Removal and Threshold Pricing: Truthful Two-sided Markets with Multi-dimensional Participants*

Moran Feldman[1] and Rica Gonen[2]

1. Dept. of Mathematics and Computer Science,
moranfe@openu.ac.il
2. Dept. of Management and Economics,
ricagonen@gmail.com
The Open University of Israel,
The Dorothy de Rothschild Campus,
1 University Road, Raanana 43107

Abstract. We consider mechanisms for markets that are two-sided and (in addition) have agents with multi-dimensional strategic spaces on at least one side. The agents of the market are strategic, and act to optimize their own utilities. The mechanism designer, on the other hand, aims to optimize a social goal, *i.e.*, the gain from trade. We focus on one example of this setting which is motivated by the foreseeable future form of online advertising.

Online advertising currently supports some of the most important Internet services, including: search, social media and user generated content sites. To overcome privacy concerns, it has been suggested to introduce user information markets through information brokers into the online advertising ecosystem. Such markets give users control over which data get shared in the online advertising exchange. We describe a model for the above foreseeable future form of online advertising, and design two mechanisms for the exchange of this model: a deterministic mechanism which is related to the vast literature on mechanism design through trade reduction and allows agents with a multi-dimensional strategic space, and a randomized mechanism which can handle a more general version of the model. We provide theoretical analyses for our mechanisms, and study their performance using simulations with real-world data collected by an online advertising system.

1 Introduction

Billions of transactions are carried out via exchanges at every given day, and the number of transactions and exchanges continues to grow as the need for competitiveness promotes adoption. The design of one-sided incentive compatible

* This work is supported by the Horizon 2020 funded project TYPES (Project number: 653449. Call Identifier H2020-DS-2014-1).

(truthful bidding) mechanisms for exchanges is relatively well understood. However, incentive compatible multi-sided mechanisms present a significantly more challenging problem as they introduce more sophisticated requirements such as budget balance.

More specifically, we are interested in designing exchanges (mechanisms) for multi-sided markets with strategic agents. The agents of the market are strategic, and act to optimize their own utilities. The mechanism designer, on the other hand, aims to optimize a social goal, *i.e.*, the gain from trade (the difference between the total value of the sold goods for the buyers and the total costs of these goods for the sellers). The design of such mechanisms raises a few interesting questions. Can the mechanism maintain simultaneously different desirable economic properties such as: individual rationality (IR)—participants do not lose by participation, incentive compatibility (IC)—agents are incentivized to report their true information to the mechanism and budget balance (BB)—the mechanism does not end up with a loss. Moreover, can the mechanism maintain these properties while suffering only a bounded loss compared to the optimal gain from trade? Finally, can this be done when **all** the agents have a **multi-dimensional strategic space**?[1]

The above questions can be studied in the context of many multi-sided markets. We focus on one such market, and leave the consideration of other multi-sided markets for future work. The market we consider is motivated by online advertising in its foreseeable future form. Online advertising currently supports some of the most important Internet services, including: search, social media and user generated content sites. However, the amount of information that companies collect about users increasingly creates privacy concerns in society as a whole, and even more so in the European society. In recent years EU regulators have actively been looking for solutions to guarantee that users' privacy is preserved. In particular, the EU regulators have been looking for tools that enable the end user to configure their privacy settings so that only information allowed by the end-user is collected by online advertising platforms.

The market we study is induced by a new solution we suggest for the above privacy issue. In this solution mediators serve as the interface between end-users and the other agents in the online advertising market. Each user informs her mediator of the attributes she is willing to reveal, and her cost, *i.e.*, the compensation she requires for every ad she views. The mediator then tries to "sell" the user on the advertising market based *only* on the attributes she agreed to reveal, and, if successful, pays her the appropriate cost out of the amount he got from the sell.

As revealing more attributes is likely to result in a more profitable sale, our solution provides incentives for users to share their information with the advertising market while allowing users to retain control of the amount of information they would like to share. Notice that the fact that our solution motivates users to participate in the advertising market, and even to provide more precise infor-

[1] We often refer to agents with a multi-dimensional strategic space as multi-dimensional agents.

mation for targeting campaigns, means that our solution improves the efficiency of the advertising system and the digital economy as a whole (in addition to answering the privacy concerns discussed above). This is in sharp contrast to other natural approaches for dealing with privacy issues, such as cryptography based approaches, which reduce the amount of information available to the advertisers but give them nothing in return.

The advertising market induced by the above solution has mediators on one side, and advertisers on the other side. Each mediator has a set of users associated with him, and he is trying to assign these users to advertisers using the market. Each one of the users has a non-negative cost which she has to be paid if she is assigned to an advertiser. The mediators themselves have no cost of their own, however, each of them has to pay his users their cost if they are assigned to advertisers. Thus, the utility of a mediator is the amount paid to him minus the total cost of his users that are assigned. Finally, each advertiser has a positive capacity determining the number of users she is interested in targeting, and she gains a non-negative value from every one of the users assigned to her (as long as her capacity is not exhausted). Thus, the advertiser's utility is her value multiplied by the number of users assigned to her (as long as this number does not exceed her capacity) minus her total payment.

A mechanism for the above market knows the mediators and the advertisers, but has no knowledge about their parameters or about the users. The objective of the mechanism is to find an assignment of users to advertisers that maximizes the gain from trade. In addition, the mechanism also decides how much to charge (pay) each advertiser (mediator). In order to achieve these goals, the mechanism receives reports from the advertisers and mediators. Each advertiser reports her capacity and value, and each mediator reports the number of his users and their costs. The mediators and advertisers are strategic, and thus, free to send incorrect reports. In other words, an advertiser may report incorrect capacity and value, and a mediator may report any subset of his users and associate an arbitrary cost with each user. We say that an advertiser is *truthful* if she reports correctly her capacity and value. A mediator is considered *truthful* if he reports to the mechanism his true number of users and the true costs of these users. Notice that we assume that the costs of the users are known to their corresponding mediators, *i.e.*, the users are non-strategic. This assumption is reasonable given the high speed of the online advertising market, especially compared to the speed at which a private user can change her contract with her mediator.

To better understand the design challenge raised by this market, we observe that even if our setting is reduced to a single buyer-single seller exchange, still it is well known from [18] that maximizing gain from trade while maintaining individual rationality and incentive compatibility perforce to run into deficit (is not budget balanced). A well known circumvention of [18]'s impossibility is [17]'s trade reduction for a simple setting of double sided auctions. In [17]'s setting, trade is conducted between multiple strategic sellers offering identical goods to multiple strategic buyers, where each seller is selling a single good and each buyer is interested in buying a single good. The result of [17] relaxes the requirement for optimal trade by means of a *trade reduction*. The trade reduction leads to an individually rational, incentive compatible and budget balance mechanism. Following [17]'s work, several other mechanisms were designed using the technique of trade reduction. However, all the trade reduction mechanisms suggested in the literature to date allow only agents with single dimensional strategic spaces (even in settings where agents can hold multiple items).[2]

1.1 Our Contribution

Given that existing trade reduction solutions do not apply in our setting, we describe new double-sided mechanisms able to handle mediators and advertisers with multi-dimensional strategic spaces. Our mechanisms guarantee desirable economic properties, and at the same time yield a gain from trade approximating the optimal gain from trade. If being truthful is a dominant strategy[3] for each advertiser and each mediator (regardless of other agents' strategies), then the mechanism is *incentive compatible* (IC); if no advertiser and no mediator can have a negative utility by participating truthfully in the mechanism than it is *individually rational* (IR). Our objective is to construct mechanisms that are IC, IR and budget balanced (BB).

We first study a special case of our setting where the advertisers' capacities are publicly known (however, these capacities need not be all equal). The set of users of each mediator, on the other hand, remains unknown to the mechanism (*i.e.*, the mechanism only learns about it through the mediator's report). For this case we present a deterministic mechanism named "Price by Removal Mechanism" (PRM) that works as follows: for every mediator find a threshold cost, and remove users of the mediator whose cost is above this threshold. Add a dummy advertiser with value that is the maximum threshold cost computed for the mediators and a capacity that is equal to the total number of users remaining. Assign the non-removed users to the advertisers using a VCG auction [22,10,16] in which the users are the goods and the bidders are the advertisers. Price the mediators according to their threshold cost, and price the advertisers according to the prices of the VCG auction describe above.

The method used to calculate the threshold costs of the above mechanism induces its properties. We prove that, for appropriately chosen threshold costs, the above mechanism is IC, IR, BB and provides a non-trivial approximation for the optimal gain from trade. More formally, if τ is the size of the optimal trade, and γ is an upper bound, known to the mechanism, on the maximum capacity of any agent (mediator or advertiser), then:

Theorem 1. *PRM is BB, IR, IC and* $\left(1 - \frac{5\gamma}{\tau}\right)$*-competitive.*

[2] The sole exception for this is the mechanism developed by Segal-Halevi et al. [20], which was developed independently in parallel to the current work.

[3] Here and throughout the paper, a reference to domination of strategies should be understood as a reference to weak domination. We never refer to strong domination.

An online advertising system constructed based on PRM and beta tested with real users and real advertising campaigns[4] allowed us to collect real-world data to study PRM performance empirically. Interestingly, our empirical simulation

shows that although the practical performance of PRM is significantly better than its theoretical one, both performances exhibits a similar dependence on the ratio γ/τ.

PRM generalizes the trade reduction ideas used so far in the literature for single dimensional strategic agents, but is much more involved. Intuitively, PRM differs from previous ideas by the following observation. A trading set is the smallest set of agents that is required for trade to occur. In the existing literature for single dimensional strategic agents a trade reduction mechanism makes a binary decision regarding every trading set of the optimal trade, *i.e.*, either the trading set is removed as a whole, or it is kept. On the other hand, dealing with multi-dimensional agents requires PRM to remove only parts of some trading sets, and thus, requires it to make non-binary decisions.

Our deterministic mechanism PRM handles one type of multi-dimensional agents (the mediators) and one type of single dimensional strategic space agents (the advertisers). In order to enrich our strategic space even further, and allow advertisers to have multi-dimensional strategic spaces as well, we present also a randomized mechanism termed "Threshold by Partition Mechanism" (TPM). TPM applies to our general setting, *i.e.*, we no longer assume that any capacity is known to the mechanism, and it works as follows: divide the set of mediators uniformly at random into two sets (M_1 and M_2) and divide the set of advertisers uniformly at random, as well, into two sets (A_1 and A_2). Then use the optimal trade for M_2 and A_2 to produce threshold cost and threshold value that allow BB pricing as well as the needed reduction in trade for M_1 and A_1. Analogously, use the optimal trade for M_1 and A_1 to produce threshold cost and threshold value that allow BB pricing as well as the needed reduction in trade for M_2 and A_2.

The above description of TPM is not complete since the use of threshold cost and value from one pair (M_i, A_i) to reduce the trade in the other pair might create an unbalanced reduction. To overcome this issue we create two random low priority sets: one of advertisers and the other of mediators. Then, whenever the reduction in trade is unbalanced, we remove additional low priority mediators or advertisers in order to restore balance (which can be done with high probability). The following theorem shows that the above mechanism is IC, IR, BB and provides a non-trivial approximation for the optimal gain from trade. The parameter α is an upper bound, known to the mechanism, on the ratio between the maximum capacity of any agent (mediator or advertiser) and the size of the optimal trade.[5]

Theorem 2. *TPM is BB, IR, IC and* $(1 - 28\sqrt[3]{\alpha} - 20e^{-2/\sqrt[3]{\alpha}})$*-competitive.*

[4] As part of a Horizon 2020 project. Details about this project are deferred to the camera ready version due to the double blind review process.

[5] The parameters γ and α both bound the maximum capacity of the agents. Moreover, they are formally related by the formula $\alpha = \gamma/\tau$. We chose to formulate Theorems 1

We note that TPM is universally truthful, *i.e.*, its IC property holds for every given choice of the random coins of the mechanism. Observe also that the competitive ratio of TPM approaches 1 when α approaches 0, *i.e.*, when the market is large enough to make the market power of all agents very low. Unfortunately, when the market is not large enough to make α very small, the theoretical competitive ratio guaranteed by Theorem 2 is not so good, and might even be meaningless. Nevertheless, our simulations suggest that in practice the performance of TPM is quite good even for moderate size markets.

Both our mechanisms suffer from a common drawback, namely their need to have access to a good bound on the maximum market power of any agent (which is captured by the parameters γ and α). From a practical point of view we believe this is a minor issue, as the mechanism can usually use the large quantity of historical data available to it to estimate the necessary bound very well.

1.2 Related Work

From a motivational point of view the most closely related literature to our work consists of works that involve mediators and online advertising markets, such as [1,14,21]. These works differ from ours in two crucial points. First, despite being motivated by the online display ads network exchange, the models studied by these works are actually auctions (*i.e.*, one-sided mechanisms). Thus, they need not deal with the challenges and impossibility integrated by the double-sided structure of our market and the requirement to keep it from running into a deficit. Second, our focus is maximization of the gain from trade, unlike the above works which focus on revenue maximization.

Another related work involving both markets and mediators studies the phenomenon of markets in which individual buyers and sellers trade through intermediaries, who determine prices via strategic considerations [4]. An essential difference between the model of [4] and our model is that [4] does not assume private values for the agents, and therefore, the impossibility of [18] does not apply in its model.

We now move our attention to the above mentioned literature on trade reduction and multi-sided markets. Deterministic mechanisms using trade reduction as a mean to achieve IC, IR and BB were described for various settings [17,8,2,15,9,3]. Moreover, for a variant of the setting of [17,3], [19] obtained a randomized mechanism achieving IC, IR and strong budget balance (*i.e.*, it is BB and leaves no surplus for the market maker). The mutual grounds of all these settings is that all agents participating in the trade have a single dimension strategic space. This idea was captured by [15] which provided a single trade reduction procedure applicable to all the above settings. In addition, [15] also defined a class of problems that can be solved by its suggested trade reduction procedure. Essentially this classification is based on partitioning the agents participating in the trading set into equivalence classes.

and 2 in terms of the parameter that the mechanism corresponding to each theorem assumes access to.

As pointed out in the previous subsection, both our presented mechanisms extend significantly on the existing trade reduction literature. More specifically, even when all advertisers have known equal capacities (while mediators can still have a variable number of users), fitting our model into the classification of [15] still requires each mediator to have his own equivalence class (because a mediator with many users can always replace a mediator with a few users within a trading set, but the reverse is often not true). It follows that [15]'s trade reduction procedure might remove all the trade, and thus, achieves only a trivial gain from trade approximation.

Recent related research on maximizing gain from trade in two-sided markets was published by [6,7,11,20]. The most relevant work to ours is the work by [20]. As marked earlier in the section, our work was developed independently in parallel to [20], and was posted on the archive shortly before theirs. To the best of our knowledge both papers were the first to tackle multi-dimensional agents in the context of two-sided markets. Though the papers use some similar ideas, their results are incomparable. Firstly, [20] only presents a randomize mechanism, while we present both deterministic and randomized mechanisms. Secondly, in settings in which both our randomized mechanism and the mechanism of [20] apply, each mechanism achieves a superior competitive ratio for a different range of the parameters.

Last but not least, we note that any result for our objective function applies also to a different objective function known as the social welfare, and therefore, our mechanisms can also be used for maximization of the social welfare in similar multi-sided markets. Hence, our work is also related to the works of [5,13,12] which study the maximization of the social welfare objective in combinatorial and matroid models of multi-sided markets.

2 Notation and Basic Observations

We begin this section with a more formal presentation of our model. Our 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 has to be paid if she is assigned to an advertiser. 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 he has to forward to his assigned users; 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, then the utility of m is $t - \sum_{p \in P(m)} x(p) \cdot c(p)$. 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 $n \cdot v(a) - t$.

A mechanism for our model accepts reports from the advertisers and mediators, and based on these reports outputs an assignment of users to advertisers. In addition, the mechanism also decides how much to charge (pay) each advertiser (mediator). The objective of the mechanism is to output an assignment of users to advertisers that maximizes the gain from trade.

For ease of the presentation, it is useful to associate a set $B(a)$ of $u(a)$ slots with each advertiser $a \in A$. We then 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 the value $v(b)$ of 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

$$\texttt{GfT}(S) = \sum_{(p,b) \in S} [v(b) - c(p)] \ .$$

In addition to the above notation, we would like to define two additional shorthands that we use occasionally. 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 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 mediators of M'.

The presentation of our mechanisms is simpler when the values of slots and the costs of users are all unique. Clearly, this is extremely unrealistic as all the slots of a given advertiser have the exact same value in our model. Thus, we simulate uniqueness using a tie-breaking rule.

Observation 3 *When the slots are ordered in an increasing (or decreasing) value order, all the slots of a single advertiser are always consecutive.*

2.1 Canonical Assignment

Given a set $B' \subseteq B$ of slots and a set $P' \subseteq P$ of users, the canonical assignment $S_c(P', B')$ is the assignment constructed by the following process. 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 used frequently by the mechanisms we describe in the next sections. In some places we refer to the user or slot at location i of a canonical solution $S_c(P', B')$. By using this expression we mean user p_i or slot b_i, respectively. Additionally, the term $|S_c(P, B)|$ is used very often in our proofs, and thus, it is useful to define the shorthand $\tau = |S_c(P, B)|$.

The following lemma shows that the above definition of τ is consistent with the use of τ in Section 1.1 as the size of the optimal trade.

Lemma 1. *The canonical assignment $S_c(P', B')$ maximizes $\texttt{GfT}(S_c(P', B'))$ among all the possible assignments of users of P' to slots of B'.*

The proof of Lemma is omitted due to space limitations.

3 Deterministic Mechanism

In this section we describe the deterministic mechanism "Price by Removal Mechanism" (PRM) for our model. Recall that PRM assumes public knowledge of the advertisers' capacities. Accordingly, we assume throughout this section that the capacities of the advertisers are common knowledge (or that the advertisers are not strategic about them). We also assume that PRM has access to a value $\gamma \geq 1$ such that:

$$u(a) \leq \gamma \quad \forall\, a \in A \qquad \text{and} \qquad |P(m)| \leq \gamma \quad \forall\, m \in M \ .$$

In other words, γ is an upper bound on how large can be the capacity of an advertiser or the number of users of a mediator. Informally, γ can be understood as a bound on the importance every single advertiser or mediator can have.

A description of PRM is given as Mechanism 1.1. Notice that Mechanism 1.1 often refers to parameters of the model that are not known to the mechanism (*i.e.*, values of advertiser, the number of users of mediators and the costs of users). Whenever this happens, this should be understood as referring to the reported values of these parameters.

Mechanism 1.1: Price by Removal Mechanism (PRM)

1. For every mediator $m \in M$, if the canonical assignment $S_c(P \setminus P(m), B)$ is of size more than 4γ, denote by p_m the user at location $|S_c(P \setminus P(m), B)| - 4\gamma$ of the canonical assignment $S_c(P \setminus P(m), B)$, and let c_m be the cost of p_m. Otherwise, set c_m to $-\infty$.

2. For every mediator $m \in M$, let $\hat{P}(m)$ be the set of users of mediator m whose cost is less than c_m. Intuitively, $\hat{P}(m)$ is the set of users of mediator m that the mechanism tries to assign to advertisers.

3. Assign the users of $\bigcup_{m \in M} \hat{P}(m)$ to the advertisers using a VCG auction. More specifically, the users of $\bigcup_{m \in M} \hat{P}(m)$ are the items sold in the auction, and the bidders are the advertisers of A plus a dummy advertiser a_d whose value and capacity are $v(a_d) = \max_{m \in M} c_m$ and $u(a_d) = \sum_{m \in M} |\hat{P}(m)|$, respectively.

4. Charge every non-dummy advertiser by the same amount she is charged (as a bidder) by the VCG auction.

5. For every user p assigned by the VCG auction, pay c_m to the mediator m of p.[6]

Remark: It can be shown that the existence of the dummy advertiser never affects the behavior of Mechanism 1.1, and thus, one can safely omit it from the mechanism. Nevertheless, we keep this advertiser in the above description of the mechanism since its existence simplifies our proof that the mechanism is BB.

[6] Note that m is BB as he forwards to each of his assigned users her cost—which is less than c_m.

Let us recall our result regarding PRM (the proof of Theorem 1 is deferred to Appendix **??**).

Theorem 1. *PRM is BB, IR, IC and* $\left(1 - \frac{5\gamma}{\tau}\right)$*-competitive.*

PRM improves over the trade reduction procedure of [15] by guarantying a non-trivial competitive ratio even in the presence of agents (mediators) with a multi-dimensional strategic space. However, the competitive ratio guaranteed by PRM is slightly worse than the competitive ratio guaranteed by the procedure of [15] when the agents have single dimensional strategic spaces (which is $1-2/\tau$). The next paragraph gives an intuitive explanation why it does not seem possible to improve the competitive ratio of PRM to match the competitive ratio of $1-2/\tau$ guaranteed by [15]'s procedure.

Observe that PRM needs to maintain IC for both mediators and advertisers. In order to maintain IC for the mediators the price for each mediator has to be computed while all his users are reduced from the trade. This is why the **real** reduction in trade might be larger by up to γ compared to the reduction 4γ explicitly stated by PRM, which explains the gap between the competitive ratio of $1 - 5\gamma/\tau$ and the explicit reduction of 4γ in the mechanism. It remains to understand why the explicit reduction is set to 4γ. The most significant difficulty in guaranteeing IC for the advertisers is that an advertiser might change her report in order to affect the costs $\{c_m\}_{m \in M}$ of the mediators, and through them, manipulate the items offered in the VCG auction. PRM tackles this difficulty by guaranteeing that advertisers performing such manipulations are assigned no users. Since the users of mediator m are removed when c_m is calculated, securing this guarantee requires that advertisers having a slot in one of the last γ locations of the optimal trade are assigned no users. As an advertiser might have up to γ slots, this translates into a requirement that an advertiser whose earliest slot is in one of the last 2γ locations of the optimal trade is assigned no users. Moreover, to simplify the proof our analysis in fact requires that an advertiser whose earliest slot is in one of the last 3γ locations of the optimal trade is assigned no users. Taking into account, again, the fact that the users of mediator m are removed when c_m is calculated, guaranteeing the last property requires a trade reduction of 4γ. To summarize, maintaining the advertisers' IC imposes an explicit trade reduction of 3γ that we present as 4γ for the sake of simplicity, and maintaining the mediators' IC implies that the real trade reduction is larger by up to γ compared to the explicit one. Thus, a competitive ratio of $1 - \frac{4\gamma}{\tau}$ seems to be inevitable in order to allow multi-dimensional strategic spaces.

4 Randomized Mechanism

In this section we describe the randomized mechanism "Threshold by Partition Mechanism" (TPM) for our model. Unlike the mechanism PRM from Section 3, TPM need not assume public knowledge about the advertisers' capacities, *i.e.*, the advertisers now have multi-dimensional strategy spaces. On the other hand, TPM assumes access to a value $\alpha \in [|S_c(P,B)|^{-1}, 1]$ such that we are guaranteed that:

$$\frac{u(a)}{|S_c(P,B)|} \leq \alpha \quad \forall\, a \in A \qquad \text{and} \qquad \frac{|P(m)|}{|S_c(P,B)|} \leq \alpha \quad \forall\, m \in M \enspace .$$

In other words, α is an upper bound on how large can be 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 is related to the value γ from Section 3 by the equation $\alpha = \gamma/\tau$, and thus, α, like γ, can be informally understood as a bound on the importance of every single advertiser or mediator. It is important to note that α is well-defined only when $|S_c(P,B)| > 0$, and thus, we assume this inequality is true throughout the rest of the section.

A description of TPM is given as Mechanism 1.2. Notice that Mechanism 1.2 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 1.2: Threshold by Partition Mechanism (TPM)

1. Let M_L be a set of mediators containing each mediator $m \in M$ with probability $\min\{17\sqrt[3]{\alpha}, 1\}$, independently. Similarly, A_L is a set of advertisers containing each advertiser $a \in A$ with probability $\min\{17\sqrt[3]{\alpha}, 1\}$, independently. Intuitively, the subscript L in M_L and A_L stands for "low priority".

2. Let σ_A be an arbitrary order over the advertisers that places the advertisers of A_L after all the other advertisers and is independent of the reports received by the mechanism. Similarly, σ_M is an arbitrary order over the mediators that places the mediators of M_L after all the other mediators and is independent of the reports received by the mechanism.

3. Partition the mediators of M into two sets M_1 and M_2 by adding each mediator $m \in M$ with probability $1/2$, independently, to M_1 and otherwise to M_2. Similarly, partition the advertisers of A into two sets A_1 and A_2 by adding each advertiser $a \in A$ with probability $1/2$, independently, to A_1 and otherwise to A_2. The rest of the algorithm explains how to assign users of mediators from M_1 to slots of advertisers from A_1, and how to charge advertisers of A_1 and pay mediators of M_1. Analogous steps, which we omit, should be added for handling the advertisers of A_2 and the mediators of M_2.

4. Let $\hat{p}$ and $\hat{b}$ be the user and slot at location $\lceil (1-4\sqrt[3]{\alpha}) \cdot |S_c(P(M_2), B(A_2))| \rceil$ of the canonical solution $S_c(P(M_2), B(A_2))$. If $(1-4\sqrt[3]{\alpha}) \cdot |S_c(P(M_2), B(A_2))| \leq 0$, 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 ∞. Using $\hat{p}$ and $\hat{b}$ define now two sets

$$\hat{P} = \{p \in P(M_1) \mid c(p) < c(\hat{p})\} \qquad \text{and} \qquad \hat{B} = \{b \in B(A_1) \mid v(b) > v(\hat{b})\} \enspace .$$

It is important to note that $\hat{P}$ and $\hat{B}$ are empty whenever $\hat{p}$ and $\hat{b}$ are dummy user and slot, respectively.

5. While there are unassigned users in $\hat{P}$ and unassigned slots in $\hat{B}$ do the following:
- Let m be the earliest mediator in σ_M having unassigned users in $\hat{P}$.
- Let a be the earliest advertiser in σ_A having unassigned slots in $\hat{B}$.
- Assign the unassigned user of $\hat{P} \cap P(m)$ with the lowest cost to an arbitrary unassigned slot of $\hat{B} \cap B(a)$, charge a payment of $v(\hat{b})$ to advertiser a and transfer a payment of $c(\hat{p})$ to mediator m.[7]

Let us recall our result regarding TPM.

Theorem 2. *TPM is BB, IR, IC and* $(1 - 28\sqrt[3]{\alpha} - 20e^{-2/\sqrt[3]{\alpha}})$*-competitive.*

Intuitively, the analysis of TPM works by exploiting concentration results showing that the canonical assignments $S_c(P(M_1), B(A_1))$ and $S_c(P(M_2), B(A_2))$ are quite similar. This similarity allows us to use information from $S_c(P(M_2), B(A_2))$ to set the payments charged to advertisers of $B(A_1)$ and payed to mediators of $P(M_1)$, and vice versa, while keeping a reasonable competitive ratio. The advantage of setting the payments this way is that it reduces the control agents have on the payments they have to pay or are paid, which helps the mechanism to be IC.

5 Experiments

We have used simulations with real-world data to study the empirical performance of our mechanisms. The real-world data was collected by an online advertising system constructed based on PRM and beta tested with real users and real advertising campaigns as part of a Horizon 2020 project. Let us begin this section by describing the simulations we used to study the performance of our deterministic mechanism (PRM). Our simulations included 30 advertisers, each associated with a campaign. For every campaign the collected data included a maximal CPC (cost-per-click) value and a minimal one, and the value of advertiser associated with this campaign was chosen uniformly at random between the minimal and maximal CPC. For each execution of the simulation, we picked a random upper bound between 1 and 65 on the capacity of the advertisers in this execution, and then picked for every advertiser a uniformly random capacity between 1 and the above mentioned upper bound. This method for picking the advertisers' capacities might look a bit odd compared to the more natural approach of picking a random capacity between 1 and 65 for each advertiser directly, but it has the advantage of generating simulation inputs with diverse γ values, and thus, allows us to study the dependence of the empirical performance of PRM on γ.

The users for our simulations were created based on data fed by 328 real users of the above mentioned online advertising system constructed based on

[7] Note that m is paid $c(\hat{p})$ for the assignment of each one of his users. Hence, m is always BB since the membership of p in $\hat{P}$ implies $c(p) < c(\hat{p})$ (and $c(p) \leq c(\hat{p})$ when the costs are compared as numbers).

PRM. Accordingly, our simulations included 328 users, each corresponding to a one of the above real users. The cost of each user in our simulations was chosen at random out of a range of "reasonable" costs reported by the corresponding real user, and the users were then assigned at random to 30 mediators.

After constructing 500 random inputs using the technique described above, we executed PRM on each one of them independently. The results of these executions are summarized in Figure 1. Each execution of PRM is depicted in this figure using two dots: an orange one and a blue one. The x-axis of both these dots is the γ/τ ratio of the input of this execution. The y-axis value of the orange dot is the theoretical competitive ratio that is guaranteed for PRM given this γ/τ ratio, and the y-axis value of the blue dot is the competitive ratio obtained in practice by PRM for this input. Naturally, the blue dots are located above their orange counterparts. Moreover, one can observe that most of the gaps between pairs of corresponding dots are roughly of the same size. In other words, the empirical competitive ratio of PRM for every input is roughly equal to 0.2 plus the theoretical guarantee of PRM for the γ/τ ratio of this input.

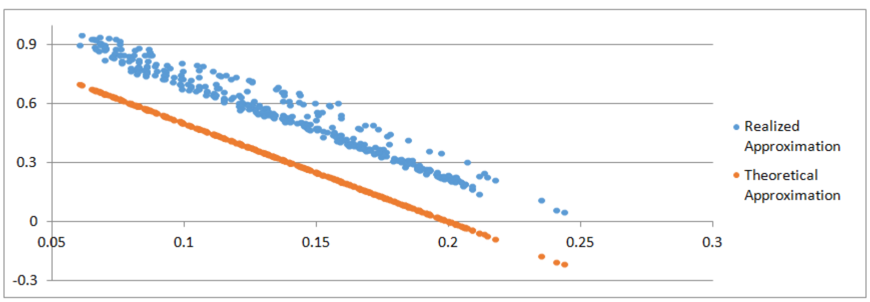

Fig. 1: Simulation results for PRM. The orange dots represent the theoretical performance guarantee for the inputs on which the simulations were run: the y-axis value of the dot is the competitive ratio and the x-axis value is the ratio γ/τ in the input corresponding to the dot. Similarly, the blue dots represent the performance obtained by the mechanism in reality on the same inputs.

It is natural to try to use the above input generation technique also for simulations of our randomized mechanism (TPM). Unfortunately, this cannot be done because TPM cannot work with such thin markets generated by the above technique. More specifically, the theoretical guarantee of TPM, as given by Theorem 2, is meaningful only for very small values of α, which usually can arise only from very large markets. This strongly suggests that TPM is unlikely to produce good results when applied to thin markets such as the ones we considered when studying PRM. Nevertheless, we conjectured that TPM should work well in practice for moderate size markets (which correspond to moderately small values of α), despite the fact that the theoretical analysis fails to show that. To test this conjecture, we needed to generate inputs of moderate size for TPM. Each one of these inputs consisted of 20000 users, where each one of these users was associated with a random one of the above mentioned 328 real users, and the cost of each simulated user was picked by taking a random cost out of the range of reasonable costs selected by its corresponding real user. The generated users were than grouped into equal size groups, and each group of users was assigned to a different mediator. The number of users per group was a parameter that we varied between simulations in order to produce inputs with different α values. Similarly, we created some simulation advertisers, where each one of these advertisers was associated with a random one out of the above mentioned 30 campaigns, and the value of each advertiser was chosen as a uniformly random value between the minimum and maximum CPC values of its corresponding campaign. The number of advertisers was another parameter that we varied between simulations in order to produce different α values, but regardless of this number, the capacity of the advertisers was set to a value that made the total number of slots of all the advertisers equal to 20000.

The results of the simulations of TPM on the inputs generated by the above technique are summarized in Figure 2. As is evident from this figure, TPM achieves in these simulations roughly 30% of the optimal trade even when α is as large as roughly 10^{-2}. Moreover, the empirical competitive ratio of TPM improves rapidly for smaller values of α, reaching roughly 80% for $\alpha \approx 10^{-3}$. Thus, our simulations support our conjecture that TPM works well for moderate size markets (and moderately small α values), despite the fact that we do not have a meaningful theoretical guarantee for this range of α values.

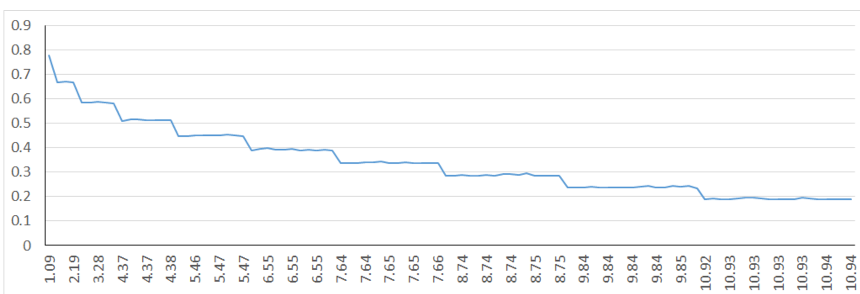

Fig. 2: The empirical competitive ratio obtained in our simulations of TPM as a function of α. The values on the y-axis are the competitive ratios, and the values on the x-axis are α times 10^3. Each data point in this graph averages 500 independent executions of TPM.

To better understand the gap between the theoretical and empirical performance of TPM, we ran experiments also for values of α for which the theoretical guarantee of TPM is meaningful. The inputs for these simulations were generated using a method which is identical to the method used to generate the inputs for the previous simulations, except for two changes. First, to allow for the production of very small α values with a reasonable number of users and advertisers, we set both the number of users per mediators and the number of slots per advertiser to 4. Second, to generate various values of α, we varied the number of users and slots in the inputs from 10^4 to 10^5. The results of the simulations of TPM on the inputs generated by this technique are summarized in Figure 3. Observe that the α values for these inputs roughly range between 10^{-4} and 10^{-5}. Since all these α values are much smaller than the moderately small α values of the previous batch of simulations, it should be no surprise that TPM achieves an empirical competitive ratio of almost 1 for all α values in Figure 3. In contrast, the theoretical guarantee for this range of α values ranges between 0% for $\alpha \approx 5 \cdot 10^{-5}$ and roughly 40% for $\alpha \approx 10^{-5}$. Comparing Figures 2 and 3, we can deduce the following crude rule of thumb: the empirical performance of TPM for a given value of α is similar to its theoretical performance for α values that are smaller by two orders of magnitude.

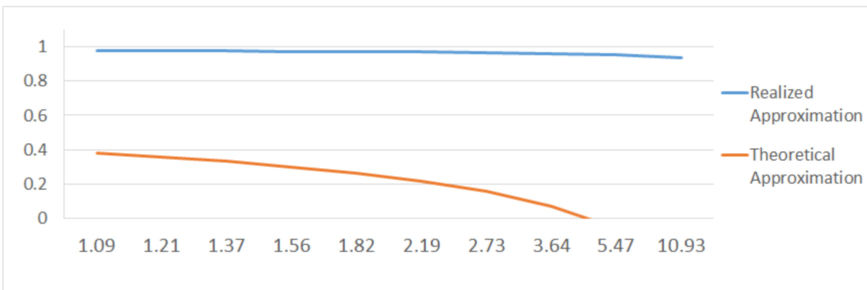

Fig. 3: The empirical competitive ratio obtained in our simulations of TPM as a function of α for low values of α, and the corresponding theoretical guarantee. The values on the y-axis are the competitive ratios, and the values on the x-axis are α times 10^5. Each data point in this graph averages 500 independent executions of TPM.

6 Conclusions

We considered mechanisms for two-sided markets that interact with strategic agents where at least one side of the market has agents with multi-dimensional strategic spaces. In particular, we explored the question of how many sides of the market can have agents with multi-dimensional strategic spaces, while still allowing for mechanisms that both maintain the desired basic economic properties

and suffer only a

bounded loss compared to the socially optimal outcome. To answer this question we presented two mechanisms. One mechanism that is deterministic and allows one side to have agents with multi-dimensional strategic spaces, and another mechanism that is randomized and allows two sides to have agents with multi-dimensional strategic spaces.

Our mechanisms significantly extend the literature on trade reduction—a technique used to achieve the desired basic economic properties in a multi-sided market. While all the previous trade reduction solutions dealt with agents having single dimensional strategic spaces, our deterministic algorithm performs a non-binary trade reduction which leads to the first trade reduction solution applying to multi-dimensional strategic spaces.

From a more practical point of view, our study is motivated by a foreseeable future form of online advertising in which users are incentivized to share their information via participation in a mediated online advertising exchange. As such our simulations use real-world data collected by an online advertising system constructed based on our deterministic mechanism and beta tested with real users and real advertising campaigns. Our empirical study of the deterministic mechanism as well as the randomized one suggests that the empirical performance of both markets is significantly better than the theoretical guarantee and can reach near optimal performance under the relevant market settings.

References

1. Ashlagi, I., Monderer, D., Tennenholtz, M.: Mediators in position auctions. Games and Economic Behavior **67**, 2–21 (2009)
2. Babaioff, M., Walsh, W.E.: Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. In: EC. pp. 65–74. ACM, New York, NY, USA (2003)
3. Babaioff, M., Nisan, N., Pavlov, E.: Mechanisms for a spatially distributed market. Games and Economic Behavior **66**(2), 660–684 (2009)
4. Blume, L., Easley, D., Kleinberg, J., Tardos, E.: Trading networks with price-setting agents. In: EC. pp. 13–16. ACM, New York, NY, USA (2007)
5. Blumrosen, L., Dobzinski, S.: Reallocation mechanisms. In: EC. pp. 617–640. ACM, New York, NY, USA (2014)
6. Blumrosen, L., Mizrahi, Y.: Approximating gains-from-trade in bilateral trading. In: In Proceedings of the 12th Conference on Web and Internet Economics (WINE). 400–413. (2016)
7. BRUSTLE, J., CAI, Y., WU, F., ZHAO, M.: Approximating gains from trade in two-sided markets via simple mechanisms. In: In Proceedings of the 18th ACM Conference on Economics and Computation (EC). 589–590 (2017)
8. Chaib-draa, B., Müller, J.: Multiagent Based Supply Chain Management (Studies in Computational Intelligence). Springer-Verlag New York, Inc., Secaucus, NJ, USA (2006)
9. Chen, R.R., Roundy, R.O., Zhang, R.Q., Janakiraman, G.: Efficient auction mechanisms for supply chain procurement. Manage. Sci. **51**(3), 467–482 (March 2005)
10. Clarke, E.H.: Multipart pricing of public goods. Public Choice **2**, 17–33 (1971)
11. Colini-Baldeschi, R., Goldberg, P., de Keijzer, B., Leonardi, S., Turchetta, S.: Fixed price approximability of the optimal gain from trade. In: Wine (2017)

12. Colini-Baldeschi, R., Goldberg, P.W., de Keijzer, B., Leonardi, S., Roughgarden, T., Turchetta, S.: Approximately efficient two-sided combinatorial auctions. In: EC. pp. 591–608. ACM, New York, NY, USA (2017)
13. Colini-Baldeschi, R., de Keijzer, B., Leonardi, S., Turchetta, S.: Approximately efficient double auctions with strong budget balance. In: SODA. pp. 1424–1443. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2016)
14. Feldman, J., Mirrokni, V.S., Muthukrishnan, S., Pai, M.M.: Auctions with intermediaries: extended abstract. In: EC. pp. 23–32. ACM, New York, NY, USA (2010)
15. Gonen, M., Gonen, R., Pavlov, E.: Generalized trade reduction mechanisms. In: EC. pp. 20–29. ACM, New York, NY, USA (2007)
16. Groves, T.: Incentives in teams. Econometrica **41**, 617–631 (1973)
17. McAfee, R.P.: dominant strategy double auction. Journal of Economic Theory **56**, 434–450 (1992)
18. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. Journal of Economic Theory **29**, 265–281 (1983)
19. Segal-Halevi, E., Hassidim, A., Aumann, Y.: SBBA: A strongly-budget-balanced double-auction mechanism. In: SAGT. pp. 260–272. Springer International Publishing AG, Cham, Switzerland (2016)
20. Segal-Halevi, E., Hassidim, A., Aumann, Y.: Muda: A truthful multi-unit double-auction mechanism. In: Proceeding of AAAI (2018)
21. 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, SC (2014)
22. Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. Journal of Finance **16**, 8–37 (1961)