

Multi-Unit Auctions: Beyond Roberts

Shahar Dobzinski^{*}
Department of Computer Science
Cornell University, Ithaca NY 14853
shahar@cs.cornell.edu

Noam Nisan[†]
School of Computer Science and Engineering
Hebrew University
noam@cs.huji.ac.il

ABSTRACT

We exhibit incentive compatible multi-unit auctions that are not affine maximizers (i.e., are not of the VCG family) and yet approximate the social welfare to within a factor of $1 + \epsilon$. For the case of two-item two-bidder auctions we show that these auctions, termed Triage auctions, are the only scalable ones that give an approximation factor better than 2. “Scalable” means that the allocation does not depend on the units in which the valuations are measured. We deduce from this that any scalable computationally-efficient incentive-compatible auction for m items and $n \geq 2$ bidders cannot approximate the social welfare to within a factor better than 2. This is in contrast to arbitrarily good approximations that can be reached under computational constraints alone, and in contrast to the existence of incentive-compatible mechanisms that achieve the optimal allocation.

Categories and Subject Descriptors

F.2.8 [Analysis of Algorithms and Problem complexity]: Miscellaneous

General Terms

Theory

Keywords

Multi-Unit Auctions, Incentive Compatibility

^{*}Supported by an Alfred P. Sloan Foundation Fellowship and a Microsoft Research New Faculty Fellowship.

[†]Supported by a grant from the Israeli Science Foundation, and by the Google Inter-university center for Electronic Markets and Auctions.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'11, June 5–9, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0261-6/11/06 ...\$10.00.

1. INTRODUCTION

Background

The field of *Algorithmic Mechanism Design* [27] designs mechanisms for achieving various computational goals, under the assumption of rational selfishness of the involved parties. The notions used are taken from the economic field of Mechanism Design, and a basic notion is that of *incentive compatibility* – where rational players are motivated to act *truthfully*. For background and survey see part II of [28]. This paper will consider only the simplest and most robust notion of incentive compatibility, that of dominant strategies in quasi-linear settings with independent private values. The typical question in the field asks for a computationally-efficient incentive compatible mechanism that implements a certain type of outcome, usually the approximate optimization of some target “social” goal. There are two variants of this challenge, the first considers situations where incentive compatibility itself is hard to achieve and the computational efficiency is just an additional burden, with the prime example being approximate minimization of the makespan in scheduling problems [27]. The second variant focuses on cases where each of the two constraints of incentive compatibility and computational efficiency can be achieved separately, and the challenge is to get them simultaneously, with the prime example being approximate welfare maximization in various types of combinatorial auctions [24].

While there has been much work and some progress on these types of challenges, with particular emphasis on the problems mentioned above of combinatorial auctions (e.g., [22, 20, 3, 15, 23, 16, 13, 6]) and scheduling (e.g., [7, 21, 2]), the basic challenge remains mostly unanswered. As noted in [22], the main issue turns out to be the richness of the domain of player’s valuations, i.e., of their private information. On one extreme are single-dimensional domains where the private information of each participant is captured by a scalar (or domains very close to it, e.g., [24]). For these types of problems, incentive-compatible mechanisms are well characterized by a certain monotonicity condition and, in most cases, the challenge of reconciling incentives with computational efficiency has been met [24, 1, 5, 9, 8]. On the other extreme are problems which are “fully dimensional” (or close to fully dimensional, e.g., [29, 17]) where there is no structure on valuations, in which case a key theorem of Roberts [30] characterizes incentive compatible mechanisms as “affine maximizers” “on a sub-range” – simple generalizations of the VCG mechanism. While such affine maximizers on a sub-range are not completely powerless in polynomial time, in

most cases this characterization implies impossibility of good computationally efficient truthful mechanisms. Most interesting problems, including those mentioned above, lie in an intermediate range where the valuation spaces are neither single dimensional nor fully dimensional, a range for which very little is known. The main problem seems to be the lack of a good characterization of incentive compatibility in these intermediate ranges¹. In particular, the key unknown is whether any useful truthful non-VCG mechanisms exist in the intermediate range².

Multi-unit Auctions

As mentioned, the paradigmatic problems for the reconciliation of computational constraints with incentive constraints are the various subclasses of combinatorial auctions. In this paper we consider the simplest variant which exhibits this tension: multi-unit auctions. In this problem there are m identical items for sale among n bidders, where each bidder i has a valuation function $v_i : \{0..m\} \rightarrow \mathbb{R}$, where $v_i(k)$ denotes player i 's value for receiving k items. The valuations v_i are assumed to be monotone non-decreasing (free disposal) with $v_i(0) = 0$ (normalization). Key and implicit here is that there are no externalities: the value of bidder i depends only on what he gets rather than also on the allocation to the others. The optimization goal is to find an allocation of items to the bidders, where bidder i gets s_i items, with $\sum_i s_i \leq m$, that maximizes social welfare $\sum_i v_i(s_i)$.

The problem becomes computationally challenging when the number of items m is “large”, i.e., when the running time of the mechanism is not allowed to be polynomial in m but rather just in $\log m$. There are two variant models in this case, the first assumes that the valuation functions are given as “black boxes” that the algorithm may query³, and the second assumes that the valuation functions are given in some succinct bidding language. Finding the optimal allocation is essentially a knapsack problem and is computationally hard in both models: in the black-box model it requires exponentially many queries, and in the succinct representation model, it is NP-hard. Just like Knapsack, the optimal social welfare can be approximated arbitrarily well (in both models) and has an FPTAS: approximation ratio of $1+\epsilon$ obtained in time that is polynomial in n , $\log m$, and ϵ^{-1} . This FPTAS does not imply any incentive compatible approximation though, and the question boils down to what degree of approximation can be obtained in an incentive compatible way in polynomial time.

Already in Vickrey’s seminal paper [32] multi-unit auctions were considered, restricted to the case of downward sloping valuations, i.e., $v_i(k+1) - v_i(k) \leq v_i(k) - v_i(k-1)$ for all $0 < k < m$. For this case the optimal allocation can be found efficiently, as an “equilibrium price” exists, which can be found by binary search (together with the optimal allocation it implies), and attaching the Vickrey payments – the point of his paper – gives incentive compatibility. For general valuations the exact optimum is computationally hard

¹The “weak monotonicity” [4, 31] characterization is from the point of view of a single player and is not specific enough to be useful in this regard.

²With a single positive exception for certain multi-unit combinatorial auctions [3].

³The usual query assumed is a “value query”, asking for $v_i(k)$ for some k , but most lower bounds hold for any queries, as they apply in the communication complexity model.

to achieve, so further work considered approximations. The single-dimensional “single minded” case was shown to have a truthful FPTAS [5], improving an earlier 2-approximation [25]. In addition, a PTAS exists for somewhat richer valuations that can be described using certain “bidding languages” (e.g., k -minded bidders for a fixed k) [14]. The general case was studied in [14] where a truthful 2-approximation was obtained using a maximal-in-range VCG mechanism⁴. It was also shown there that no computationally-efficient maximal in range VCG mechanism can achieve a better approximation ratio.

The main open problem was whether there exist *non-VCG* truthful mechanisms with a better approximation ratio. Some evidence [22, 17] supported the conjecture that that there are no such mechanisms: truthful mechanisms *for two players that always allocate all items* must be affine maximizers⁵. It should be emphasized that the question regards deterministic mechanisms, as a randomized FPTAS was obtained by [12]⁶.

Our Results

Given the evidences and our own intuition, we were surprised to find that there are non-VCG mechanisms that provide arbitrarily good approximation ratios:

Theorem: For every $\epsilon > 0$, there exists a truthful $(1 + \epsilon)$ -approximation mechanism for multi-unit auctions of m items between two players which is not an affine maximizer.

We call these mechanisms *Triage mechanisms* as they split the valuation domain into three sub-domains, depending on the ratios $v(1)/v(m)$ and $v(m-1)/v(m)$. Their payment structure mimics VCG prices with two exceptions: in the “low sub-domain”, the price for a single item is decreased to a specific fraction of the value of all items, and in the “high sub-domain”, the payments of all non-empty bundles are increased, by the same amount, in a specific linear way. This family of mechanisms is parameterized by three parameters (specifying a weight and the “high thresholds” for both players), with all other parameters uniquely determined by them. We also exhibit two other families of finitely approximating incentive compatible mechanisms, but their approximation factor is worse.

Our next, and main, result shows that these Triage mechanisms are the only *scalable* incentive compatible mechanisms with a good approximation ratio for the case of two items and two players. Scalability means that the auction’s allocation rule does not depend on the “units” in which the valuations are given: multiplying all valuations by the same

⁴The situation in multi-unit auctions mirrors, with different parameters, that of other types of combinatorial auctions where there is a gap between the computationally achievable approximation ratio and the best known computationally efficient truthful mechanism known for the multi-parameter case, which is a maximal-in-range VCG mechanism (see, e.g., [15, 18, 14]).

⁵The driving force of these and similar characterization results is the annulment of the no-externalities condition as everything not won by one player must be given to the other player.

⁶This again mirrors the situation in other types of combinatorial auctions where randomized mechanisms obtain better approximation ratios than those obtained by the known deterministic ones [11, 12].

positive constant does not change the allocation⁷. Triage mechanisms and the other mechanisms mentioned above are all scalable.

Theorem: A scalable truthful c -approximation mechanism, for $c < 2$, in a multi-unit auction of two items among two bidders must be a Triage mechanism for some choice of parameters.

This is the first characterization of truthfulness in an auction domain or, more generally, in a domain with no externalities. Our novel approach is radically different than previous characterization results (e.g., [30]): we analyze the *payment functions* of the players, rather than the allocation rule directly. The proof is quite involved and reveals the properties of the payment functions (monotonicity, continuity, invertibility, linearity, etc.) gradually, one property after the other. The proof also makes repeated use of the *approximation guarantee* of the mechanism, in contrast to previous results that characterized *all* mechanisms in a certain domain, and were not conditioned on the approximation ratio.

Triage mechanisms are affine maximizers on the “middle sub-domain” and we show that this extends to auctions of an arbitrary number of items among two players.

Theorem: A scalable truthful c -approximation mechanism, for $c < 2$, in a multi-unit auction of $m > 2$ items among two bidders, must be identical to an affine maximizer with VCG payments on the sub-domain where $v_i(1) = 0$ and $v_i(m-1) = v_i(m-2)$ for every player i .

Adopting the point of view of economics, our theorem can be interpreted as follows: Green and Laffont [19] characterize efficient (read: welfare-maximizing) mechanisms and show that VCG is the unique efficient mechanism. We relax the efficiency requirement to “approximate efficiency” and (almost completely) characterize all truthful (scalable) mechanisms in the multi-unit auction domain.

Interestingly, the theorem is not proved by direct characterization, but rather by reducing the characterization problem to the two-item case. We achieve this by introducing a new technical tool that enables us to use our two-item characterization as a black box: *induced mechanisms*. The technique might be of independent interest: it hints that in general characterizing truthful mechanisms may require only the characterization of small instances. The theorem immediately implies computational hardness, a first-of-a-kind result for an auction domain:

Theorem: Fix a model of computation in which finding the exact social-welfare maximizing allocation of m items between two players is computationally hard, even with valuations restricted to $v_i(1) = 0$ and $v_i(m-1) = v_i(m-2)$. Then, getting a scalable truthful c -approximation, for $c < 2$, of the social welfare in a multi-unit auction of m items among any $n \geq 2$ bidders, is also computationally hard.

This implies an exponential lower bound on communication in the black-box model [14] and implies NP-hardness in the succinct representation model, with, e.g., the bidding lan-

guage allowing valuations to be specified by boolean circuits [22]⁸.

Very recently a different approach was introduced for proving the impossibility of polynomial-time truthful mechanisms for combinatorial auctions with submodular bidders that use only value queries [10]. However, we do not know how to apply the technique of [10] to multi-unit auctions. Also note that, unlike [10], the results in this paper are not restricted to a specific type of query. Furthermore, we believe that obtaining characterizations of truthful mechanisms, whenever possible, is of interest regardless of computational considerations.

The main open problem is to get rid of the scalability assumption which we believe is not really necessary for all our theorems. We note that our reduction to the two-item case from an arbitrary number of items does not require scalability, so the hurdle is really just in characterizing the two-item two-bidder case. The fixed small size would perhaps suggest a direct attack, perhaps even a computer-assisted one, but obviously we were not able to do so.

Organization

In Section 2 the setting and basic definitions are given. Triage auctions (and two additional families of auctions) are discussed in Section 3. Section 4 characterizes two-item two-bidder truthful and scalable mechanisms. Finally, Section 5 provides a characterization of mechanisms for any number of items.

2. PRELIMINARIES

The Setting

In a multi-unit auction there is a set of m identical items, and a set $N = \{1, 2, \dots, n\}$ of bidders. Each bidder i has a valuation function $v_i : [m] \rightarrow \mathbb{R}^+$, which is normalized ($v_i(0) = 0$) and non-decreasing. Denote by V the set of all normalized and non-decreasing valuations. An allocation of the items $\vec{s} = (s_1, \dots, s_n)$ is a vector of non-negative integers such that $\sum_i s_i \leq m$ (we say that an allocation (s_1, \dots, s_n) is *infeasible* if $\sum_i s_i > m$). Denote the set of allocations by S . The goal is to find an allocation that maximizes the welfare: $\sum_i v_i(s_i)$.

In most of this paper we concentrate in the case where $n = 2$. For convenience, we name the bidders Alice and Bob. We usually denote Alice’s valuation by v , and Bob’s by u . When $m = 2$ we sometimes use the notation $(v(1), v(2))$ to denote Alice’s valuation and $(u(1), u(2))$ to denote Bob’s.

Truthfulness

The reader is referred to [26] for the (standard) proofs missing in this subsection. An n -bidder mechanism for multi-unit auctions is a pair (A, p) where $A : V^n \rightarrow S$ and $p = (p^{(1)}, \dots, p^{(n)})$, where for each i , $p^{(i)} : V^n \rightarrow \mathbb{R}$.

DEFINITION 2.1. Let (A, p) be a mechanism. (A, p) is truthful if for all i , all v_i, v'_i and all v_{-i} we have that:

$$v_i(A(v_i, v_{-i})) - p^{(i)}(v_i, v_{-i}) \geq v_i(A(v'_i, v_{-i})) - p^{(i)}(v'_i, v_{-i})$$

⁸As expected, the theorem does not imply hardness for, say, single minded bidders, since finding the welfare-maximizing allocation among two single-minded bidders is computationally easy.

⁷In terms of pure computation, scalability comes for free as one can always scale all inputs by the largest value. We also note that the truthful randomized FPTAS of [12] is scalable.

It is well known that an algorithm (for multi-unit auctions) is truthful if and only if each bidder is presented with a payment for each bundle t that does not depend on bidder i 's valuation (i.e., $p^{(i)} : V^{n-1} \rightarrow \mathbb{R}$). Denote this payment by $p_t^{(i)}(v_{-i})$. Each bidder is allocated a bundle that maximizes his profit: $v_i(t) - p_t^{(i)}(v_{-i})$ (this is called the ‘‘taxation principle’’ – we will sometimes say that these payments are *induced* by v_{-i}). We note that we may assume without loss of generality that for $t > t'$, $p_t^{(1)}(v) \geq p_{t'}^{(1)}(v)$ (‘‘payment monotonicity’’): otherwise, we have a mechanism with the same allocation rule by using $p_t^{(1)}(v) = p_{t'}^{(1)}(v)$ and the appropriate tie-breaking between bundles t and t' when $u(t) = u(t')$.

We assume that the mechanisms are *individually rational*: the payment of a bidder with an identically zero valuation is 0. Every mechanism can be made individually rational as follows: let x be the number of items a bidder receives when his valuation is identically zero and subtract $p_x^{(i)}(v_{-i})$ from all prices induced by v_{-i} . Notice that the relative preference of bundles by i has not changed (since we subtract the same constant from all prices), thus the allocation rule of the mechanism may remain the same. This normalization implies, together with the price uniqueness theorem (e.g., [26]) that in our setting there is exactly one set of prices that implement each truthful mechanism.

The following definition and proposition are standard:

DEFINITION 2.2. *A is an affine maximizer if there exist a set of allocations \mathcal{R} , a constant $\alpha_i \geq 0$ for each $i \in N$, and a constant $\beta_{\vec{s}} \in \mathfrak{R}$ for each $\vec{s} \in S$, such that $A(v_1, \dots, v_n) \in \arg \max_{\vec{s}=(s_1, \dots, s_n) \in \mathcal{R}} (\sum_i \alpha_i v_i(s_i)) + \beta_{\vec{s}}$. A is called welfare maximizer⁹ if $\beta_{\vec{s}} = 0$ for each $\vec{s} \in S$.*

PROPOSITION 2.3. *Let A be an affine maximizer (in particular, welfare maximizer). There are payments p such that (A, p) is a truthful mechanism.*

Notice that when A is a two-bidder welfare maximizer, the payments are as follows: there is a constant $w > 0$ such that for each t , a valuation v of Alice and a valuation u of Bob, $p_t^{(2)}(v) = w(v(m) - v(m - t))$ and $p_t^{(1)}(u) = (u(m) - u(m - t))/w$. We sometimes use a table notation to denote a 2-item instance. This notation is illustrated below for the 2-bidder welfare maximizer case (notice that each bidder's pavements depend only on the valuation of the other bidder):

Number of items	Alice's value	Alice's payment
One	$v(1)$	$(u(2) - u(1))/w$
Two	$v(2)$	$u(2)/w$

Number of items	Bob's value	Bob's payment
One	$u(1)$	$w(v(2) - v(1))$
Two	$u(2)$	$w \cdot v(2)$

Scalability

This paper considers two definitions of scalability.

⁹Some papers reserve the term ‘‘welfare maximizer’’ for affine maximizers in which all α_i 's are equal, and use the term ‘‘weighted welfare maximizer’’ if they are not. For simplicity, we use a slightly different terminology in this paper.

DEFINITION 2.4. *An auction is allocation scalable if multiplying the valuations of all bidders by the same positive factor does not change the allocation.*

DEFINITION 2.5. *An auction is payment scalable if for each bidder i , valuations of the other bidders v_{-i} , and $\alpha > 0$, $\alpha \cdot p^{(i)}(v_{-i}) = p^{(i)}(\alpha \cdot v_{-i})$.*

We now show that every allocation scalable mechanism is also payment scalable, and thus in this paper we use the term *scalable* to denote the less restrictive notion of scalability – payment scalability.

PROPOSITION 2.6. *Let A be an allocation scalable mechanism. Then, A is also payment scalable.*

PROOF. We prove the proposition for the case of $n = 2$ but the proof easily extends to $n > 2$ bidders. Fix a valuation u of Bob. Let $B_t(u)$ be the set of all valuations v that assign Alice t items in input (v, u) . Formally, $B_t(u) = \{v | A_1(v, u) = t\}$. We say that t is *in the range* of u if $B_t \neq \emptyset$.

We claim that t is in the range of u if and only if t is in the range of $\alpha \cdot u$. To see that, consider t that is in the range of u . We have that t is also in the range of $\alpha \cdot u$ since $A(v, u) = A(\alpha \cdot v, \alpha \cdot u)$. The ‘only if’ direction is symmetric.

We now show that for every t, t' in the range of u and $\alpha > 0$, $\alpha(p_t^{(i)}(u) - p_{t'}^{(i)}(u)) = p_t^{(i)}(\alpha \cdot u) - p_{t'}^{(i)}(\alpha \cdot u)$ (for bundles not in the range we set the payment to be equal to the payment of the next bigger bundle that is in the range, and use the appropriate tie-breaking rule). Fix some t in the range of u such that there exists $v \in B_t(u)$ that is on the border of $B_t(u)$ (in the usual topological sense). If there is no such point, Alice is always assigned t and we are done by letting $p_t^{(i)}(\alpha \cdot u) = 0$ for every $\alpha \neq 0$. Assume otherwise. There exists at least one $t' \neq t$ which is in every ϵ -neighborhood of v and in $B_{t'}$, since v is on the border. Thus we have $p_t^{(i)}(u) - p_{t'}^{(i)}(u) = v(t) - v(t')$. From scalability, we have that $\alpha \cdot v$ is on the border of $B_t(u)$ with t' playing the same role. We have that $p_t^{(i)}(\alpha \cdot u) - p_{\alpha t'}^{(i)}(u) = \alpha(v(t) - v(t')) = \alpha(p_t^{(i)}(u) - p_{t'}^{(i)}(u))$.

We continue similarly. Fix $t'' \neq t, t'$ where there exists $v \in B_{t''}(u)$, and v is on the border of $B_t \cup B_{t''}$. Thus, in every ϵ -neighborhood of v there exists a valuation v' for which $v' \in B_t$ or $v' \in B_{t''}$. Without loss of generality assume that $v' \in B_t$ (otherwise, switch the roles of t and t'). By using scalability similarly to the arguments above, we get that $v(t) - v(t'') = p_t^{(i)}(u) - p_{t''}^{(i)}(u)$, and consequently we also have $\alpha(v(t) - v(t'')) = p_t^{(i)}(\alpha \cdot u) - p_{t''}^{(i)}(\alpha \cdot u) = \alpha(p_t^{(i)}(u) - p_{t''}^{(i)}(u))$. The proof continues similarly until all bundles in the range are considered. \square

3. THE TRIAGE AUCTION

We present three families of truthful mechanisms for multi-unit auctions that provide a bounded approximation ratio for multi-unit auctions with two bidders. Each of the families contain mechanisms that are not affine maximizers. The first family, the Triage auction, includes mechanism that guarantee an approximation ratio of $1 + \epsilon$, and the next sections show that triage auctions are the only two-item truthful and scalable mechanisms that provide an approximation ratio better than 2. The other two families – shifted welfare maximizers and fractions auctions – provide an approximation of almost 2. We postpone their description to the full

version. To the best of our knowledge all previously known finitely-approximating mechanisms are either affine maximizers or are essentially single-parameter mechanisms (i.e., each bidder either receives all items, or no items at all).

We describe the mechanisms by specifying the payment functions of the bidders (recall that each function depends only on the other bidder's valuation). Truthfulness is obvious since each bidder is allocated a bundle that maximizes his profit, and we are left only with proving feasibility and analyzing the approximation ratio.

DEFINITION 3.1. *The Triage auction is parameterized by three parameters, w, θ_A, θ_B , for $w > 0$, $0 \leq \theta_A, \theta_B \leq 1$, and $\theta_A \geq 1 - \theta_B$. The payment functions are:*

- $p_m^{(2)}(v) = wv(m)$ if $v(1) < \theta_A v(m)$, and $p_m^{(2)}(v) = \frac{wv(1)}{\theta_A}$ otherwise.
- For $2 \leq k \leq m - 1$, $p_k^{(2)}(v) = p_m^{(2)}(v) - w \cdot v(m - k)$.
- $p_1^{(2)}(v) = p_m^{(2)}(v) - wv(m - 1)$ if $v(m - 1) > (1 - \theta_B)v(m)$, and $p_1^{(2)}(v) = p_m^{(2)}(v) - w(1 - \theta_B)v(m)$ otherwise (notice that in the latter case we have in fact $p_m^{(2)}(v) = wv(m)$).

and

- $p_m^{(1)}(u) = w^{-1}u(m)$ if $u(1) < \theta_B u(m)$, and $p_m^{(1)}(u) = \frac{w^{-1}u(1)}{\theta_B}$ otherwise.
- For $2 \leq k \leq m - 1$, $p_k^{(1)}(u) = p_m^{(1)}(u) - w^{-1} \cdot u(m - k)$.
- $p_1^{(1)}(u) = p_m^{(1)}(u) - w^{-1}u(m - 1)$ if $u(m - 1) > (1 - \theta_A)u(m)$, and $p_1^{(1)}(u) = p_m^{(1)}(u) - w^{-1}(1 - \theta_A)u(m)$ otherwise (again, in the latter case $p_m^{(1)}(u) = w^{-1}u(m)$).

THEOREM 3.2. *The (w, θ_A, θ_B) -Triage auction is feasible. The $(1, \theta_A, \theta_B)$ -Triage auction provides an approximation ratio of $\max(\frac{1}{\theta_A}, \frac{1}{\theta_B})$.*

We remark that when $w = \theta_A = \theta_B = 1$ we get the VCG mechanism. We also note that Section 5 gives a proof that that Triage auctions require exponential time to run.

Before proving the theorem we introduce an important definition:

DEFINITION 3.3. *Fixing the other bidder's valuation, we say that a bundle of s items is in the winning set of a bidder, if this bundle maximizes his profit.*

For the algorithms we present in this paper, we assume that the algorithm chooses an allocation (s, t) with the maximal value such that s is in the winning set of Alice and t is in the winning set of Bob.

PROOF. We will use the following claim several times:

CLAIM 3.4. *For triage auctions with $w = 1$, for each optimal allocation $(k, m - k)$, Alice's winning set contains at least one of the following bundles: k items, one item, or no items. Similarly, Bob's winning set contains at least one of the following bundles: $m - k$ items, one item, or no items.*

PROOF. We will show that the equation $v(k) - p_k^{(1)}(u) \geq v(t) - p_t^{(1)}(u)$ holds for $t \neq 0, 1$. The equation implies that Alice prefers k items over t items, thus if t is in the winning set so does k , as needed. To see that the equation holds, observe that for each $t \neq 1, 0$ we have that $p_k^{(1)}(u) - p_t^{(1)}(u) \geq u(m - k) - u(m - t)$. Since $(k, m - k)$ is an optimal allocation, we also have that $v(k) + u(m - k) \geq v(t) + u(m - t)$. Together we have that $v(k) - p_k^{(1)}(u) \geq v(t) - p_t^{(1)}(u)$, for $t \neq 0, 1$. \square

LEMMA 3.5. *The (w, θ_A, θ_B) -Triage auction is feasible.*

PROOF. We first note that without loss of generality we may assume that $w = 1$ (since multiplying one bidder's payments by w and dividing the other's by the same w maintains feasibility).

Consider first an optimal allocation $(k, m - k)$ where $k \neq 0, m$. We claim that in this case the mechanism is feasible: by Claim 3.4, Alice's winning set contains at least one of the following bundles: 0, 1, or k items. Similarly, Bob's winning set contains at least one of the following bundles: 0, 1, $m - k$. Thus there is a feasible allocation (s, t) such that s is in the winning set of Alice and t is in the winning set of Bob.

Thus from now on it suffices to assume that in all optimal allocations at least one bidder is assigned the empty bundle. We therefore assume that $(m, 0)$ is an optimal allocation (the case where $(0, m)$ is an optimal allocation is symmetric). By Claim 3.4 Bob's winning set contains at least one of the following: the empty bundle or the bundle of one item (the optimality of $(m, 0)$ implies that $v(m) \geq u(m)$, therefore if the bundle of m items is in the winning set, so is the empty bundle). Thus, we only have to show that if Bob's winning set contains only the bundle of 1 item, then Alice's winning set contains bundles that have less than m items.

If Bob's winning set contains only the bundle of one item, this implies that $u(1) > p_1^{(2)}(v)$. The definition of triage auction implies that $u(1) > v(m) - v(m - 1)$ or $u(1) > v(m)\theta_B$, depending on the ratio between $v(m - 1)$ and $v(m)$. The first case cannot happen since it implies that $u(1) + v(m - 1) > v(m)$, which is false since we assumed that $(m, 0)$ is an optimal allocation. In the second case we have that $p_m^{(1)} = \frac{u(1)}{\theta_B}$, since $u(1) > v(m)\theta_B \geq u(m)\theta_B$. Therefore, the bundle of m items is not in Alice's winning set since $v(m) < \frac{u(1)}{\theta_B}$. \square

LEMMA 3.6. *The $(1, \theta_A, \theta_B)$ -Triage auction provides an approximation ratio of $\max(\frac{1}{\theta_A}, \frac{1}{\theta_B})$.*

PROOF. Let $(k, m - k)$ be an optimal allocation. In case Alice has the bundle of k items in her winning set and Bob has the bundle of $m - k$ items in his winning set then the approximation ratio is 1. By Claim 3.4, the only other cases to consider are when at least one of the bidders (without loss of generality Alice) does not have these bundles in his winning set.

The first case we consider is when Alice's winning set does not contain the bundle of k items, but contains the empty bundle (in particular, we have that $k \neq 0$). Since the empty bundle has a zero profit, it means that the profit from the bundle of k items is negative: $v(k) < p_k^1(u)$. By the definition of the payment function we either have that $p_k^1(u) \geq u(m) - u(m - k)$ or that $p_k^1(u) = \frac{u(1)}{\theta_B} - u(m - k)$.

The first option does not occur since otherwise $v(k) < u(m) - u(m - k)$. In other words, the social welfare of the

allocation $(0, m)$ is bigger than the social welfare of the allocation $(k, m - k)$, a contradiction to the optimality of the latter. Thus, the second option occurs and in particular:

$$v(m) < \frac{u(1)}{\theta_B} \quad (1)$$

Since Alice's profit for the bundle of k items is negative, (i.e., $v(k) < \frac{u(1)}{\theta_B} - u(m - k)$), we have, in particular, that $v(m) < \frac{u(1)}{\theta_B}$. Therefore, to prove an approximation ratio of $\frac{1}{\theta_B}$, it suffices to show that that Bob has at least one non-empty bundle in his winning set. Suppose not, i.e.: $u(1) < p_1^{(2)}(v) \leq \theta_B v(m)$ (the last inequality is from the definition of the payment function – the payment for one item is always at most $\theta_B v(m)$). But then, $\frac{u(1)}{\theta_B} \leq v(m)$, a contradiction to (1).

We are left with considering the case where the only bundle in Alice's winning set is the bundle of one item. We start by showing that $u(m - 1) \leq (1 - \theta_A)u(m)$. Suppose for contradiction that $u(m - 1) > (1 - \theta_A)u(m)$. Since the bundle of one item maximizes the profit: $v(1) - p_1^{(1)}(u) > v(k) - p_k^{(1)}(u)$. Using the definition of the payment function: $v(1) + u(m - 1) > v(k) + u(m - k)$. In other words, the allocation $(k, m - k)$ is not optimal, a contradiction. Thus we have established that

$$u(m - 1) \leq (1 - \theta_A)u(m) \quad (2)$$

Recall that the bundle of one item is profitable for Alice. Using (2) and the definition of the payment function, we claim that:

$$v(1) \geq \theta_A u(m) \quad (3)$$

Since the bundle of one item maximizes the profit, using the definition of the payment function: $v(1) + u(m - 1) > v(k) + u(m - k)$. Using (2) we have that:

$$v(1) + (1 - \theta_A)u(m) > v(k) + u(m - k) \quad (4)$$

Thus, the approximation ratio is no worse than $\frac{v(k) + u(m - k)}{v(1)} \leq \frac{v(1) + (1 - \theta_A)u(m)}{v(1)} \leq 1 + \frac{(1 - \theta_A)u(m)}{\theta_A u(m)} \leq \frac{1}{\theta_A}$ (where the leftmost inequality is due to (4), and the middle one is due to (3)). \square

4. CHARACTERIZATION OF SCALABLE TWO-ITEM AUCTIONS

This section is devoted to proving the following characterization result:

THEOREM 4.1 (TWO-ITEM CHARACTERIZATION). *The only feasible, scalable and truthful auctions with an approximation ratio strictly better than 2 for two identical goods and two bidders are triage auctions for some (w, θ_A, θ_B) .*

We now provide a brief road map to the proof of the theorem. Very differently from Roberts' theorem proof, we analyze the *payment functions* of the bidders (rather than the allocation rule) and show that the payment functions are identical to the payment functions of some triage auction. Notice that in the two-item case, the payment functions of a triage auction are defined using three different regions that correspond to the ratio between the value for two items and the value for one item: high, mid, and low. The proof of the theorem is quite involved and for readability we divide it into subsections that roughly correspond to these regions.

Subsection 4.1 gives an alternative definition of the triage auction, for the special case where $m = 2$, that is easier for us to work with. Subsection 4.2 characterizes the payment function for two items. The results of subsection 4.2 hold for any scalable mechanism with a bounded approximation ratio, not just ones with an approximation ratio better than 2. The next subsections are devoted to characterizing the payment functions for one item. Subsection 4.3 defines and “separates” the high-range from the mid and low ranges: it shows that, roughly speaking, a valuation that is not in the high range induces payment (for one item) that is also not in the high range. Due to lack of space we defer the next sections of the proof to the full version, and keep here only the first simpler parts: the next subsection (in the full version) proves some basic properties, like continuity, of the payment function in the low and mid range. The central part of the proof shows that the payment functions in the mid range are equivalent to the payment functions of weighted VCG. We conclude the proof with analyzing the value of the transition points between the high and mid range, and with characterizing the high range.

4.1 An Alternative Description of the Triage Auction with Two Items

DEFINITION 4.2. *Let $p, q : [0, 1] \rightarrow \mathbb{R}^+$ and $f, g : [0, 1] \rightarrow [0, 1]$ be real valued functions and r, s be two variables that take values in $[0, 1]$. The scalable mechanism based on p and q is given by the following table.*

Number of items	Alice's value	Alice's payment
One	rv	$g(s) \cdot q(s) \cdot u$
Two	v	$q(s) \cdot u$

Number of items	Bob's value	Bob's payment
One	su	$f(r) \cdot p(r) \cdot v$
Two	u	$p(r) \cdot v$

PROPOSITION 4.3. *For any $p, q : [0, 1] \rightarrow \mathbb{R}^+$ and $f, g : [0, 1] \rightarrow [0, 1]$, the scalable mechanism based on them is scalable and truthful (but may be infeasible and allocate more than 2 items). Any truthful scalable mechanism (even a non-feasible one as long as it allocates at most two items to any bidder) is equivalent to one based on some functions.*

The proof of the proposition can be found in the full version. We now give an alternative (equivalent) definition of the triage auction, for the $m = 2$ case.

DEFINITION 4.4. *The (w, θ_A, θ_B) -trriage auction for $w > 0$ and $0 \leq \theta_A, \theta_B \leq 1$, $\theta_A \geq 1 - \theta_B$, is the scalable mechanism based on:*

- For $r \leq 1 - \theta_B$: $f(r) = \theta_B$, and $p(r) = w$.
- For $1 - \theta_B \leq r \leq \theta_A$: $f(r) = 1 - r$, and $p(r) = w$.
- For $r \geq \theta_A$: $f(r) = 1 - \theta_A$, and $p(r) = wr / \theta_A$.

and

- For $s \leq 1 - \theta_A$: $g(s) = \theta_A$, and $q(s) = w^{-1}$.
- For $1 - \theta_A \leq s \leq \theta_B$: $g(s) = 1 - s$, and $q(s) = w^{-1}$.
- For $s \geq \theta_B$: $g(s) = 1 - \theta_B$, and $q(s) = w^{-1}s/\theta_B$.

4.2 Characterizing the Payment for Two Items

The results in this section hold for any scalable and truthful mechanism with a bounded approximation ratio. We usually prove the theorem only for the function p . The proof for q is symmetric.

LEMMA 4.5. *The function p is monotone non-decreasing.*

PROOF. Assume towards contradiction that for some $r' > r$ we have $p(r') < u < u' < p(r)$. Since $u > p(r')$, on inputs $(r', 1)$ and $(0, u)$ Bob must win both items, so Alice cannot win anything. Notice that $(r, 1)$ wins nothing against $(0, u'(1 + \epsilon))$ (by payment scalability, $(0, u'(1 + \epsilon))$ induces bigger payments than $(0, u')$, and Alice did not win any items with the bigger valuation $(r, 1)$), but also Bob does not win both items since $u'(1 + \epsilon) < p(r)$ for small enough $\epsilon > 0$, so the total welfare achieved is 0 contradicting finite approximation. \square

LEMMA 4.6 (WEIGHTING). $p(0) \cdot q(0) = 1$.

PROOF. Consider the following input:

Number of items	Alice's value	Alice's payment
One	0	?
Two	v	$uq(0)$

Number of items	Bob's value	Bob's payment
One	0	?
Two	u	$vp(0)$

The only allocations that give finite approximation ratio on inputs of the form $(0, v)$ and $(0, u)$ are those that give both items to one of the bidders. If $u < vp(0)$ then Bob does not win two items; whereas if $u > vp(0)$ then he wins both items, and dually for Alice. So we get a contradiction to feasibility if $u > vp(0)$ and $v > uq(0)$, i.e., if $p(0)q(0) < 1$. On the other hand, if $u < vp(0)$ and $v < uq(0)$, i.e., if $p(0)q(0) > 1$, then we get a total welfare of 0, contradicting finite approximation ratio. \square

At this point we are ready to give a more precise definition of the payment function. We start with the low range, i.e., when $r < g(0)$. In particular we show that the function is constant in this range.

LEMMA 4.7 (LOW RANGE). *If $r < g(0)$ then $p(r) = p(0)$.*

PROOF. Assume that $p(r) \neq p(0)$ then, using monotonicity, let $p(r) > u' > u > p(0)$. On input $(0, 1)$ and $(0, u)$ Bob gets both items (since $u > p(0)$) and so Alice must get none. On inputs $(r, 1)$ and $(0, u')$ Alice gets at most 1 item (since, by the scalability of the payments, the payment induced by Bob for two items has increased), but since $u' < p(r)$ Bob does not get two items, and so for finite approximation, Alice must get an item, so $r \geq u'g(0)q(0) \geq p(0)q(0)g(0) = g(0)$. \square

The following claim will be helpful in analyzing the payment function in the high range:

CLAIM 4.8. $p(r) \geq r/(g(0)q(0))$.

PROOF. Let $u > p(r)$, then on input $(r, 1)$ and $(0, u)$ Bob gets both items. Alice's payment for a single item is $uq(0)q(0)$ which for feasibility must be at least r . Since this holds for all $u > p(r)$ we get that $r \leq p(r)g(0)q(0)$ as required. \square

For the high range ($r > g(0)$) we show that the payment grows in a specific linear way:

LEMMA 4.9 (HIGH RANGE). *If $r > g(0)$ then $p(r) = r/(g(0)q(0))$.*

PROOF. We will prove the contra-positive which by the previous claim assumes $p(r) > u > r/(g(0)q(0))$. Consider the following input:

Number of items	Alice's value	Alice's payment
One	r	$uq(0)g(0)$
Two	1	$uq(0)$

Number of items	Bob's value	Bob's payment
One	0	?
Two	u	$p(r)$

In this case Bob cannot win both items so he gets a value of 0. By the choice of u , Alice's payment for two items is $uq(0) > r/g(0)$ is greater than 1, thus she cannot win two items. Thus for finite approximation she must win one item and thus $r \geq uq(0)g(0)$ and since this is true for every $u < p(r)$, we have $r \geq p(r)g(0)q(0)$, contradiction. \square

At this point we have completed the required characterization of p and q .

DEFINITION 4.10. *Let $w = p(0)$, $\theta_A = g(0)$, and $\theta_B = f(0)$.*

LEMMA 4.11 (SUMMARY OF SUBSECTION). *For $r \leq \theta_A$ we have that $p(r) = w$ and for $r \geq \theta_A$ we have $p(r) = wr/\theta_A$. Similarly, for $s \leq \theta_B$ we have that $q(s) = w^{-1}$ and for $s \geq \theta_B$ we have $q(s) = w^{-1}s/\theta_B$.*

PROOF. The low range lemma states the required fact for $r < \theta_A$. The high range lemma states the required fact for $r > \theta_A$, taking into account the inverse lemma, $p(0)q(0) = 1$, the same holds for q , replacing w with w^{-1} , again relying on $p(0)q(0) = 1$. For $r = \theta_A$ we observe that $p(r) = w$ since p is a monotone function and approaches w above and below w . \square

4.3 Separating the high range

In this section we show that for $r \leq \theta_A$ we have that $f(r) \leq \theta_B$. Similarly it follows that for $s \leq \theta_B$ we have that $g(s) \leq \theta_A$. Note that by the previous section $r > \theta_A$ if and only if $p(r) > w$, and this last condition is what drives this section.

At this point we separate into two cases, according to whether $r > wf(r)$. We start with the easy case: we show that if the payment for one item is "too high" then we do not get the required approximation ratio.

LEMMA 4.12 (CASE I). *If $r \leq \theta_A$ and $r \leq wf(r)$ then $f(r) \leq \theta_B$.*

PROOF. Assume by way of contradiction that $f(r) > \theta_B$ and so $q(f(r)) > w^{-1}$, and for ϵ small enough consider the input:

Number of items	Alice's value	Alice's payment
One	r	?
Two	1	$w(1 - \epsilon)q(f(r)) > 1$

Number of items	Bob's value	Bob's payment
One	$f(r)(1 - \epsilon)w$	$f(r)w$
Two	$(1 - \epsilon)w$	w

Notice that Bob gets negative utility from taking an item or two items and thus takes nothing. Alice gets negative utility from taking two items so can take at most a single item. The total welfare is thus at most r , whereas the social optimum is at least $r + f(r)(1 - \epsilon)w$. Since $r \leq wf(r)$ this is a contradiction to better than 2-approximation. \square

For the second case we first need to prepare two lemmas and a corollary.

LEMMA 4.13 (WEAK ONE-SIDE INVERSE). *If $r \leq \theta_A$ then for any $\delta > 0$, $wg(f(r) - \delta)q(f(r) - \delta) \geq r$.*

PROOF. Assume to the contrary $wg(f(r) - \delta)q(f(r) - \delta) < r$ and consider the following input:

Number of items	Alice's value	Alice's payment
One	r	$wg(f(r) - \delta)q(f(r) - \delta)(1 + \epsilon) < r$
Two	1	?

Number of items	Bob's value	Bob's payment
One	$(f(r) - \delta)w(1 + \epsilon)$	$wf(r)$
Two	$w(1 + \epsilon)$	w

Bob takes two items. However, when ϵ is small enough, Alice gets positive utility from one item so she will take (at least) a single item, contradicting feasibility. \square

LEMMA 4.14. *For $r > \theta_A$ we have that $r > f(r)p(r)$.*

PROOF. Consider the input:

Number of items	Alice's value	Alice's payment
One	r	?
Two	1	$q(f(r))p(r)(1 - \epsilon)$

Number of items	Bob's value	Bob's payment
One	$f(r)p(r)(1 - \epsilon)$	$f(r)p(r)$
Two	$p(r)(1 - \epsilon)$	$p(r)$

Bob has negative utility for either one item or two items. Since $p(r) > w$ and $q(f(r)) \geq w^{-1}$, Alice has negative utility

for two items, as long as ϵ is small enough. Thus the total welfare is at most r , whereas the social optimum is at least $r + f(r)p(r)(1 - \epsilon)$, so for better than 2-approximation we must have $r > f(r)p(r)$. \square

COROLLARY 4.15. *If $f(r) > \theta_B$ then $f(r) > g(f(r))q(f(r))$.*

PROOF. This is the previous lemma with the roles of the players switched and with $s = f(r)$. \square

We are now ready to handle the second case:

LEMMA 4.16 (CASE II). *If $r \leq \theta_A$ and $r > wf(r)$ then $f(r) \leq \theta_B$.*

PROOF. Assume towards contradiction that there exists $\delta > 0$ such that $f(r) - \delta > \theta_B$. Combining the weak one-sided inverse lemma and the previous corollary we have that $f(r) - \delta > g(f(r) - \delta)q(f(r) - \delta) \geq r/w$; Contradiction. \square

Which concludes this subsection:

LEMMA 4.17. *For $r \leq \theta_A$ we have $f(r) \leq \theta_B$. For $s \leq \theta_B$ we have $q(s) \leq \theta_B$.*

PROOF. The Case I and Case II lemmas cover all possibilities for f ; for g the situation is symmetric. \square

5. CHARACTERIZING MECHANISMS FOR ANY NUMBER OF ITEMS

We showed that every two-item scalable mechanism that provides an approximation ratio better than 2 is a triage auction. This section gives an almost complete characterization for truthful and scalable mechanisms that guarantee an approximation ratio better than 2 for any number of items. In particular this section's characterization implies that truthful and scalable mechanisms for multi-unit auctions cannot guarantee an approximation ratio better than 2 in polynomial-time.

The two-item characterization is used as a black box to characterize mechanisms for more items. Importantly, the scalability assumption is not used in this section. In other words, proving that triage auctions are the only truthful mechanisms (scalable or not) that provide an approximation ratio better than 2 in multi-unit auctions with only two items, would immediately imply our characterization result for any number of items, and in particular would imply an unconditional lower bound on the power of all polynomial time truthful mechanisms. All missing proofs appear in the full version.

5.1 Induced Mechanisms: Definition and Basic Properties

Our main working horses will be induced mechanisms. Induced mechanisms allow us to define a two-item mechanism given an m -item mechanism. By leveraging our two-item characterization, we show that the induced two-item mechanisms are triage auctions. We then study the relationship between all induced mechanisms and prove that many of them must be welfare maximizers. We show that this implies that the m -item mechanism we started with must have a very specific form, as needed.

DEFINITION 5.1. *Let l_1, h_1 be such that $1 \leq l_1 < h_1 \leq m$. The (l_1, h_1) -extension of a two-item valuation v , denoted v^{l_1, h_1} , is defined as follows: for every $k < l_1$, $v^{l_1, h_1}(k) = 0$. For every $h_1 > k \geq l_1$, $v^{l_1, h_1}(k) = v(1)$. For every $k \geq h_1$, $v^{l_1, h_1}(k) = v(2)$.*

DEFINITION 5.2 (INDUCED MECHANISM). Let A be a mechanism for multi-unit auctions with m items and 2 bidders. Let l_1, h_1, l_2, h_2 be positive integers with the following constraints: $l_1 < h_1 \leq m$, $l_2 < h_2 \leq m$, $l_1 + l_2 \leq m$, $l_1 + h_2 > m$ and $l_2 + h_1 > m$. Define the induced mechanism A^{l_1, h_1, l_2, h_2} for 2 items as follows: given two valuations v and u run A with the (l_1, h_1) -extended valuation v^{l_1, h_1} and the (l_2, h_2) -extended valuation $u^{(l_2, h_2)}$. Let (a_1, a_2) be the output allocation of A and (p_1, p_2) be the payments the bidders are charged in A . If $a_1 < l_1$ then let $a'_1 = 0$, if $l_1 \leq a_1 < h_1$ then let $a'_1 = 1$, otherwise let $a_1 = 2$. If $a_2 < l_2$ then let $a'_2 = 0$, if $l_2 \leq a_2 < h_2$ then let $a'_2 = 1$, otherwise let $a_2 = 2$. The output of A^{l_1, h_1, l_2, h_2} on v and u is (a'_1, a'_2) . Alice's payment is p_1 and Bob's payment is p_2 .

PROPOSITION 5.3. Let A be a truthful and scalable mechanism for multi-unit auctions with m items and 2 bidders that provides an approximation ratio of α . Let A^{l_1, h_1, l_2, h_2} be an induced mechanism. A^{l_1, h_1, l_2, h_2} is feasible, truthful, scalable, and provides an approximation ratio of α .

In this section we denote the $p^{(2)}$ function of A (the payments induced by Alice) by f and by f^{l_1, h_1, l_2, h_2} the $p^{(2)}$ function of the induced mechanism A^{l_1, h_1, l_2, h_2} . We denote by g the $p^{(1)}$ function of A (the payments induced by Bob) and by g^{l_1, h_1, l_2, h_2} the $p^{(1)}$ function of the induced mechanism A^{l_1, h_1, l_2, h_2} . As a corollary of Proposition 5.4 we get the following relationship between the payment functions of A and its induced mechanisms.

COROLLARY 5.4. Let v be a two-item valuation and let $v^{(l_1, h_1)}$ be its (l_1, h_1) -extension. Let l_2, h_2 be such that A^{l_1, h_1, l_2, h_2} is an induced mechanism. $f_{l_2}(v^{h_1, l_1}) = f_1^{l_1, h_1, l_2, h_2}(v)$ and $f_{h_2}(v^{h_1, l_1}) = f_2^{l_1, h_1, l_2, h_2}(v)$.

Symmetrically, let u be a two-item valuation and let $u^{(l_2, h_2)}$ be its (l_2, h_2) -extension. Let l_1, h_1 be such that A^{l_1, h_1, l_2, h_2} is an induced mechanism. $g_{l_1}(u^{h_2, l_2}) = g_1^{l_1, h_1, l_2, h_2}(u)$ and $g_{h_1}(u^{h_2, l_2}) = g_2^{l_1, h_1, l_2, h_2}(u)$.

5.2 Relations between Induced Mechanisms

Let A be a scalable and truthful mechanism for multi-unit auctions for m items with an approximation ratio better than 2. By our characterization and the discussion above we have that all induced mechanisms of A are triage auctions. Denote the parameters of the triage mechanism A^{l_1, h_1, l_2, h_2} by $\theta_A^{l_1, h_1, l_2, h_2}, \theta_B^{l_1, h_1, l_2, h_2}, w^{l_1, h_1, l_2, h_2}$. The point of this subsection is to study the relations between the parameters of the induced triage mechanisms of A .

CLAIM 5.5. Let A be a truthful and scalable mechanism for multi unit auctions with m items that provides an approximation ratio better than 2. Let A^{l_1, h_1, l_2, h_2} and $A^{l_1, h_1, l_2, h_2'}$ be two induced mechanisms of A . Then, $w^{l_1, h_1, l_2, h_2} = w^{l_1, h_1, l_2, h_2'}$, $\theta_A^{l_1, h_1, l_2, h_2} = \theta_A^{l_1, h_1, l_2, h_2'}$, and $\theta_B^{l_1, h_1, l_2, h_2} = \theta_B^{l_1, h_1, l_2, h_2'}$. Symmetrically, let A^{l_1, h_1, l_2, h_2} and A^{l_1, h_1', l_2, h_2} be two induced mechanisms of A . Then, $w^{l_1, h_1, l_2, h_2} = w^{l_1, h_1', l_2, h_2}$, $\theta_A^{l_1, h_1, l_2, h_2} = \theta_A^{l_1, h_1', l_2, h_2}$, and $\theta_B^{l_1, h_1, l_2, h_2} = \theta_B^{l_1, h_1', l_2, h_2}$.

CLAIM 5.6. Let A be a truthful and scalable mechanism for multi unit auctions with m items that provides an approximation ratio better than 2. Let A^{l_1, h_1, l_2, h_2} and A^{l_1, h_1, l_2', h_2} be two induced mechanisms of A . Then, $\theta_A^{l_1, h_1, l_2, h_2} = \theta_A^{l_1, h_1, l_2', h_2}$

and $w^{l_1, h_1, l_2, h_2} = w^{l_1, h_1, l_2', h_2}$. Symmetrically, let A^{l_1, h_1, l_2, h_2} and $A^{l_1, h_1, l_2, h_2'}$ be two induced mechanisms of A . Then, $\theta_B^{l_1, h_1, l_2, h_2} = \theta_B^{l_1, h_1, l_2, h_2'}$ and $w^{l_1, h_1, l_2, h_2} = w^{l_1, h_1, l_2, h_2'}$.

We now use the claims to prove that all induced mechanisms share the same w . Thus, after proving it we may denote the w^{l_1, h_1, l_2, h_2} parameter of every induced mechanism A^{l_1, h_1, l_2, h_2} by (the same) w .

LEMMA 5.7. Let A be a truthful and scalable mechanism for multi unit auctions with m items that provides an approximation ratio better than 2. Let A^{l_1, h_1, l_2, h_2} and A^{l_1, h_1', l_2, h_2} be two induced mechanisms of A . Then, the following equality holds: $w^{l_1, h_1, l_2, h_2} = w^{l_1, h_1', l_2, h_2}$.

5.3 Some Induced Mechanisms are Welfare Maximizers

The heart of this section is here. We show that “many” of the induced triage auctions take the simple form of welfare maximizers. The crux is that the payment function for some items is simultaneously the payment function of one item for one induced mechanism, and the payment function for two items for another. Simple algebra then gives us that some θ 's must equal to 1. We are then able to specify the payment functions of “simple” valuations.

LEMMA 5.8. Let A be a truthful and scalable mechanism for multi unit auctions with m items with an approximation ratio better than 2. Let $A^{l_1, m, l_2, m}$ where $l_1, l_2 \geq 2$. Then, $\theta_A^{l_1, m, l_2, m} = \theta_B^{l_1, m, l_2, m} = 1$.

DEFINITION 5.9. A valuation v is l -simple if there exists some $0 < l < m$ such that for every $k < l$, $v(k) = 0$, and for every $l \leq k < m$ we have that $v(k) = v(l)$.

COROLLARY 5.10. Let $l \geq 2$. For every l -simple valuation v we have that $f_m(v) = wv(m)$ and for all $1 < t < m - 1$ such that $l + t \leq m$ we have that $f_t(v) = w(v(m) - v(m - t))$. Similarly, for every l -simple valuation u we have that $f_m(u) = u(m)/w$ and for all $1 < t < m - 1$ such that $l + t \leq m$ we have that $g_t(u) = (u(m) - u(m - t))/w$.

PROOF. For the first part, observe that $A^{l, m, t, m}$ is an induced mechanism for every t such that $l + t \leq m$ and apply Lemma 5.8. The proof of the second part is similar. \square

5.4 Concluding the Characterization

We are now ready to obtain our final characterization. We give an almost complete description of the payment functions for valuations where the value of one item is 0. Lemma 5.11 provides the payment function for m items, and Lemma 5.12 provides the payment functions for smaller bundles.

LEMMA 5.11. For each v where $v(1) = 0$, $f_m(v) = wv(m)$. Symmetrically, for each u where $u(1) = 0$, $g_m(u) = \frac{u(m)}{w}$.

LEMMA 5.12. Let v be a valuation with $v(1) = 0$. For every $k \neq 1, m - 1$ we have that $f_k(v) = w(v(m) - v(m - k))$. Similarly, for every valuation u with $u(1) = 0$ we have, for every $k \neq 1, m - 1$ that $g_k(u) = w^{-1}(u(m) - u(m - k))$.

Our final characterization result is:

DEFINITION 5.13. A valuation v is degenerate if $v(1) = 0$ and $v(m - 1) = v(m - 2)$.

THEOREM 5.14. (Characterization of mechanisms for any number of items) *Let A be a truthful and scalable two-bidder mechanism for $m > 2$ items that provides an approximation ratio better than 2. There exists a constant $w > 0$ such that for all degenerate v and u and on all inputs (v, u) A outputs a solution with value $\max_k(v(k) + wu(m - k))$.*

Acknowledgements

The first author wishes to thank Itai Ashlagi and Ron Lavi for several stimulating early discussions. The fractions auction of Section 3 was obtained during these discussions. We thank Bobby Kleinberg for helpful discussions. We thank Sigal Oren for numerous very helpful comments and ideas throughout the work on this paper.

6. REFERENCES

- [1] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In FOCS'01.
- [2] Itai Ashlagi, Shahar Dobzinski, and Ron Lavi. An optimal lower bound for anonymous scheduling mechanisms. In EC'09.
- [3] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In TARK'03.
- [4] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu'alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, July 2006.
- [5] Patrick Briest, Piotr Krysta, and Berthold Vöcking. Approximation techniques for utilitarian mechanism design. In STOC'05.
- [6] Dave Buchfuhrer, Shaddin Dughmi, Hu Fu, Robert Kleinberg, Elchanan Mossel, Christos Papadimitriou, Michael Schapira, Yaron Singer, and Chris Umans. Inapproximability for vcg-based combinatorial auctions. In SODA'10.
- [7] George Christodoulou, Elias Koutsoupias, and Angelina Vidali. A lower bound for scheduling mechanisms. In SODA'07.
- [8] George Christodoulou and Annamária Kovács. A deterministic truthful PTAS for scheduling related machines. In SODA'10.
- [9] Peerapong Dhangwatnotai, Shahar Dobzinski, Shaddin Dughmi, and Tim Roughgarden. Truthful approximation schemes for single-parameter agents. In FOCS'08.
- [10] Shahar Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In STOC'11.
- [11] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In APPROX'07.
- [12] Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. In FOCS'09.
- [13] Shahar Dobzinski and Noam Nisan. Limitations of vcg-based mechanisms. Preliminary version in STOC'07.
- [14] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. In EC'07.
- [15] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In STOC'05.
- [16] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In STOC'06.
- [17] Shahar Dobzinski and Mukund Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In EC'08.
- [18] Uriel Feige. On maximizing welfare where the utility functions are subadditive. In STOC'06.
- [19] J. Green and J.J. Laffont. Characterization of satisfactory mechanism for the revelation of preferences for public goods. *Econometrica*, pages 427–438, 1977.
- [20] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.
- [21] Elias Koutsoupias and Angelina Vidali. A lower bound of $1+\phi$ for truthful scheduling mechanisms. In MFCS'07.
- [22] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In FOCS'03.
- [23] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In FOCS'05.
- [24] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. In *JACM* 49(5), pages 577–602, Sept. 2002.
- [25] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In AAAI'02.
- [26] Noam Nisan. 2007. Introduction to Mechanism Design (for Computer Scientists). In “Algorithmic Game Theory”, N. Nisan, T. Roughgarden, E. Tardos and V. Vazirani, editors.
- [27] Noam Nisan and Amir Ronen. Algorithmic mechanism design. In STOC'99.
- [28] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [29] Christos Papadimitriou, Michael Schapira, and Yaron Singer. On the hardness of being truthful. In FOCS'08.
- [30] Kevin Roberts. The characterization of implementable choice rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences. Papers presented at the first European Summer Workshop of the Economic Society*, pages 321–349. North-Holland, 1979.
- [31] Michael E. Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In EC'05.
- [32] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.