

Mechanism Design for Personalized Recommender Systems

Qingpeng Cai
Tsinghua University, China
cq14@mails.tsinghua.edu.cn

Aris Filos-Ratsikas
University of Oxford, UK
Aris.Filos-Ratsikas@cs.ox.ac.uk

Chang Liu
Alibaba Group, China
qingsong.lc@alibaba-inc.com

Pingzhong Tang
Tsinghua University, China
kenshinping@gmail.com

ABSTRACT

Strategic behaviour from sellers on e-commerce websites, such as faking transactions and manipulating the recommendation scores through artificial reviews, have been among the most notorious obstacles that prevent websites from maximizing the efficiency of their recommendations. Previous approaches have focused almost exclusively on machine learning-related techniques to detect and penalize such behaviour. In this paper, we tackle the problem from a different perspective, using the approach of the field of *mechanism design*. We put forward a game model tailored for the setting at hand and aim to construct *truthful* mechanisms, i.e. mechanisms that do not provide incentives for dishonest reputation-augmenting actions, that guarantee good recommendations in the worst-case. For the setting with two agents, we propose a truthful mechanism that is optimal in terms of social efficiency. For the general case of m agents, we prove both lower and upper bound results on the efficiency of truthful mechanisms and propose truthful mechanisms that yield significantly better results, when compared to an existing mechanism from a leading e-commerce site on real data.

CCS Concepts

•Information systems → Recommender systems; •Theory of computation → Algorithmic mechanism design; •Applied computing → E-commerce infrastructure;

Keywords

Mechanism design; Reputation systems; Approximation

1. INTRODUCTION

When a buyer signs in an e-commerce website (e.g., Amazon or eBay or Taobao), the website returns a list of recommended product-seller pairs that the buyer might be interested in. This recommendation is usually personalized, i.e. it is based on several factors related to the buyer, such as the buyer's demographic and past

browsing or purchase history. The appropriate choice of product-seller pair to be suggested to a buyer of certain characteristics is selected by a ranking algorithm, which can be thought of as a systematic way to allocate the whole amount of buyer impressions. It is in the platform's best interest to allocate the buyer impressions in a way that yields high click-through rates (CTRs) and high click-conversion rates (CVRs), typically by giving better display slots (i.e., higher rankings on the webpage) to sellers with higher reputation, more historical transactions or those that best match the buyer's characteristics. As a result, all these websites incorporate a *reputation system* (e.g. see [6]) in their designs, that records the sellers' reputation and historical transactions and rewards those with higher scores via their ranking algorithms. We will refer to such scores as *recommendation scores*. A well-designed reputation system encourages sellers to increase their quality of service, and in turn attracts more businesses [18].

It takes time and effort for sellers to build up their reputation; in Amazon for example, some trusted, well-known sellers have accumulated more than one million reviews with positive scores as high as 97%. As a result, as it is also observed often in the industry, dishonest sellers may take a "shortcut" and hire buyers to conduct fake transactions with them as a fast way to accumulate positive feedback and increase their reputation scores and number of historical transactions. The severity of the problem is also highlighted by Amazon's recent lawsuit against sellers that were allegedly using fake reviews to boost their profits.

In fact, there has even been an emerging underground industry that provides sophisticated solutions for the sellers who want to quickly boost their reputations. Xu et al. [20] refer to such enterprises as *seller-reputation-escalation (SRE) markets*.

Current approaches, which are reflected in most of the existing literature [11, 4, 21] aim to tackle the problem by training machine learning predictive models using features of the review texts, to detect and punish fake reviews. However, Ott et al. [15] show that such deceptive statements are not easily identified either by learning algorithms or even by human readers. For example, Amazon recently sued more than one thousand sellers for conducting fake transactions, each of which was involved in several purchases; it is easily conceivable that this is only a small fraction of the number of sellers that employ such reputation-augmenting strategies. Also, in the current design of Taobao, the world's largest e-commerce website in terms of gross volume, according to a third party estimation (which is reinforced by inference from data) even after applying such a manipulation-detection engine, there is still more than 10% of the total Taobao orders that are fake. Finally, such detection

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

RecSys '16, September 15-19, 2016, Boston, MA, USA

© 2016 ACM. ISBN 978-1-4503-4035-9/16/09...\$15.00

DOI: <http://dx.doi.org/10.1145/2959100.2959135>

methods also suffer from the possibility of penalizing honest sellers, decreasing their overall experience of using the website as a platform for their transactions.

A mechanism design approach

In this paper, we aim to tackle this problem from a different perspective, using the tools from the fields of *game theory* and more specifically, *mechanism design*. Game theory is predictive in the sense that it is concerned with what the selfish or *rational* actions of the people involved in a system will lead to. For our problem, the participants or *agents* are the sellers who aim to boost their recommendation scores. The field of mechanism design, which has its roots in the pioneering works of Maskin [9] and Myerson [13] is preventive, in the sense that the rules of the system are designed appropriately, in such a way that selfish behaviour is either completely discouraged or at the very least, it is handled carefully and without severe consequences.

We model the problem described above as a variant of the *resource allocation setting* [Chapters 10,11 from [17]] where the designer (i.e., the platform) has to allocate one unit of a single divisible good. This unit can be interpreted as the number of total impressions of buyers with certain characteristics that have to be allocated among sellers,¹ or the probability that a seller is recommended to a single buyer when the buyer visits the website, or even the fraction of time for which the seller will be suggested to the appropriate buyers over a specified period of time. For example, given the first interpretation, if a seller receives an allocation of $1/3$, it means that he will receive $1/3$ of the total recommendation slots for buyers of a certain kind.

In traditional mechanism design settings, each agent has an associated *type*, which conveys information about the preferences of the agent and is *reported* to the mechanism designer, which then runs the allocation rule or the *mechanism* with the types as inputs. The type does not necessarily contain the true preferences of the agent; if a rational agent can force a better outcome by feigning a fake reported preference, he will do so. Central to the field of mechanism design is the notion of truthfulness, i.e. a guarantee that under any circumstances and regardless of the choices of the other participants, an agent will never have an incentive to report anything but his true type. The preferences of the agents are measured through utility functions [20] and a truthful mechanism ensures that an agent receives the highest possible utility by telling the truth. In fact, ensured by the well-known revelation principle [12], it is without loss of generality to consider only truthful mechanisms.² This is also the reason why we can restrict attention to truthful mechanism design throughout the paper.

Our setting is slightly different from the traditional model, in the sense that the types of the agents are the recommendation scores, which are maintained by the system and are established through the process of carrying out transactions and obtaining positive feedback. One key feature of our model is the *cost of manipulation* for the sellers. Each seller can “report” any possible type, however, he suffers from a cost by misreporting, which is the cost of hiring people to write fake reviews or using services like the SRE mar-

¹We model the total number of impressions as a continuous unit rather than an integer. Considering that in most e-commerce websites, this number is rather large, this is not an unrealistic assumption; in fact our guarantees will hold approximately with very small approximation error for any large number of discrete impressions.

²The revelation principle states that any objective implementable in dominant strategies can be implemented by a truthful mechanism. Other commonly used names for truthfulness are *incentive-compatibility* or *staregy-proofness*.

ket mentioned earlier or even the probability of getting caught and being penalized. We model the cost function to be explicitly correlated with the distance of manipulation (reported type minus the true type) as well as the current value of the true type. It is natural to assume that the higher one’s current reputation, the harder the manipulation, especially since it involves the risk of detection; being severely penalized or being removed from the market might be catastrophic for a highly-respected seller.

A seller’s utility is the difference between how much he values the current allocation and his cost if he chooses to manipulate. We note here that while our utility functions are quasi-linear in the cost, they are different from standard quasi-linear utilities in most of the work in mechanism design; the payment function in the standard quasi-linear settings is imposed exogenously by the mechanism in order to produce good incentives (like the well-known VCG mechanism [19, 3, 1]) whereas here, the cost function is associated with the manipulation only and the allocating mechanisms do not use payments. In that sense, we can view our approach as following the agenda of *approximate mechanism design without money* put forward by Procaccia and Tennenholtz [16].

Mechanism design approaches have been employed in the past in recommendation systems, but most of them [5, 7, 8, 2] are concerned with how to design reputation systems that incentivize buyers to report honest and constructive feedbacks rather than considering sellers as the selfish participants. As a notable exception, Zhang et al. [22] consider both strategic buyers and sellers, but their model incorporates a social network-type graph and reputation systems for both buyers and sellers and is, in most regards, quite different from ours.

Our results

Our goal will be to design truthful mechanisms, i.e. mechanisms that do not encourage sellers to engage in reputation-altering manipulations and at the same time maximize the socially desired outcome, i.e. make sure that buyers receive recommendations for sellers with high recommendation scores. Following the usual mechanism design terminology, we will refer to our objective as the *social welfare*.³ We will measure the performance of truthful mechanism by its *efficiency*, i.e. worst case ratio between the social welfare achieved by the mechanism over the optimal social welfare, achieved by recommending the sellers with the highest recommendation scores, ignoring potential manipulations and strategic behaviour.

Our results can be summarized as follows. For the case of two sellers, we design truthful mechanisms that are optimal among all such mechanisms in terms of efficiency for both *regular* cost functions and general cost functions. For the general case of many sellers, we design two truthful mechanisms. We provide a worst-case guarantee for the efficiency of the first one together with a general upper bound on the efficiency of any truthful mechanism, establishing that under some assumptions on the valuation and cost functions, the efficiency of the mechanism is quite close to the efficiency of the best truthful mechanism. We evaluate both mechanisms on real-life data from Taobao and show that our mechanisms significantly outperform the mechanism that Taobao currently uses. We also observe that the performance of our two mechanisms scales differently with the number of sellers and the choices for the cost

³We note here that usually the social welfare refers to the aggregate happiness of the participating agents. In our case however, although the strategic entities are the sellers, the real welfare objective is the aggregate satisfaction of the buyers which is also aligned with the interests of the e-commerce platform.

and valuation functions, showing that both are useful, for different input parameters and sizes.

2. THE MODEL

In our model, there are m sellers and one divisible unit of impression or *item* to allocate between them. Each seller is associated with a non-negative *recommendation score* v_i which is a function of a seller's reputation and propriety or fitness with respect to a buyer of certain characteristics. Let r_i denote the *recorded* recommendation score of seller i , i.e. the recommendation score that is stored in the platform's database for this seller.⁴ Note that the recorded recommendation score might be different from the real, inherent recommendation score of a seller, since the former might have been acquired through fake transactions. Let r denote the vector of recorded recommendation scores of all sellers and let r_{-i} denote the vector of recorded scores of all sellers besides seller i . We will call r a recommendation score profile. The definitions for the vectors of real recommendation scores are similar.

A *mechanism* f is a function that inputs a recorded score profile r , and outputs an allocation q , i.e. a mapping from r to $q = (q_1, \dots, q_m)$, where $q_i(r)$ denotes the fraction of the item seller i gets, which as mentioned in the introduction, can have different interpretations. Clearly, an allocation f is feasible if and only if $\forall r \forall i, q_i(r) \geq 0$, and $\forall r, \sum_{i=1}^m q_i(r) \leq 1$. Note that any feasible allocation of sellers to ad slots can be realized by an appropriate convex combination of permutations of sellers to those slots.

Each seller i has an intrinsic positive valuation $g(v_i)$ towards receiving the item that denotes how happy the seller would be if he were allocated the whole unit of impression according to his real recommendation score. In general, we model the valuation g to be a positive function that maps recommendation scores to valuations; this allows us to consider cases where the value is positively or negatively correlated with the recommendation score. A natural choice would be to set $g(v_i) = \alpha$ for some constant α , which implies that all sellers would be equally happy if they received the whole unit of buyer impression.⁵

In order to feign a fake recommendation score, the seller has to incur a cost for manipulating.

DEFINITION 1. *The cost for a seller with (real) recommendation score v to obtain a recorded score r is $c(v, r) = |r - v|h(v)$, where $h(v)$ is a positive continuous increasing function.*

Intuitively, the higher a seller's recommendation score, the more costly the manipulation. The form of the cost function also assumes that it linearly depends on the extent to which a seller can increase his recorded score. We will sometimes say that a seller *reports* a recommendation score of r , but it should be understood that he obtains that score through costly manipulations, according to the cost function defined above. Our model assumes that the shape of the cost functions is public information; this is not an unrealistic assumption since the cost of hiring fake reviewers or using SRE services can be calculated or estimated to a high degree. Furthermore it is not hard to see that without any knowledge of the cost

⁴Or more precisely, the vector of scores, since each score depends on a group of buyers of certain characteristics.

⁵In our model, it is implicitly assumed that sellers are indifferent between different advertisement slots on a website. We discuss the added difficulties introduced when considering different values $g(v_i)_j$ for different slots j when it comes to achieving truthfulness in conjunction with efficiency in Section 6; besides, there are several websites (e.g. some major flight-comparing websites) that only have one single advertisement slot or several slots that are not favourable over others.

function, we can not hope to do much in terms of truthfulness. We now define the *utility*⁶ of a seller.

DEFINITION 2. *The utility of seller i with (real) recommendation score v_i when the profile of recorded recommendation scores is r is defined as $u_i(v_i, r) = q_i(r)g(v_i) - c(v_i, r_i)$.*

As we mentioned earlier, it is without loss of generality to restrict attention to *truthful mechanisms*.

DEFINITION 3. *A mechanism f is truthful if for each seller i and for all recorded scores of all other sellers r_{-i} and for each report r_i of seller i it holds that $u_i(v_i, (v_i, r_{-i})) \geq u_i(v_i, (r_i, r_{-i}))$, i.e. the seller does not have any incentives to try to fake his real recommendation score.*

By the definition above, when analyzing a truthful mechanism, we will use v to denote the input to the mechanism, since the recorded scores and the real scores are the same.

Ideally, one would be interested in finding a truthful allocation mechanism which maximizes the social welfare (among all such mechanisms) for every instance of the problem. There are several obstacles to doing this. First, the space of available scores (the type space) is continuous and hence there are infinitely many input instances that one would have to consider. Secondly, if we assume that the recommendation scores come from some discrete set, then one idea would be to adopt a linear programming-based approach where the truthfulness constraint of Definition 3 would be a constraint of the linear program. However, such an approach would require us to write one constraint for each possible pair of scores, resulting in a number of constraints not manageable even for a relatively small number of recommendation scores. Given that in platforms with many sellers, we need to maintain many different possible scores in order to distinguish between them, it does not seem that such an approach would work.

Instead, we will aim to design mechanisms that perform well with respect to all possible inputs. As we will see, the performance of our mechanisms will be limited by the worst-case instances but the experimental evaluation suggests that they perform much better on typical inputs. We define the *efficiency* of a mechanism f as the ratio between its social welfare and the algorithmic optimum (i.e. the best welfare one can achieve without imposing truthfulness constraints) in the worst case, i.e.,

$$E(f) = \min_{v > 0} \frac{\sum_{i=1}^m q_i(v)v_i}{v_{(1)}}$$

We remark here that our efficiency notion is the same one as the approximation ratio for truthful mechanisms used in the literature of algorithmic mechanism design [14, 16]. We will be interested in designing mechanisms that have the maximum efficiency among all truthful mechanisms.

3. MECHANISMS FOR 2 SELLERS

In order to explain our approach better, we will start from the design of truthful mechanisms for two sellers; this will allow us to demonstrate some of the concepts of mechanism design in a simpler environment plus, the mechanisms that we will present in the following section for the general case of m sellers will be very similar in spirit. We will first consider the case of *regular cost functions* and then extend our analysis to the case of general cost functions.

⁶We remark here that while the cost function is *single-peaked* [10] in the recommendation score, the dependence of the utility on the allocation and the cost might give rise to more complicated structures.

DEFINITION 4. A cost function c is regular if $h(v)/g(v)$ is non-increasing and integrable, i.e. if we let $H(v) = \int h(v)/g(v)dv$, then c is regular if $H(v)$ is concave.

By the definition above, since $H'(v) = h(v)/g(v) > 0$, it holds that $H(v)$ is an increasing function.

Optimal mechanism for regular cost functions

In this section we present a truthful mechanism for two sellers, prove that it is optimal among all truthful mechanisms for the case of *regular cost functions* and actually the mechanism is derived from the deduction in the proof. The mechanism is the following one.

MECHANISM 1. Consider the recommendation score profile v and let v_l denote the larger value, and v_s denote the smaller value. Let $q_j(v)$ be the allocation of the seller with value v_j for $j \in \{l, s\}$. Then allocate the item as follows:

$$q_l(v) = \min \left\{ H(v_l) - H(v_s) + \frac{1}{2}, 1 \right\}, \quad q_s(v) = 1 - q_l(v)$$

As an example, when $v_l = v_s$, $q_l(v) = q_s(v) = 1/2$ and each seller receives half of the total impression. It is easy to see that the mechanism is feasible, and the intuition why this mechanism is truthful is that manipulations are not desirable because H is a concave function. We formalize this intuition in the following theorem.

THEOREM 1. Mechanism 1 is truthful.

PROOF. Without loss of generality, (by symmetry), we prove that seller 1 does not have an incentive to report a fake score, given an arbitrary recorded score of seller 2. Clearly, seller 1 has no incentive to report a score smaller than his real score because he will receive a smaller fraction of the item in that case and therefore we consider two cases.

Case 1. $v_1 \leq r_2$: The utility of seller 1 by reporting truthfully is

$$u_1(v_1, (v_1, r_2)) = \left(\max \left\{ -H(r_2) + H(v_1) + \frac{1}{2}, 0 \right\} \right) g(v_1).$$

If seller 1 reports r_1 such that $v_1 \leq r_1 \leq r_2$, then his utility $u_1(v_1, (r_1, r_2))$ is

$$\left(\max \left\{ -H(r_2) + H(r_1) + \frac{1}{2}, 0 \right\} \right) g(v_1) - (r_1 - v_1)h(v_1).$$

We have that the difference in utility $\delta_u = u_1(v_1, (r_1, r_2)) - u_1(v_1, (v_1, r_2))$ is at most

$$\delta_u \leq (H(r_1) - H(v_1))g(v_1) - (r_1 - v_1)h(v_1),$$

$H(v)$ is a concave function, and the derivative of it is $h(v)/g(v)$. By concavity and by the inequality above, we get that $\delta_u \leq 0$. If seller 1 chooses to report r_1 such that $r_1 > r_2$, then his utility $u_1(v_1, (r_1, r_2))$ is

$$\left(\min \left\{ H(r_1) - H(r_2) + \frac{1}{2}, 1 \right\} \right) g(v_1) - (r_1 - v_1)h(v_1).$$

Then, combining the formulas for the utility of truthful reporting and of misreport r_1 , we obtain again that the difference in utility $\delta'_u = u_1(v_1, (r_1, r_2)) - u_1(v_1, (v_1, r_2))$ is at most

$$\delta'_u \leq (H(r_1) - H(v_1))g(v_1) - (r_1 - v_1)h(v_1) \leq 0.$$

Case 2. $v_1 > r_2$. The utility of seller 1 by reporting truthfully is

$$u_1(v_1, (v_1, r_2)) = \left(\min \left\{ -H(r_2) + H(v_1) + \frac{1}{2}, 1 \right\} \right) g(v_1).$$

If seller 1 reports r_1 such that $v_1 < r_1$, his utility $u_1(v_1, (r_1, r_2))$ is

$$\left(\min \left\{ -H(r_2) + H(r_1) + \frac{1}{2}, 1 \right\} \right) g(v_1) - (r_1 - v_1)h(v_1).$$

Similarly to before, we get that the difference in utility $\delta_u = u_1(v_1, (r_1, r_2)) - u_1(v_1, (v_1, r_2))$ is at most

$$\delta_u \leq (H(r_1) - H(v_1))g(v_1) - (r_1 - v_1)h(v_1) \leq 0.$$

This completes the proof. \square

In the following, we will prove the worst-case efficiency guarantee of Mechanism 3 and the fact that it is optimal among all truthful mechanisms. For the latter part, we will need the next lemma, that provides a necessary condition for a mechanism to be truthful.

LEMMA 1. Let f be a truthful mechanism. It holds that

$$- \text{for all } v_1 \geq v_2, q_1(v_1, v_2) \leq q_1(v_2, v_2) + H(v_1) - H(v_2).$$

$$- \text{for all } v_2 \geq v_1, q_2(v_1, v_2) \leq q_2(v_1, v_1) + H(v_2) - H(v_1).$$

PROOF. By symmetry, we only prove the first statement of the lemma. If seller 1 has score v_2 and seller 2 reports score r_2 , we have that for any $\delta > 0$ it must hold that

$$(q_1(v_2 + \delta, r_2) - q_1(v_2, r_2))g(v_2) \leq \delta h(v_2),$$

otherwise seller 1 will have an incentive to misreport $v_2 + \delta$, i.e.,

$$\forall \delta > 0, v_2 > 0 : \frac{q_1(v_2 + \delta, r_2) - q_1(v_2, r_2)}{\delta} \leq \frac{h(v_2)}{g(v_2)}.$$

Because the cost function c is regular, the function $h(v)/g(v)$ is integrable, and hence $q_1(v_1, v'_2) - q_1(v_2, v'_2) \leq \int_{v_2}^{v_1} h(v)/g(v)dv$, i.e., $\forall v_1 \geq v_2, q_1(v_1, v_2) \leq q_1(v_2, v_2) + H(v_1) - H(v_2)$. \square

We are now state the following theorem. The proof is omitted due to lack of space; we refer the reader to the full version.

THEOREM 2. Let

$$E(H) = \min_{v_1 \geq v_2} \left(\frac{v_2}{v_1} + \frac{v_1 - v_2}{v_1} \left(\frac{1}{2} + H(v_1) - H(v_2) \right) \right) \quad (1)$$

The efficiency of Mechanism 1 is $E(M_1) = \min \{E(H), 1\}$, which is optimal among all truthful mechanisms.

Optimal mechanism for general cost functions

In this section we present a truthful mechanism with two sellers and general cost functions, and prove it is optimal among all truthful mechanisms. The idea is to extend the idea we used in the previous section and find a decreasing function that is below $h(v)/v$ that is “as large as possible”. For ease of reference, we will use say that a function g_1 is *not larger than* function g_2 if for all $v \in (0, +\infty)$ it holds that $g_1(v) \leq g_2(v)$. Given the cost function $c(v, r)$, for all $v > 0$, we define the function

$$h_1(v) = \min_{0 < t \leq v} \frac{h(t)}{g(t)}.$$

Let $H_1(v) = \int h_1(v)dv$ and hence $H_1(v)$ is a concave function.

Note that $h_1(v)$ is a decreasing function not larger than $h(v)/g(v)$, and moreover, it holds that $h_1(v) = h(v)/g(v)$ when $h(v)/g(v)$ is decreasing. The following lemma states that function $h_1(v)$ is the largest decreasing function which is not larger than $h(v)/g(v)$.

LEMMA 2. For any decreasing function $h_2(v)$ not larger than $h(v)/g(v)$, for all $v > 0$, it holds that $h_2(v) \leq h_1(v)$.

PROOF. For any score v_1 such that $0 < v_1 \leq v$, we have that $h_2(v) \leq h_2(v_1) \leq h(v_1)/g(v_1)$, i.e.

$$h_2(v) \leq \min_{0 < t \leq v} \frac{h(t)}{g(t)} = h_1(v).$$

□

Using the concave function $H_1(v)$ defined above and the same intuition of the design of Mechanism 1, we obtain the following optimal truthful mechanism for general cost functions.

MECHANISM 2. Consider the recommendation score profile v and let v_l denote the larger value, and v_s denote the smaller value. Let $q_j(v)$ be the allocation of the seller with value v_j for $j \in \{l, s\}$. Then allocate the item as follows:

$$q_l(v) = \min \left\{ H_1(v_l) - H_1(v_s) + \frac{1}{2}, 1 \right\}, \quad q_s(v) = 1 - q_l(v).$$

The following theorem establishes the truthfulness of the mechanism.

THEOREM 3. Mechanism 2 is truthful.

The proof of Theorem 3 follows from very similar arguments as the ones used in the proof of Theorem 1 (see full version) and the following lemma.

LEMMA 3. Recall that $H_1(v) = \int h_1(v)dv$. For any $0 < v_1 \leq r_1$, we have that

$$H_1(r_1) - H_1(v_1) \leq \frac{h(v_1)}{g(v_1)}(r_1 - v_1).$$

PROOF. Because H_1 is concave, for any $0 < v_1 \leq v'_1$, it holds that $H_1(r_1) - H_1(v_1) \leq h_1(v_1)(r_1 - v_1)$, and furthermore $h_1(v_1) \leq h(v_1)/g(v_1)$, which proves the lemma. □

Finally, the following theorem establishes that Mechanism 2 is optimal among all truthful mechanisms for two agents, for general cost functions.

THEOREM 4. The efficiency of Mechanism 2 is equal to $\min\{E(H_1), 1\}$ which is optimal among all truthful mechanisms.

Again, Theorem 4 can be proved using arguments very similar to those used in the proof of Theorem 2, together with the following lemma. The proof of the lemma is similar to the proof of Lemma 1 and we omit it due to lack of space.

LEMMA 4. Let f be a truthful mechanism. It holds that

- for all $v_1 \geq v_2$, $q_1(v_1, v_2) \leq q_1(v_2, v_2) + H_1(v_1) - H_1(v_2)$.
- for all $v_2 \geq v_1$, $q_2(v_1, v_2) \leq q_2(v_1, v_1) + H_1(v_2) - H_1(v_1)$.

4. MECHANISMS FOR MANY SELLERS

In this section we consider the more general setting where we have m sellers (with $m > 2$) that we are allocating the unit of impression to. Our main contribution of this section is the design of a truthful mechanism whose efficiency (a) approaches optimality among truthful mechanisms when the cost of manipulation approaches 0 and (b) is strictly better than the obvious truthful mechanism, that allocates the unit uniformly to the sellers. As we will see in Section 5, in typical instances of the problem, our mechanism will significantly outperform the uniform allocation, which makes the extra effort of analyzing its properties clearly justified.

MECHANISM 3. Let i be the seller with the highest recommendation score, i.e. $i = \arg \max_j v_j$ and j be the seller with the second highest recommendation score. The allocation of is:

$$q_i(v) = \min \left\{ \frac{1}{m} + (H(v_i) - H(v_j)), 1 \right\} \text{ and}$$

$$q_k(v) = \frac{1 - q_i(v)}{m - 1} \text{ for all } k \neq i.$$

It is easy to check that the mechanism is feasible. Note that Mechanism 3 is in fact a generalization of Mechanism 1; the fraction of the item that the seller with the highest score receives is determined by the difference between the highest score and the second highest score, and other sellers split the remainder of the item evenly.

For ease of exposition, we will analyze the mechanism in the setting of regular cost functions. We can obtain similar mechanisms with analogous efficiency guarantees for general cost functions by using similar lemmas as the ones that we employed in Section 3.

THEOREM 5. Mechanism 3 is truthful.

PROOF. Again, without loss of generality, it suffices to prove that seller 1 does not have an incentive to misreport his recommendation score. Obviously, seller 1 has no incentive to report a score smaller than v_1 because then, he will get a smaller fraction of the item. Let $r = (v_1, r_2, \dots, r_m)$, and $r' = (r_1, r_2, \dots, r_m)$. We consider three cases.

Case 1 $v_1 \geq r_i$ for all $i \neq 1$: The utility of seller 1 by telling the truth is

$$u_1(v_1, r) = \min \left\{ \frac{1}{m} + H(v_1) - H(r_2), 1 \right\} g(v_1).$$

If seller 1 reports a score r_1 such that ($r_1 > v_1$) to report, his utility $u_1(v_1, r')$ becomes

$$\min \left\{ \frac{1}{m} + H(r_1) - H(r_2), 1 \right\} g(v_1) - (r_1 - v_1)h(v_1).$$

and for the difference in utility $\delta_u = u_1(v_1, r') - u_1(v_1, r)$ it holds that $\delta_u \leq (H(r_1) - H(v_1))g(v_1) - (r_1 - v_1)h(v_1) \leq 0$.

Case 2. $r_2 > v_1 \geq r_3 \geq \dots \geq r_m$ Seller 1's utility from telling the truth

$$u_1(v_1, r) = \frac{\max \{0, (m-1)/m + H(v_1) - H(r_2)\}}{m-1} g(v_1).$$

If seller 1 reports r_1 such that $r_2 > r_1 > v_1$, his utility $u_1(v_1, r')$ becomes

$$\frac{\max \{0, (m-1)/m + H(r_1) - H(r_2)\}}{m-1} g(v_1) - (r_1 - v_1)h(v_1).$$

and for the difference in utility $\delta_u = u_1(v_1, r') - u_1(v_1, r)$ it holds that $\delta_u \leq (H(r_1) - H(v_1))g(v_1)/(m-1) - (r_1 - v_1)h(v_1) \leq 0$. If seller 1 reports r_1 such that $r_1 \geq r_2$, his utility $u_1(v_1, r')$ is

$$\min \left\{ \frac{1}{m} + H(r_1) - H(r_2), 1 \right\} g(v_1) - (r_1 - v_1)h(v_1).$$

and the difference $\delta_u = u_1(v_1, r') - u_1(v_1, r)$ is at most

$$\left(\frac{H(r_1) - H(v_1)}{(m-1)} - \frac{(m-2)H(r_2)}{(m-1)} \right) g(v_1) - (r_1 - v_1)h(v_1) \leq (H(r_1) - H(v_1))g(v_1) - (r_1 - v_1)h(v_1) \leq 0.$$

Case 3. $r_2 \geq r_3 > v_1$. Seller 1's utility from telling the truth is

$$u_1(v_1, r) = \frac{\max \{0, (m-1)/m + H(r_3) - H(r_2)\}}{m-1} g(v_1).$$

By the construction of the Mechanism, seller 1 can only affect the allocation outcome only when his reported score is the highest score or the second highest score. If he reports r_1 such that $r_2 > r_1 > r_3$, his utility $u_1(v_1, r')$ becomes

$$\frac{\max\{0, (m-1)/m + H(r_1) - H(r_2)\}}{m-1} g(v_1) - (r_1 - v_1)h(v_1).$$

and the difference in utility $\delta_u = u_1(v_1, r') - u_1(v_1, r)$ is at most

$$\begin{aligned} & \frac{(H(r_1) - H(r_3))g(v_1)}{(m-1)} - (r_1 - v_1)h(v_1) \\ & \leq \frac{(H(r_1) - H(v_1))g(v_1)}{(m-1)} - (r_1 - v_1)h(v_1) \leq 0. \end{aligned}$$

If he reports r_1 such that $r_1 > r_2 \geq r_3$, his utility becomes

$$\min\left\{\frac{1}{m} + H(r_1) - H(r_2), 1\right\} g(v_1) - (r_1 - v_1)h(v_1).$$

and the difference $\delta_u = u_1(v_1, r') - u_1(v_1, r)$ is at most

$$\begin{aligned} & \left(\frac{H(r_1) - H(r_3)}{(m-1)} - \frac{(m-2)H(r_2)}{(m-1)}\right) g(v_1) - (r_1 - v_1)h(v_1) \\ & \leq (H(r_1) - H(v_1))g(v_1) - (r_1 - v_1)h(v_1) \leq 0. \end{aligned}$$

This concludes the proof of the theorem. \square

The following theorem establishes the efficiency guarantee of Mechanism 3. Note that this a theoretical lower bound, over all possible instances of the problem. We will see in Section 5 that on real-life instances, Mechanism 3 outperforms the worst-case bound considerably. We omit the proof of the lemma due to lack of space.

THEOREM 6. *Let $a = h(0)/g(0)$. The efficiency of Mechanism 3 is at least*

$$\min\left\{\frac{2}{m}, \frac{1}{m} + \frac{m-2}{m-1} \cdot a\right\}.$$

An upper bound for all truthful mechanisms

We conclude the section with an upper bound on the efficiency of any truthful mechanism. To prove the bound, similarly to the machinery needed for the proof of Lemma 1, we firstly obtain a necessary condition for truthfulness for the case of m sellers. Again, we omit the proof due to lack of space.

LEMMA 5. *Let f be a truthful mechanism. For every $v_1 \geq v_2$,*

$$q_1(v_1, v_2, \dots, v_2) \leq q_1(v_2, v_2, \dots, v_2) + H(v_1) - H(v_2).$$

Using Lemma 5 we prove the following theorem. Notice the implications of the theorem; when $c = 0$, in which case there is no cost for manipulating, then the best thing that we can hope for with truthful mechanisms is a uniform allocation between sellers. As c goes to infinity the incentive for manipulation is too small and truthful mechanisms that approach algorithmic optimality are possible. Furthermore, notice that if the functions $h(\cdot)$ and $g(\cdot)$ are “smooth” enough (for example when h and g are constant functions) and their values at 0 are small enough the efficiency guarantee of Mechanism 3 is very close to that of the best possible truthful mechanism.

THEOREM 7. *Let $c = h(1)/g(1)$. The efficiency of any truthful mechanism f is at most $\frac{1}{m} + 2\sqrt{c\left(c + \frac{m-1}{m}\right)} - 2c$.*

PROOF. Assume without loss of generality that $v_1 = \max_i v_i$. By definition, $E(f) = \min_{v>0} (\sum_{i=1}^m q_i(v)v_i)/v_1$. Consider the profiles $v = (v_1, v_2, \dots, v_2)$ and $v' = (v_2, v_2, \dots, v_2)$. It holds that

$$E(f) \leq \min_{v_1 \geq v_2 > 0} \left(\frac{v_2}{v_1} + \frac{(v_1 - v_2)}{v_1} q_1(v)\right).$$

By Lemma 5, we have that $q_1(v) \leq q_1(v') + H(v_1) - H(v_2)$. Then

$$E(f) \leq \min_{v_1 \geq v_2 > 0} \left(\frac{v_2}{v_1} + \frac{(v_1 - v_2)}{v_1} (q_1(v') + H(v_1) - H(v_2))\right).$$

Similarly, for all sellers i it holds that

$$E(f) \leq \min_{v_1 \geq v_2 > 0} \left(\frac{v_2}{v_1} + \frac{(v_1 - v_2)}{v_1} (q_i(v') + H(v_1) - H(v_2))\right).$$

By feasibility, we have that $\sum_{i=1}^m q_i(v') \leq 1$ and therefore

$$E(f) \leq \min_{v_1 \geq v_2 > 0} \left(\frac{v_2}{v_1} + \frac{(v_1 - v_2)}{v_1} \left(\frac{1}{m} + H(v_1) - H(v_2)\right)\right).$$

Letting $a = v_1/v_2$, the right-hand side of the inequality above can be written as

$$\min_{v_1 > 0, a \geq 1} \left(\frac{1}{a} + \frac{a-1}{a} \left(\frac{1}{m} + H(v_1) - H(v_1/a)\right)\right).$$

Now observe that

$$H(v_1) - H\left(\frac{v_1}{a}\right) \leq \left(v_1 - \frac{v_1}{a}\right) \frac{h(v_1/a)}{g(v_1/a)},$$

and hence the efficiency can be upper bounded as

$$E(f) \leq \min_{v_1=a, a \geq 1} \left(\frac{1}{a} + \frac{a-1}{a} \left(\frac{1}{m} + (a-1)c\right)\right),$$

which implies the bound of the theorem. \square

5. EXPERIMENTS

Up until now, we have been discussing the worst-case theoretical guarantees of the mechanisms that we designed. In this section, we will evaluate the performance of our mechanisms empirically, using real data from Taobao, the primary online marketplace in China and one of the biggest e-commerce websites in the world. In particular, because the number of sellers and buyers in Taobao is very large, we gather information about the transactions and buyers’ data from 2047 randomly sampled sellers with respect to buyers of a certain demographic (female buyers, of ages between 20 and 30) that occurred within the past year. The number of transaction orders after deleting buyers that have been detected to fake transactions in this dataset is 11599033, thus we think doing experiments on this dataset is without loss of generality.

As we explained in the model section, we will interpret the item as the total number of buyer impressions for this buyer category. The recommendation scores of the sellers are calculated as follows. First, for each seller, we calculate the number of transaction orders he could have made if he were allocated all buyer impressions for this buyer category by machine learning methods and then, we scale these numbers appropriately to make sure they lie in the range $[0, 1]$; the latter is just a convention but it is in accordance with usual conventions in reality, where the scores are usually % percentages. The social welfare achieved by a mechanism is the total (normalized) number of transactions resulted from the impressions being allocated by to the mechanism.

As we mentioned earlier, there exist fake transactions in the input data that need to be taken into consideration. For this reason, first we compose a “blacklist” of buyers that have been detected to fake transactions in the past by the Alibaba group, the company that owns Taobao. Then, we remove these fake orders from the input data and estimate the real recommendation score for each seller. As a result, we can construct a data generator D for any seller in Taobao with associated recommendation scores, i.e, the distribution of real recommendation score is uniformly drawn from the real scores of 2047 sellers.

For our experiments, we will evaluate the efficiency of Mechanism 3, the mechanism used by Taobao and the following mechanism which we can show to be truthful (proof omitted due to lack of space).

MECHANISM 4. Recall the definition of H . Let $H(0) = 1/m$. Seller i is allocated the fraction $q_i(v) = H(v_i) / \left(\sum_{j=1}^m H(v_j) \right)$.

Note that while we do not provide a worst-case lower bound for the efficiency of Mechanism 4, as we will see shortly, the mechanism actually performs very well on the real-life inputs that we generate.

The recommendation algorithm that Taobao uses works as follows: when a single buyer visits the system, the algorithm ranks sellers according to their recommendation scores associated to buyers of the visitor’s characteristics. Then, it picks a certain number of sellers from the top of the ranking and suggests these sellers to the buyer. Unfortunately, the exact allocation rule to the selected sellers based on their scores is not public information. However, we can infer the allocation rule using machine learning methods from the data, and we can simulate the Taobao mechanism with input scores of any sellers by this rule; since we are interested in this particular data set, our implementation will be an accurate approximation.

We compare the mechanisms for different sample sizes. For each sample size m , we first use our data generator to generate recommendation scores of artificial sellers, i.e, the score of each seller is i.i.d drawn from D . We then compare our two mechanisms against the Taobao mechanism, as well as the uniform mechanism that gives each seller a $1/m$ fraction of the item; the later comparison is useful to demonstrate that although the worst-case bounds of Mechanism 3 are comparable to those of the uniform mechanism, in reality, Mechanism 3 significantly outperforms the uniform allocations. We repeat those experiments 3000 times and we calculate the average efficiency (i.e. the average ratio between the social welfare of the mechanisms and the welfare of the algorithmic optimal).

We consider different choices for the valuation functions $g(v)$ and the function $h(v)$ associated with the cost function. Figure 1 shows the comparisons of the average efficiency as a function of the sample size, for the case when $g(v) = 1$ and $h(v) = 1$. This simple case corresponds to the assumption that all sellers would be equally satisfied by receiving the whole amount of impressions ($g(v) = 1$) and that a sellers’ cost is simply the distance between his recorded score and his real score. As we can see, both of our mechanisms outperform Taobao’s algorithm for up to 4000 sellers and Mechanism 4 is actually much better for all input sizes we consider. Both Mechanism 3 and Mechanism 4 outperform the uniform allocation by a lot although, as the number of sellers grows large, Mechanism 3 seems to converge to its theoretical guarantee.

Figure 2 makes the same assumption on the valuation function $g(v)$ but assumes a “steep” scaling of the cost function, to model instances where the cost of manipulation is much lower than the valuation of a seller. In this case, Mechanism 3 exhibits a poorer performance; on the other hand, Mechanism 4 performs exceptionally well, relatively to the other mechanisms.

The high performance of Mechanism 4 can be explained by the fact that the functions $g(v)$ and $h(v)$ are constant functions, and the mechanism performs better when “smooth” enough functions are considered. To demonstrate this even more clearly, we consider the case where $g(v) = v$ and $h(v) = \alpha v + \alpha$ where α is either 1 or 10; the results of the experiments are summarized in Figure 3. As we can see, in the first case Mechanism 3 outperforms Mechanism 4 for input sizes up to roughly 3000 sellers and the Taobao algorithm for sizes up to slightly less than 6000 sellers. Note that in the second case, when the function $h(v)$ is more “steep”, Mechanism




Figure 1: The average efficiency of the four mechanisms with $h(v) = 1$ and $g(v) = 1$.




Figure 2: The average efficiency of the four mechanisms with $h(v) = 10^{-4}$ and $g(v) = 1$.

4 performs worse and is outperformed by Mechanism 3 for input sizes up to roughly 6000 sellers. It is also worth pointing out that both mechanisms outperform Taobao’s algorithm, which seems to not scale well with quickly increasing functions $h(v)$ either.

Based on our observation, we can draw the following conclusions for real-life instances: (a) Both of our mechanisms outperform Taobao’s algorithm for the most part and the (truthful) uniform mechanism on all occurrences and (b) Mechanism 4 seems to be much better in terms of scaling with the number of sellers when the model functions are relatively “smooth” and Mechanism 3 is preferable when these functions are quickly increasing. Overall, the experiments seem to be indicative that the truthful mechanisms we present can yield better results than existing approaches.

6. CONCLUSION AND FUTURE WORK

There are many interesting directions for future work. As we mentioned earlier, we assume that the value $g(v)$ denotes the satisfaction that a seller would experience for receiving any slot on a webpage. This is a fair assumption in many personalized recommender systems with limited ad-slots for instance, but ideally, we would like to assume that sellers also have preferences over the different slots. Then, we would have to model each slot as a different item j and each seller would have a different value v_{ij} for each one of them. However, this multi-dimensionality introduces added dif-




Figure 3: The average efficiency of the four mechanisms with $h(v) = v+1$ and $g(v) = v$ and $h(v) = 10v+10$ and $g(v) = v$.

difficulties in the design of truthful mechanisms with good efficiency guarantees. The task is challenging but certainly worth exploring.

Another interesting future direction is to impose certain allocative constraints on the amount of impressions. Currently, we assume that the social optimum would be to give the unit to the seller with the highest score; it seems natural to assume that the platform imposes some *fairness* constraints as well, making sure that at least respected sellers receive a certain fraction of an impression. The challenge then would be to design truthful mechanisms that obey the allocative constraints (which can render the quest for truthfulness quite harder) and those mechanisms would be compared to the best possible allocation among those that respect the constraints.

Finally, since Mechanism 4 seems to outperform the other mechanisms significantly, at least for the case of “smooth” functions, a future task would be to provide a worst-case theoretical guarantee on the efficiency of the mechanism.

Acknowledgments

We thank Yao Lu from Alibaba Group for the advice on revising the paper and her contribution to the experimental section. Qingpeng Cai and Pingzhong Tang were supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the NSFC Grant 61033001, 61361136003, 61303077, a Tsinghua Initiative Scientific Research Grant and a National Youth 1000-talent program. Aris Filos-Ratsikas was supported by the ERC Advanced Grant 321171 (ALGAME) and acknowledges support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which this work was performed and the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council.

7. REFERENCES

- [1] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 2:19–33, 1971.
- [2] B. Faltings. Using incentives to obtain truthful information. In *Agents and Artificial Intelligence*, pages 3–10. Springer, 2013.
- [3] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.
- [4] N. Jindal and B. Liu. Opinion spam and analysis. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, pages 219–230. ACM, 2008.
- [5] S. Johnson, J. W. Pratt, and R. J. Zeckhauser. Efficiency despite mutually payoff-relevant private information: The finite case. *Econometrica: Journal of the Econometric Society*, pages 873–900, 1990.
- [6] R. Jurca and B. Faltings. Obtaining reliable feedback for sanctioning reputation mechanisms. *Journal of Artificial Intelligence Research (JAIR)*, 29:391–419, 2007.
- [7] R. Jurca and B. Faltings. Truthful opinions from the crowds. *ACM SIGecom Exchanges*, 7(2):3, 2008.
- [8] R. Jurca, B. Faltings, et al. Mechanisms for making crowds truthful. *Journal of Artificial Intelligence Research*, 34(1):209, 2009.
- [9] E. S. Maskin. Mechanism design: How to implement social goals. *The American Economic Review*, pages 567–576, 2008.
- [10] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [11] A. Mukherjee, A. Kumar, B. Liu, J. Wang, M. Hsu, M. Castellanos, and R. Ghosh. Spotting opinion spammers using behavioral footprints. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 632–640. ACM, 2013.
- [12] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [13] R. B. Myerson. *Mechanism design*. Center for Mathematical Studies in Economics and Management Science, Northwestern University, 1988.
- [14] N. Nisan, T. Roughgarden, Éva Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [15] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 309–319. Association for Computational Linguistics, 2011.
- [16] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 177–186. ACM, 2009.
- [17] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game theoretic and Logical Foundations*. Cambridge Uni. Press, 2009.
- [18] G. Swamyathan, K. C. Almeroth, and B. Y. Zhao. The design of a reliable reputation system. *Electronic Commerce Research*, 10(3-4):239–270, 2010.
- [19] W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37, 1961.
- [20] H. Xu, D. Liu, H. Wang, and A. Stavrou. E-commerce reputation manipulation: The emergence of reputation-escalation-as-a-service. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1296–1306. International World Wide Web Conferences Steering Committee, 2015.
- [21] K.-H. Yoo and U. Gretzel. Comparison of deceptive and truthful travel reviews. *Information and communication technologies in tourism 2009*, pages 37–47, 2009.
- [22] J. Zhang, R. Cohen, and K. Larson. Combining trust modeling and mechanism design for promoting honesty in e-marketplaces. *Computational Intelligence*, 28(4):549–578, 2012.