



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

Moran Feldman<sup>1</sup>(✉), Gonen Frim<sup>1</sup>, and Rica Gonen<sup>2</sup>

<sup>1</sup> Department of Mathematics and Computer Science, The Open University of Israel,  
Ra'anana, Israel

{moranfe,gonenr}@openu.ac.il

<sup>2</sup> Department of Management and Economics, The Open University of Israel,  
Ra'anana, Israel  
gonen987@gmail.com

**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 (2018), we introduce a new online advertising market model which uses information brokers to give users such control. Unlike Feldman and Gonen (2018), 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 regimes for which our theoretical guarantee is weak or irrelevant.

**Keywords:** Dynamic mechanisms · Mutli-sided markets  
Online advertising market

## 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 [7]. Such concerns were actively raised by EU regulators in recent years in efforts to find solutions to guarantee users' privacy. Recently privacy concerns have also reached the U.S. Senate and Congress as a response to Facebook's information leak to Cambridge

Analytica. It was evident in Facebook’s hearing before the U.S. Senate, particularly in Senator Schatz’s line of questioning [8], that Facebook is expected to develop tools to enable end users to configure their privacy settings and that the notion of a data fiduciary was put forward to apply pressure to Facebook in this area.

Based on this motivation, and extending principles suggested by [10], 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.<sup>1</sup> 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 [10], 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 in 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. From a practical point of view, an assignment of a user  $p$  to an advertiser  $a$  means that when  $p$  views interstitial advertising, it will be from advertiser  $a$ . Given this meaning for an assignment, it is natural to assume 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 [2, 19]. 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 man-

---

<sup>1</sup> The use of mediators is necessary because it should not be possible to link an information portfolio offered for sell on the market to any particular user, which prevents users from interacting directly with the market.

aged to achieve these goals simultaneously. Wurman et al. [20] presented a mechanism incentivizing truthful reporting from either the buyers or the sellers, but not simultaneously from both. A different mechanism given by Blum et al. [4] maximizes the social welfare of buyers and non-selling sellers (as opposed to maximizing the gain from trade).<sup>2</sup> Finally, Bredin et al. [5] 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 [16]. 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 incentivize agents to report truthful information to the mechanism. Additionally, 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 all the agents arrive.

## 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 loses by participating. However, in a dynamic setting it is natural to require also that an agent never loses 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)

---

<sup>2</sup> Blum et al. [4] has multiple different objectives (maximizing profit, liquidity and welfare). We refer to the objective which is most relevant to our work.

if the agent’s utility can only increase over time when the agent is truthful.<sup>3</sup> 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 [3], we say that a mechanism is *user-side incentive compatible* if truthfulness is a dominant strategy<sup>4</sup> 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  $\alpha$  is an upper bound, known to the mechanism, on the market importance of any single agent. Formally,  $\alpha$  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.

---

<sup>3</sup> 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 Sect. 2.

<sup>4</sup> Here and throughout the paper, a reference to domination of strategies should always be understood as a reference to weak domination.

**Theorem 1.** *OPM is budget balanced<sup>5</sup>, 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 the competitive ratio guaranteed by Theorem 1 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 latter 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 the current theoretical analysis shows. We note that for large markets, such as the market we study in this work,  $\alpha$  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  $\alpha$ . In addition, in order to demonstrate the need for an involved solution such as OPM, we compare OPM’s practical performance to that of a straw-man mechanism. This comparison demonstrates that our mechanism performs significantly better than the straw-man mechanism even for moderate size markets. 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 [1, 18]. 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 [13, 15, 17], 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 dynamic matching problems with an adversarial arrival order. This body of work was originated by the work of [14] 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, [6]. We note that none of these works refers to strategic considerations.

---

<sup>5</sup> A mechanism is *budget balanced* if the amount it charges (from the advertisers) is at least as large as the amount it pays.

## 2 Model and Definitions

Let us now present the exact details of the model we consider. This 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  $t - \sum_{p \in P(m)} x(p) \cdot c(p)$ .<sup>6</sup> 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 Sect. 1, we assume the entities (i.e., the mediators and advertisers) arrive in a uniformly random order. A mechanism for this model knows the total number of entities,<sup>7</sup> 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 arriv-

---

<sup>6</sup> 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.

<sup>7</sup> 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,  $[0, 1]$ ). See [11] for more details.

ing 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, it is also allowed to recommend for each mediator how much of the payment he received to forward to each one of his user. It is important to observe that, since the utility functions of the mediators are not affected by the forwarding of payments to the users, 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  $S \subseteq P \times B$  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:  $\text{GfT}(S) = \sum_{(p,b) \in S} [v(b) - c(p)]$ .

Finally, we define two additional useful shorthands. Given a set  $A' \subseteq A$  of advertisers, let  $B(A') = \bigcup_{a \in A'} B(a)$  be 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  $b_1, b_2, \dots, b_{|B'|}$  and the users of  $P'$  in an increasing cost order  $p_1, p_2, \dots, p_{|P'|}$ . Then, for every  $1 \leq i \leq \min\{|B'|, |P'|\}$ ,  $S_c(B', P')$  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 [10] 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  $m \in M: \max\{u(a), |P(m)|\} \leq \alpha \cdot |S_c(P, B)|$ . In other words,  $\alpha$  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  $\alpha$  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  $[(1 - 2r^{-1} \cdot \sqrt[3]{\alpha}) \cdot |S_c(P(M_T), B(A_T))|]$  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  $\infty$ . 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  $\sigma_E$  be the sequence of the entities that arrived so far after the observation phase. Initially  $\sigma_E$  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  $\sigma_E$ .
  - 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  $\sigma_E$  having unassigned assignable slots (users), do:
    - Let  $a$  ( $m$ ) be the first advertiser (mediator) in  $\sigma_E$  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  $c(p)$ .<sup>a</sup>

---

<sup>a</sup> 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  $c(\hat{p})$ ).

---

We would like to note that OPM is based on a mechanism of [10] 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 [10] 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 2.** *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.*

At this point we would like to prove the following restatement of Theorem 1.

**Theorem 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  $r = \min\{1/2, 4\sqrt[3]{\alpha}\}$  the competitive ratio of OPM is at least:  $1 - 9.5\sqrt[3]{\alpha} - 10e^{-2/\sqrt[3]{\alpha}}$ .*

One part of Theorem 1 is proved by the following observation.

**Observation 3.** *OPM is budget balanced.*

*Proof.* 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  $p$ .

Consider an arbitrary ordered pair  $(p, b)$  from the assignment produced by OPM. The fact that  $p$  is assigned implies that  $c(p) < c(\hat{p})$ , and thus,  $\hat{p}$  is not a dummy user (since  $c(\hat{p}) = -\infty$  when  $\hat{p}$  is a dummy user). Similarly, the fact that a user is assigned to  $b$  implies that  $v(b) > v(\hat{b})$ , and thus,  $\hat{b}$  is not a dummy slot (since  $v(\hat{b}) = \infty$  when  $\hat{b}$  is a dummy slot).

Recall that the fact that  $\hat{p}$  and  $\hat{b}$  are not dummy user and slot, respectively, implies that  $\hat{p}$  and  $\hat{b}$  are matched by the canonical assignment  $S_c(P(M_T), B(A_T))$ . Since a canonical assignment never assigns a user  $p'$  to a slot  $b'$  when  $c(p') > v(b')$ , we get  $c(\hat{p}) < v(\hat{b})$ . The proof now completes by observing that the advertiser of  $b$  is charged  $v(\hat{b})$  for the assignment of  $p$  to  $b$ , and the mediator of  $p$  is paid only  $c(\hat{p})$  for this assignment.  $\square$

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

**Observation 4.** *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.*

*Proof.* Clearly the invariant holds during the observation phase because only mediators and advertisers that arrive after the observation phase contribute assignable users and slots, respectively. Next, assume the invariant held before the arrival of some mediator  $m$  which arrives after the observation phase, and let us prove that it holds also after the arrival of  $m$ . If before the arrival of  $m$  case (2) of the invariant held, then this case also holds after the arrival of  $m$  since  $m$  contributes no new slots. On the other hand, if case (1) held before the arrival of  $m$ , then OPM assigns the assignable users of  $m$  to assignable slots of advertisers that have already arrived till one of two things happen: either all the assignable slots of advertisers that have already arrived get assigned (and thus, case (2) of the invariant now holds), or all the assignable users of  $m$  get assigned (and thus, case (1) of the invariant holds again). It remains to prove that if the invariant held before the arrival of an advertiser  $a$  which arrives after the observation phase, then it also holds after her arrival. However, this proof is analogous to the above proof for mediators, and thus, we omit it.  $\square$

### 3.1 The Incentive Properties of OPM

In this section we prove the incentive parts of Theorem 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. Due to space constraints, we defer the proof of this lemma to the full version of this paper.

**Lemma 1.** *For every user  $p$ , assuming the mediator  $m$  of  $p$  is truthful, OPM is continuously individually rational for  $p$ , and truthfulness is a dominant strategy for her.*

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

**Lemma 2.** *For every mediator  $m$ , assuming the users of  $m$  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 it. 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 4 implies that OPM assigns no users of  $m$  as long as there are mediators appearing earlier in  $\sigma_E$  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  $\sigma_E$  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  $m$ .  $\square$

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

### 3.2 The Competitive Ratio of OPM

In this section we analyze the competitive ratio of OPM. Throughout the section we use the letter  $\tau$  as a shorthand for  $|S_c(P, B)|$ . We also define  $\tilde{P}$  ( $\tilde{B}$ ) as the set of the users (slots) at locations 1 to  $\lceil(1 - 6r^{-1} \cdot \sqrt[3]{\alpha})\tau\rceil$  of the canonical assignment  $S_c(P, B)$  ( $\tilde{P}$  and  $\tilde{B}$  are defined to be empty when  $1 - 6r^{-1} \cdot \sqrt[3]{\alpha} \leq 0$ ). The following observation shows that most of the gain from trade of the canonical assignment  $S_c(P, B)$  comes from the users and slots of  $\tilde{P}$  and  $\tilde{B}$ , respectively. For convenience, let us denote by  $P_o$  the set of users that are assigned by  $S_c(P, B)$ , and by  $B_o$  the set of slots that are assigned some user by  $S_c(P, B)$ .

**Observation 5.**  $\sum_{b \in \tilde{B}} v(b) - \sum_{p \in \tilde{P}} c(p) \geq (1 - 6r^{-1} \cdot \sqrt[3]{\alpha}) \cdot \text{GfT}(S_c(P, B))$ .

*Proof.* If  $1 - 6r^{-1} \cdot \sqrt[3]{\alpha} \leq 0$ , then both  $\tilde{B}$  and  $\tilde{P}$  are empty, and the inequality that we need to prove holds since its left hand side is 0 and its right hand side is non-positive (recall that  $S_c(P, B)$  is an assignment of users from  $P$  to slots of  $B$  maximizing the gain from trade, and thus, its gain from trade is at least 0 since  $\text{GfT}(\emptyset) = 0$ ). Thus, we may assume in the rest of the proof that  $1 - 6r^{-1} \cdot \sqrt[3]{\alpha} > 0$ .

Since  $\tilde{B}$  contains the  $\lceil(1 - 6r^{-1} \cdot \sqrt[3]{\alpha})\tau\rceil$  slots with the largest values among the slots of  $B_o$ , we get:

$$\sum_{b \in \tilde{B}} v(b) \geq \lceil(1 - 6r^{-1} \cdot \sqrt[3]{\alpha})\tau\rceil \cdot \frac{\sum_{b \in B_o} v(b)}{\tau}.$$

Similarly, since  $\tilde{P}$  contains the  $\lceil(1 - 6r^{-1} \cdot \sqrt[3]{\alpha})\tau\rceil$  users with the lowest costs among the users of  $P_o$ , we get:

$$\sum_{p \in \tilde{P}} c(p) \leq \lceil(1 - 6r^{-1} \cdot \sqrt[3]{\alpha})\tau\rceil \cdot \frac{\sum_{c \in P_o} c(p)}{\tau}.$$

Combining the two inequities gives:

$$\begin{aligned} & \sum_{b \in \tilde{B}} v(b) - \sum_{p \in \tilde{P}} c(p) \\ & \geq [(1 - 6r^{-1} \cdot \sqrt[3]{\alpha})_\tau] \cdot \frac{\sum_{b \in B_o} v(b) - \sum_{p \in P_o} c(p)}{\tau} \\ & = [(1 - 6r^{-1} \cdot \sqrt[3]{\alpha})_\tau] \cdot \frac{\text{GfT}(S_c(P, B))}{\tau} \\ & \geq (1 - 6r^{-1} \cdot \sqrt[3]{\alpha}) \cdot \text{GfT}(S_c(P, B)). \quad \square \end{aligned}$$

Observation 5 shows that one can prove a competitive ratio for OPM by relating the gain from trade of the assignment it produces to the gain from trade obtained by assigning the users of  $\tilde{P}$  to the slots  $\tilde{B}$ . The following lemma is a key lemma we use to relate the two gains. In order to state this lemma we need some additional definitions. Consider the following two sets.

$$\hat{P} = \{p \in P(M \setminus M_T) \mid c(p) < c(\hat{p})\}$$

and

$$\hat{B} = \{b \in B(A \setminus A_T) \mid v(b) > v(\hat{b})\}.$$

Intuitively,  $\hat{P}$  is the set of the assignable users, and  $\hat{B}$  is the set of the assignable slots. It is important to note that  $\hat{P}$  and  $\hat{B}$  are both empty whenever  $\hat{p}$  and  $\hat{b}$  are dummy user and slot, respectively. We also define two additional sets  $A_L$  and  $M_L$  as follows. Let  $f$  be a random variable distributed according to the binomial distribution  $\mathcal{B}(|A \setminus A_T| + |M \setminus M_T|, \min\{16r^{-1} \cdot \sqrt[3]{\alpha}, 1\})$ , and let  $L$  be the set of the last  $f$  entities in  $\sigma_E$  (or equivalently, the last  $f$  entities to arrive). The sets  $A_L$  and  $M_L$  are then defined as  $A_L = A \cap L$  and  $M_L = M \cap L$ .

**Lemma 4.** *There exists an event  $\mathcal{E}$  of probability at least  $1 - 10e^{-2/\sqrt[3]{\alpha}}$  such that  $\mathcal{E}$  implies the following:*

- (i)  $\tilde{B} \setminus B(A_T) \subseteq \hat{B}$
- (ii)  $\tilde{P} \setminus P(M_T) \subseteq \hat{P}$
- (iii)  $|\hat{P} \setminus P(M_L)| \leq |\hat{B}|$
- (iv)  $|\hat{B} \setminus B(A_L)| \leq |\hat{P}|$
- (v)  $c(p) \leq \ell(P, B) \leq v(b)$  for every user  $p \in \hat{P}$  and slot  $b \in \hat{B}$ , where  $\ell(P, B)$  is a value which is independent of the random coins of OPM and obeys  $c(p) \leq \ell(P, B) \leq v(b)$  for every  $p \in P_o$  and  $b \in B_o$ .

The proof of Lemma 4 is very similar to the proof of Lemma 4.6 in [9] (which is a version of [10] with full proofs), and thus, we omit it. In the rest of this section we explain how the competitive ratio of OPM follows from Lemma 4. Let  $\hat{S}$  be the assignment produced by OPM.

**Lemma 5.** *The event  $\mathcal{E}$  implies the following inequality:*

$$\text{GfT}(\hat{S}) \geq \sum_{\substack{b \in \hat{B} \\ b \notin B(A_T \cup A_L)}} [v(b) - \ell(P, B)] + \sum_{\substack{p \in \hat{P} \\ p \notin P(M_T \cup M_L)}} [\ell(P, B) - c(p)].$$

*Proof.* Lemma 4 shows that given  $\mathcal{E}$  we have  $|\hat{P} \setminus P(M_L)| \leq |\hat{B}|$ , hence, Observation 4 implies that OPM assigns at least  $|\hat{P} \setminus P(M_L)|$  users. Additionally, since OPM assigns users of mediators from  $M_L$  only after all the assignable users of mediators from  $M \setminus (M_T \cup M_L)$  are assigned to slots we get that all the users of  $\hat{P} \setminus P(M_L)$  are assigned by  $\hat{S}$  given  $\mathcal{E}$ . On the other hand, Lemma 4 also shows that given  $\mathcal{E}$  all the users of  $\hat{P} \setminus P(M_T)$  belong to  $\hat{P}$ , and thus, the users of  $\hat{P} \setminus P(M_T \cup M_L)$  are all assigned by  $\hat{S}$ . A similar argument shows that the slots of  $\hat{B} \setminus B(A_T \cup A_L)$  are all assigned users by  $\hat{S}$  given  $\mathcal{E}$ . Finally, observe that  $\mathcal{E}$  also implies that  $c(p) \leq \ell(P, B) \leq v(b)$  for every pair  $(p, b) \in \hat{S} \subseteq \hat{P} \times \hat{B}$ .

In the rest of the proof we assume that  $\mathcal{E}$  happens. Consider an ordered pair  $(p, b) \in \hat{S}$ . Then, the contribution of  $(p, b)$  to  $\mathbf{GfT}(\hat{S})$  is:

$$v(b) - c(p) = [v(b) - \ell(P, B)] + [\ell(P, B) - c(p)].$$

By the above discussion, the two terms that appear in brackets on the right hand side of the last equation are both positive. This allows us to lower bound the gain from trade of  $\hat{S}$  as follows:

$$\begin{aligned} \mathbf{GfT}(\hat{S}) &= \sum_{(p,b) \in \hat{S}} [v(b) - c(p)] \\ &= \sum_{(p,b) \in \hat{S}} \{[v(b) - \ell(P, B)] + [\ell(P, B) - c(p)]\} \\ &\geq \sum_{\substack{b \in \hat{B} \\ b \notin B(A_T \cup A_L)}} [v(b) - \ell(P, B)] + \sum_{\substack{p \in \hat{P} \\ p \notin P(M_T \cup M_L)}} [\ell(P, B) - c(p)]. \end{aligned} \quad \square$$

**Corollary 1.** *OPM is at least  $(1 - r - 22r^{-1} \cdot \sqrt[3]{\alpha} - 10e^{-2/\sqrt[3]{\alpha}})$ -competitive.*

*Proof.* The corollary is trivial when  $r + 22r^{-1} \cdot \sqrt[3]{\alpha} + 10e^{-2/\sqrt[3]{\alpha}} > 1$ . Thus, we assume in this proof  $r + 22r^{-1} \cdot \sqrt[3]{\alpha} + 10e^{-2/\sqrt[3]{\alpha}} \leq 1$ . For every two sets  $M' \subseteq M$  and  $A' \subseteq A$  of mediators and advertisers, respectively, let  $\mathbf{Val}(M', A')$  denote the expression:

$$\sum_{b \in \hat{B} \setminus B(A')} [v(b) - \ell(P, B)] + \sum_{p \in \hat{P} \setminus P(M')} [\ell(P, B) - c(p)].$$

The definition of  $\ell(P, B)$  guarantees that  $v(b) - \ell(P, B) \geq 0$  and  $\ell(P, B) - c(p) \geq 0$  for every  $b \in \hat{B} \subseteq B_o$  and  $p \in \hat{P} \subseteq P_o$ . Thus,  $\mathbf{Val}(M', A') \leq \mathbf{Val}(\emptyset, \emptyset)$  for every two sets  $M' \subseteq M$  and  $A' \subseteq A$ . Additionally, it is well-known that the way  $t$  is chosen by OPM guarantees that every entity of  $M \cup A$  belongs to  $T$  with probability  $r$ , independently (a proof of this fact can be found, e.g., as Lemma A.1 in [12]). Similarly, every entity of  $M \cup A$  that does not belong to  $T$  is added to  $L$  with probability  $\min\{1, 16r^{-1} \cdot \sqrt[3]{\alpha}\} = 16r^{-1} \cdot \sqrt[3]{\alpha}$ , independently. Hence, every user (slot) of  $\hat{P}$  ( $\hat{B}$ ) belongs to  $\hat{P} \setminus P(M_T \cup M_L)$  ( $\hat{B} \setminus B(A_T \cup A_L)$ ) with probability

$$(1 - r)(1 - 16r^{-1} \cdot \sqrt[3]{\alpha}) \geq 1 - r - 16r^{-1} \cdot \sqrt[3]{\alpha}.$$

Therefore,

$$\begin{aligned} & \mathbb{E}[\mathbf{Val}(M_T \cup M_L, A_T \cup A_L)] \\ & \geq (1 - r - 16r^{-1} \cdot \sqrt[3]{\alpha}) \cdot \left\{ \sum_{b \in \tilde{B}} [v(b) - \ell(P, B)] + \sum_{p \in \tilde{P}} [\ell(P, B) - c(b)] \right\} \\ & = (1 - r - 16r^{-1} \cdot \sqrt[3]{\alpha}) \cdot \mathbf{Val}(\emptyset, \emptyset). \end{aligned}$$

Using Lemma 5 and the observation that OPM always produces an assignment of non-negative gain from trade, we now get:

$$\begin{aligned} \mathbb{E}[\mathbf{GfT}(\hat{S})] &= \Pr[\mathcal{E}] \cdot \mathbb{E}[\mathbf{GfT}(\hat{S}) \mid \mathcal{E}] + \Pr[-\mathcal{E}] \cdot \mathbb{E}[\mathbf{GfT}(\hat{S}) \mid -\mathcal{E}] \\ &\geq \Pr[\mathcal{E}] \cdot \mathbb{E}[\mathbf{Val}(M_T \cup M_L, A_T \cup A_L) \mid \mathcal{E}] \\ &= \mathbb{E}[\mathbf{Val}(M_T \cup M_L, A_T \cup A_L)] \\ &\quad - \Pr[-\mathcal{E}] \cdot \mathbb{E}[\mathbf{Val}(M_T \cup M_L, A_T \cup A_L) \mid -\mathcal{E}] \\ &\geq (1 - r - 16r^{-1} \cdot \sqrt[3]{\alpha}) \cdot \mathbf{Val}(\emptyset, \emptyset) - \Pr[-\mathcal{E}] \cdot \mathbf{Val}(\emptyset, \emptyset) \\ &= [(1 - r - 16r^{-1} \cdot \sqrt[3]{\alpha}) - \Pr[-\mathcal{E}]] \cdot \mathbf{Val}(\emptyset, \emptyset). \end{aligned} \quad (1)$$

Recall that  $\Pr[-\mathcal{E}] \leq 10e^{-2/\sqrt[3]{\alpha}}$  by Lemma 4. Additionally, note that Observation 5 and the fact that  $|\tilde{P}| = |\tilde{B}|$  by definition imply together:

$$\begin{aligned} & \mathbf{Val}(\emptyset, \emptyset) \\ &= \sum_{b \in \tilde{B}} [v(b) - \ell(P, B)] + \sum_{p \in \tilde{P}} [\ell(P, B) - c(p)] \\ &= \sum_{b \in \tilde{B}} v(b) - \sum_{p \in \tilde{P}} c(p) \\ &\geq (1 - 6r^{-1} \cdot \sqrt[3]{\alpha}) \cdot \mathbf{GfT}(S_c(P, A)). \end{aligned}$$

Plugging the last observations into (1) gives:

$$\begin{aligned} \mathbb{E}[\mathbf{GfT}(\hat{S})] &\geq [(1 - r - 16r^{-1} \cdot \sqrt[3]{\alpha}) - \Pr[-\mathcal{E}]] \cdot \mathbf{Val}(\emptyset, \emptyset) \\ &\geq [(1 - r - 16r^{-1} \cdot \sqrt[3]{\alpha}) - 10e^{-2/\sqrt[3]{\alpha}}] \cdot (1 - 6r^{-1} \cdot \sqrt[3]{\alpha}) \cdot \mathbf{GfT}(S_c(P, A)) \\ &\geq (1 - r - 22r^{-1} \cdot \sqrt[3]{\alpha} - 10e^{-2/\sqrt[3]{\alpha}}) \cdot \mathbf{GfT}(S_c(P, B)). \end{aligned}$$

The corollary now follows by recalling that  $S_c(P, B)$  is the assignment of users from  $P$  to slots of  $B$  which maximizes the gain from trade.  $\square$

## 4 A Straw-Man Mechanism

In this section we describe a simple straightforward dynamic straw-man mechanism for matching advertisers and mediators. This mechanism maintains the desired economic properties, and thus, it is a good candidate for comparison with OPM. In Sect. 5 we use simulations to compare the two mechanisms. Our simulations show that OPM significantly outperforms the straw-man mechanism.

---

### Mechanism 2. Straw-Man Mechanism

---

1. Let  $\sigma_E$  be an ordered subset of the entities that have arrived so far. Initially  $\sigma_E$  is empty, and entities are added to it as they arrive.

2. For every arriving entity (mediator or advertiser):
    - a. Let  $i$  denote the index of the arriving entity among the previously arriving entities of the same type (in other words, if the arriving entity is a mediator, then it is the  $i$ -th mediator to arrive, and if it is an advertiser, then it is the  $i$ -th advertiser to arrive).
    - b. If  $i$  is odd and the arriving entity is a mediator  $m$  (advertiser  $a$ ), draw a random user of  $m$  (slot of  $a$ ) and denote it by  $\hat{p}_i$  ( $\hat{b}_i$ ). We say that  $c(\hat{p}_i)$  ( $v(\hat{b}_i)$ ) is the threshold for the users of the next mediator (slots of the next advertiser). Moreover, every such user (slot) is *assignable* if her cost is less than (its value is more than) this threshold.
    - c. Otherwise (i.e., when  $i$  is even), if the arriving entity is a mediator  $m$  (advertiser  $a$ ), then add  $m$  ( $a$ ) to the end of  $\sigma_E$ . As long as  $m$  ( $a$ ) has unassigned assignable users (slots) and there is an advertiser (mediator) in  $\sigma_E$  having unassigned assignable slots (users) whose threshold is above (below) the threshold of  $m$ 's users ( $a$ 's slots), do:
      - Let  $a$  ( $m$ ) be the earliest advertiser (mediator) in  $\sigma_E$  of this kind.
      - Assign the unassigned assignable user of mediator  $m$  with the lowest cost to an arbitrary unassigned assignable slot of  $a$ , charge advertiser  $a$  with an amount equal to the threshold of her slots and pay  $m$  an amount equal to the threshold of his users.
    - d. For every mediator  $m \in \sigma_E$  of even index  $i$ , 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  $c(\hat{p}_{i-1})$ .
      - Otherwise, let  $p_i$  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  $c(p_i)$ .
- 

## 5 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  $\beta$  be the budget specified by the bid, and let  $\delta$  be one of the three CPC values specified by this bid, then the advertiser created for this CPC value has a value of  $\delta$  and a



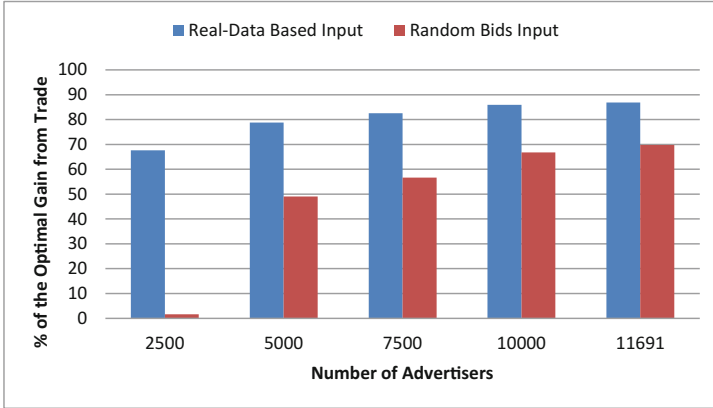
capacity of  $\frac{\beta}{100\delta}$ .<sup>8</sup> For every advertiser we also created a single mediator having the same number of users as the capacity of the advertiser. Every user of these mediators was assigned an independent cost chosen uniformly at random from the range between the smallest 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 than 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 efficient canonical assignment) are depicted in Fig. 1. In order to reduce variance and error margins, every value given by this figure (and the next ones) 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  $\alpha$ ). Moreover, these results demonstrate that our mechanism performs well (between 65% and 85% of the efficient gain from trade) on inputs derived by real world data even for moderate size markets, which is better than what can be predicted based on our theoretical result alone. While one might achieve better performance, a 65% to 85% range seems reasonable given the need to handle both an online setting and economic issues.<sup>9</sup>

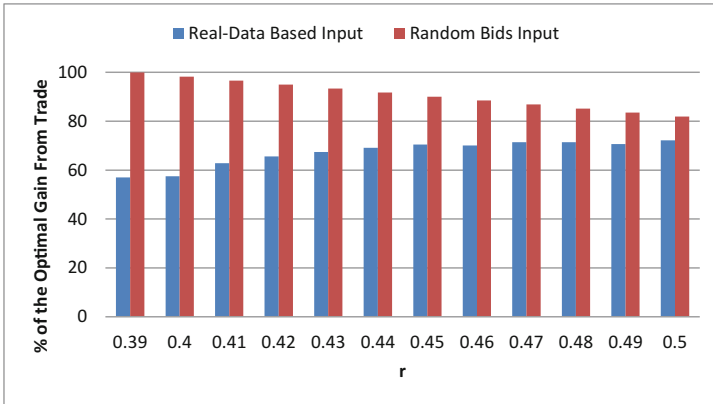
In the previous experiment, we used the value of the parameter  $r$  of OPM which was specified by the version of Theorem 1 given in Sect. 3. Our second simulation was designed to study the possibility of improving the performance of the mechanism by varying the value of  $r$ . Specifically, we repeated 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 Fig. 2. As in the previous experiment, we see again that using the value of  $r$  specified by Theorem 1 ( $1/2$  in

<sup>8</sup> 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.

<sup>9</sup> Note that the experiments did not simulate the information trading part of the model since they were intended to study the mechanism's competitive ratio. However, information exchange can occur in our model in practice. Intuitively, one can think of the users in a single execution of the mechanism as the users who agreed to sell information that, if revealed, implies that they have one particular type  $t$ . Then, if an advertiser's ad is shown to a user, the advertiser may learn that the user is of type  $t$  and the user is monetarily compensated for that.



**Fig. 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.




**Fig. 2.** The performance of OPM as a function of the value of the parameter  $r$  on markets with 11691 advertisers.

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 to be 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.

Our last simulation was designed to demonstrate the need for creating involved solutions such as OPM for solving a dynamic multi-sided market trade problem. We used the straw-man mechanism described in Sect. 4 as a bench-

mark for straightforward mechanisms solving the problem at hand and ran OPM against it. Figure 3 shows that for medium size markets OPM improves by over 60% compared to the straightforward mechanism.



**Fig. 3.** The gain from trade improvement of OPM vs. the straw-man mechanism. The number below each column specifies the number of advertisers plus mediators in the market.

## 6 Conclusion

In this paper we have presented a dynamic model for a foreseeable form of the online advertising market based on principles suggested by [10], 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 non-relevant) for smaller markets, and thus, we have complemented it with simulation results. These results suggest that OPM performs well in practice even for markets of moderate size.

The model we study assumes that all users are equally valuable for the advertisers. Handling users with different value to advertisers is not explicitly described in the model but is supported. The approach would be to assume multiple markets, one for each type of user, and every user is directed by its mediator to the market corresponding to its type (based on the information it

agrees to sell). While this reduction allows the advertisers to effectively have user-dependent valuation functions, studying mechanisms for richer models which directly allow advertisers to have such user-dependent valuation functions (without having to go through the above reduction) is an extension of our current research; which we consider for future work.

## References

1. Ashlagi, I., Monderer, D., Tennenholtz, M.: Mediators in position auctions. *Games Econ. Behav.* **67**, 2–21 (2009)
2. Babaioff, M., Dinitz, M., Gupta, A., Immorlica, N., Talwar, K.: Secretary problems: weights and discounts. In: *SODA*, pp. 1245–1254 (2009)
3. Babaioff, M., Feldman, M., Tennenholtz, M.: Mechanism design with strategic mediators. *ACM Trans. Econ. Comput.* **4**, 7:1–7:48 (2016)
4. Blum, A., Sandholm, T., Zinkevich, M.: Online algorithms for market clearing. In: *SODA*, pp. 971–980 (2002)
5. Bredin, J., Parkes, D., Duong, Q.: Chain: a dynamic double auction framework for matching patient agents. *J. Artif. Intell. Res.* **30**, 133–179 (2007)
6. Charikar, M., Henzinger, M., Nguyen, H.L.: Online bipartite matching with decomposable weights. In: Schulz, A.S., Wagner, D. (eds.) *ESA 2014*. LNCS, vol. 8737, pp. 260–271. Springer, Heidelberg (2014). [https://doi.org/10.1007/978-3-662-44777-2\\_22](https://doi.org/10.1007/978-3-662-44777-2_22)
7. Conitzer, V., Taylor, C.R., Wagman, L.: Hide and seek: costly consumer privacy in a market with repeat purchases. *Mark. Sci.* **31**(2), 277–292 (2012)
8. Facebook: CEO Mark Zuckerberg Testifies Before Senate on User Data (2018). <https://m.youtube.com/watch?v=qAZiDRonYZI>
9. Feldman, M., Gonen, R.: Markets with strategic multi-minded mediators. *CoRR abs/1603.08717* (2016). <http://arxiv.org/abs/1603.08717>
10. Feldman, M., Gonen, R.: Removal and threshold pricing: truthful two-sided markets with multi-dimensional participants. In: Deng, X. (ed.) *SAGT 2018*. LNCS, vol. 11059, pp. 163–175. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-99660-8\\_15](https://doi.org/10.1007/978-3-319-99660-8_15)
11. Feldman, M., Naor, J.S., Schwartz, R.: Improved competitive ratios for submodular secretary problems (extended abstract). In: Goldberg, L.A., Jansen, K., Ravi, R., Rolim, J.D.P. (eds.) *APPROX/RANDOM -2011*. LNCS, vol. 6845, pp. 218–229. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-22935-0\\_19](https://doi.org/10.1007/978-3-642-22935-0_19)
12. Feldman, M., Svensson, O., Zenklusen, R.: A simple  $O(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. *Math. Oper. Res.* **43**(2), 638–650 (2018)
13. Gonen, M., Gonen, R., Pavlov, E.: Generalized trade reduction mechanisms. In: *EC*, pp. 20–29. ACM, New York (2007)
14. Karp, R.M., Vazirani, U.V., Vazirani, V.V.: An optimal algorithm for on-line bipartite matching. In: *STOC*, pp. 352–358 (1990)
15. McAfee, R.P.: dominant strategy double auction. *J. Econ. Theory* **56**, 434–450 (1992)
16. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. *J. Econ. Theory* **29**, 265–281 (1983)
17. Segal-Halevi, E., Hassidim, A., Aumann, Y.: MUDA: a truthful multi-unit double-auction mechanism. In: *Proceeding of AAAI* (2018)

18. Stavrogiannis, L.C., Gerding, E.H., Polukarov, M.: Auction mechanisms for demand-side intermediaries in online advertising exchanges. In: AAMAS, pp. 5–9. International Foundation for Autonomous Agents and Multiagent Systems, Richland (2014)
19. Vaze, R., Coupechoux, M.: Online budgeted truthful matching. *SIGMETRICS Perform. Eval. Rev.* **44**(3), 3–6 (2016)
20. Wurman, P., Walsh, W., Wellman, M.: Flexible double auctions for electronic commerce: theory and implementation. *Decis. Support Syst.* **24**, 17–27 (1998)