (Feldman and Gonen 2016), and described a mechanism called `OPM` for it. `OPM` is the first mechanism for a multi-sided market that guarantees the economic properties of budget balance, incentive compatibility and individual rationality while having a non-trivial theoretical approximation guarantee. For large markets, such as the online advertising market, the theoretical competitive ratio of `OPM` approaches 1. However, this theoretical guarantee becomes much weaker (or even not relevant) for smaller markets, and thus, we have complimented it with simulation results. These results suggest that `OPM` performs well in practice even for markets of moderate size.

References

Ashlagi, I.; Monderer, D.; and Tennenholtz, M. 2009. Mediators in position auctions. *Games and Economic Behavior* 67:2–21.

Babaioff, M.; Dinitz, M.; Gupta, A.; Immorlica, N.; and Talwar, K. 2009. Secretary problems: weights and discounts. In *SODA*, 1245–1254.

Babaioff, M.; Feldman, M.; and Tennenholtz, M. 2016. Mechanism design with strategic mediators. *ACM Transactions on Economics and Computation* 4:7:1–1:48.

Blum, A.; Sandholm, T.; and Zinkevich, M. 2002. Online algorithms for market clearing. In *SODA*, 971–980.

Bredin, J.; Parkes, D.; and Duong, Q. 2007. Chain: A dynamic double auction framework for matching patient agents. *Journal of Artificial Intelligence Research* 30:133–179.

Charikar, M.; Henzinger, M.; and Nguyen, H. L. 2014. Online bipartite matching with decomposable weights. In *ESA*, 260–271.

Conitzer, V.; Taylor, C. R.; and Wagman, L. 2012. Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science* 31(2):277–292.

Feldman, M., and Gonen, R. 2016. Markets with strategic multi-minded mediators. *CoRR* abs/1603.08717.

Feldman, M.; Naor, J. S.; and Schwartz, R. 2011. Improved competitive ratios for submodular secretary problems. In *APPROX*, 218–229.

Gonen, M.; Gonen, R.; and Pavlov, E. 2007. Generalized trade reduction mechanisms. In *EC*, 20–29.

Karp, R. M.; Vazirani, U. V.; and Vazirani, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *STOC*, 352–358.

McAfee, R. P. 1992. dominant strategy double auction. *Journal of Economic Theory* 56:434–450.

Myerson, R. B., and Satterthwaite, M. A. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29:265–281.

Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2016. Mida: A multi item-type double-auction mechanism. *CoRR* abs/1604.06210.

Stavrogiannis, L. C.; Gerding, E. H.; and Polukarov, M. 2014. Auction mechanisms for demand-side intermediaries in online advertising exchanges. In *AAMAS*, 5–9.

Vaze, R., and Coupechoux, M. 2016. Online budgeted truthful matching. *CoRR* abs/1606.00581.

Wurman, P.; Walsh, W.; and Wellman, M. 1998. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems* 24:17–27.