Towards a Characterization of Truthful Combinatorial Auctions

(Extended Abstract)

Ron Lavi, Ahuva Mu'alem and Noam Nisan *
School of Engineering and Computer Science
The Hebrew University of Jerusalem
Israel
{tron,ahumu,noam}@cs.huji.ac.il

Abstract

This paper analyzes incentive compatible (truthful) mechanisms over restricted domains of preferences, the leading example being combinatorial auctions. Our work generalizes the characterization of Roberts (1979) who showed that truthful mechanisms over unrestricted *domains with at least 3 possible outcomes must be "affine maximizers". We show that truthful mechanisms for combinatorial auctions (and related restricted domains) must be "almost affine maximizers" if they also satisfy an additional requirement of "independence of irrelevant alternatives". This requirement is without loss of generality for unrestricted domains as well as for auctions between two players where all goods must be allocated. This implies unconditional results for these cases, including a new proof of Roberts' theorem. The computational implications of this characterization are severe as reasonable "almost affine maximizers" are shown to be as computationally hard as exact optimization. This implies the near-helplessness of such truthful polynomial-time auctions in all cases where exact optimization is computationally intractable.*

1 Introduction

1.1 Background

In recent years we have seen much research aimed at designing algorithms that are intended to function in environments where the inputs to the algorithm are distributed among players who have different goals. Such environments become increasingly common on the Internet, in communication networks, and in many electronic commerce situations. In such cases, the algorithms must be designed in a way that motivates the participants to reveal the true inputs rather than some carefully designed lies intended to produce results favored by the particular participant.

The most commonly studied model is taken from the economic field of mechanism design (see e.g. [23, 26, 18, 22, 6, 2, 10, 9, 29]). In this model each player has a private valuation function that assigns real values to each possible outcome of the algorithm. The algorithm motivates the players by handing payments to them. Each player is assumed to be rational in the sense of attempting to maximize his net utility, i.e. the sum of his valuation of the algorithms' outcome and his payment. Such algorithms with attached payment functions are called mechanisms. The payments should be carefully designed to ensure that each rational player will report his input – his valuation function – truthfully. Mechanisms for which reporting the truth is a dominant strategy for each player are called incentive compatible, or truthful[1]. The basic challenge in this field is that of designing *computationally efficient incentive compatible* mechanisms for the relevant computational problems.

A particularly central problem of this kind is that of combinatorial auctions. In a combinatorial auction, k items are simultaneously auctioned among n bidders. Bidders value bundle of items in a way that may depend on the combination they win, i.e. each bidder has a valuation function v_i that assigns a real value $v_i(S)$ for each possible subset of items that he may win. The goal is to find a partition $S_1 ... S_n$ of the items in a way that maximizes the total social welfare $\sum_i v_i(S_i)$. This problem has many interesting sub-cases, abstracts many complex allocation problems, and has been extensively studied. The combinatorial auction problem is NP-complete for even very simple valuation functions v_i,

* Supported by grants from the Israeli Ministry of Science and the Israeli Academy of Sciences.

[1] This model makes several assumptions that may all be relaxed. The most significant ones are the existence of "money" (quasi-linearity), the independence of the valuations (private values), and using the notion of dominant strategies rather than weaker notions of equilibria.

and may be approximated to within a factor of $O(\sqrt{k})$ (but no better [19, 14]). Experimental results have shown that many reasonable heuristics can quickly obtain optimal or approximately optimal solutions for combinatorial auctions with up-to thousands of items [33, 28, 12].

Unfortunately, it is not known how to turn such non-fully-optimal heuristics or approximation algorithms into incentive compatible mechanisms! There is only one known general method for designing incentive compatible mechanisms: the VCG payment scheme [35, 8, 13]. The key difficulty is the fact that VCG-payments ensure truthfulness only when the chosen outcome optimizes the social welfare; thus attaching VCG payments to an approximation algorithm or heuristic does *not* ensure incentive compatibility! This problem was noticed in [26, 19] and was shown to be essentially universal in [25]. Additionally, there are many computational scenarios where the optimization goal is different than optimizing the social welfare [26]. Thus, the main positive tool of mechanism design does not address the basic challenge of algorithmic mechanism design for combinatorial auctions as well as most other computationally interesting scenarios! Indeed, almost all mechanisms designed so far for computational scenarios apply only to very restricted "single dimensional" domains. Such mechanisms include e.g. scheduling to minimize the makespan [2], revenue maximization (for digital and other types of goods) [10, 31], auctioning with bounded communication [6], as well as combinatorial auctions with very restrictive bidders (see below).

1.2 Characterizing Incentive Compatibility

A general approach to this central task of designing computationally efficient incentive compatible mechanisms would be to obtain a characterization of their powers. To do this, let us get slightly more formal about the basic model.

There is a set A of possible outcomes of the mechanism, and each player has a valuation function $v_i : A \to \mathcal{R}$ that specifies his value $v_i(a)$ for each possible outcome $a \in A$, where v_i is chosen from some possible domain of valuations V_i. For each n-tuple of valuations $v = (v_1, \ldots, v_n)$, the mechanism produces some outcome $f(v)$ that may be viewed as aggregating the preferences v_i of the n players. The function f is called the *social choice function*. Additionally, the mechanism hands out payments to the players. For example, in the case of combinatorial auctions, A is the set of all possible partitions $(a_1, \ldots, a_n)$ of the items, and each V_i is the set of valuations that depend only on a_i ("no externalities") and are monotone in a_i ("free disposal"). It turns out that for each social choice function f that may be obtained by an incentive compatible mechanism there is essentially a single way to "implement it", i.e. to set the payments needed as to ensure incentive compatibility[2]. The basic question is *what social choice functions are implementable?*

The VCG mechanism mentioned above implements the social choice function that maximizes the social welfare, i.e. the social choice function $f(v) = \text{argmax}_{a \in A} \sum_i v_i(a)$. Three generalizations may be applied to the VCG payment scheme, yielding generalizations to the implemented social choice function: (a) the range may be restricted to an arbitrary $A' \subset A$; (b) different non-negative weights ω_i can be given to the different players; (c) different additive weights γ_a can be given to different outcomes. All three generalizations can be combined, yielding an implementation for any social choice function that is an *affine maximizer*:

Definition: *A social choice function f is an affine maximizer if for some $A' \subset A$, non-negative $\{\omega_i\}$, and $\{\gamma_a\}$, for all $v_1 \in V_1, \ldots, v_n \in V_n$ we have $f(v_1, \ldots, v_n) \in argmax_{a \in A'}(\sum_i \omega_i v_i(a) + \gamma_a)$.*

What other social choice functions can be implemented? A classic negative result of Roberts [30] shows that if the domain of players' valuations is unrestricted, and the range is non-trivial, then nothing more:

Theorem (Roberts, 1979): *If there are at least 3 possible outcomes, and players' valuations are unrestricted ($V_i = \mathcal{R}^{|A|}$), then any implementable[3] social choice function is an affine maximizer.*

The requirement that the valuations are unrestricted is very restrictive. In almost all interesting computational applications the domain of valuations is restricted. E.g., as mentioned, for the combinatorial auction problem the valuations are restricted in two ways: "free disposal" and no "externalities", and thus $V_i \neq \mathcal{R}^{|A|}$. Indeed, some assumption about the space of valuations is also necessary: In the extreme opposite case, the domain is so restricted as to become single dimensional, for which truthful non affine maximizers exist. A trivial example is given by an auction of a single item, where v_i is a scalar, and a truthful auction may assign the auctioned item to the bidder i with highest value of $(v_i)^i$ (for the price of $(v_j)^{j/i}$ where j is the player with second highest value of $(v_j)^j$).

More interesting examples in the context of combinatorial auctions involve "single-minded" bidders, where the valuation function is given by a single value v_i offered for a single set of items S_i [19]. While the optimization problem in this case is still NP-complete and thus affine maximization is not efficiently computable, [19] presented computationally efficient truthful approximation mechanisms for

[2] Imprecisely stated, when outcome a is chosen, each player i is charged the minimum value of $v_i(a)$ that would still result in a being chosen.

[3] Roberts, as we do here, only discusses implementation in dominant strategies in private-value environments. See [20] for a generalization to environments with inter-dependent valuations.

it. Additional mechanisms for this single-minded case were presented in [22, 1].

However, most interesting computational problems are not single dimensional either – they lie somewhere between the two extremes of "unrestricted" and "single dimensional". This intermediate range includes combinatorial auctions and many of their interesting special cases such as, multi-unit (homogeneous) auctions, or unit-demand auctions (matching). It also includes most examples of other combinatorial optimization problems such as various variants of scheduling and routing problems. Almost nothing is known about this intermediate range. The only positive example of a non-VCG mechanism for non-single-dimensional domains is for a special case of multi-unit combinatorial auctions where each bidder is restricted to demand at most a fraction of the number of units of each type [5].

It is interesting to draw parallels with the non-quasi-linear case, i.e. the model where player preferences are given by order relations $\succeq_i$ over the possible outcomes. The classic Gibbard-Satterthwaite result [11, 34] shows that, in this case, no non-trivial social choice function over an unrestricted domain is implementable. The proof shows that any implementable social choice function must essentially satisfy Arrow's condition of "Independence of Irrelevant Alternatives", and thus Arrow's impossibility result [4] applies. On the other hand, in this non-quasi-linear case, there exists much literature implementing non-trivial social choice functions over various interesting restricted domains, e.g. "single peaked domains" [7, 21].

1.3 Our results

In this paper we initiate an analysis of implementable social choice functions over restricted domains in quasi-linear environments. It is widely known that certain monotonicity requirements characterize implementable social choice functions. E.g. Roberts starts by defining a condition of "positive association of differences" (PAD) that characterizes implementable social choice functions over unrestricted domains. It turns out that this condition is usually meaningless for restricted domains. We start with a formulation of a "weak monotonicity" condition (W-MON), that provides this characterization for "usual" restricted domains (exact definitions of all terms appear in the body of the paper). We demonstrate that other natural notions are not appropriate.

Lemma: *Every implementable social choice function over every domain must satisfy W-MON. Over "usual" domains, W-MON is also a sufficient condition.*

As opposed to the case of unrestricted domains, it turns out that, for restricted domains, W-MON by itself does not imply affine maximization! A key contribution of this paper is the identification of a key additional property, *Independence of Irrelevant Alternatives (IIA)*, that will provide this implication. This property is a natural analog, in the quasi-linear setting, of Arrow's similarly named property in the non-quasi-linear setting. This condition states that if the social choice function changes its value from one outcome a to another outcome b, then this is due to a change in some player's preference between a and b.

Definition: *A social choice function f satisfies IIA if for any $v, u \in V$, if $f(v) = a$ and $f(u) = b \neq a$ then there exists a player i such that $u_i(a) - u_i(b) \neq v_i(a) - v_i(b)$.*

We show that the IIA property is equivalent to a slight, but significant, strengthening of the W-MON condition, termed "strong monotonicity". We further show that in unrestricted domains IIA may be assumed without loss of generality. This is also true in a class of domains that includes the case of combinatorial auctions with two players in which all items are always allocated. In other domains we demonstrate that IIA may not be assumed without loss of generality.

We then get to our main result: incentive compatible mechanisms that also satisfy IIA must be "almost" affine maximizers. The theorem is proved in a general setting and requires certain technical conditions.

Main Theorem: *In "certain reasonable domains", an implementable social choice function that additionally satisfies IIA and certain technical conditions must be an "almost" affine maximizer.*

The proof of this theorem is different from the one Roberts provides for unrestricted domains, and uses ideas suggested, in a somewhat different context, by Archer and Tardos [3]. This theorem applies to combinatorial auctions as well as to multi-unit (non-combinatorial) auctions. It even applies to the case of "known double minded bidders", i.e. where each bidder has only two bundles on which he may bid – showing that the mechanisms of [19, 22] regarding single-minded bidders cannot be generalized this way. For unrestricted domains, the IIA condition may be assumed without loss of generality, and therefore this yields a new proof of Roberts' theorem (the qualifications in the theorem statement all disappear in this case). For two-player auctions where all items must always be allocated, the IIA condition can similarly be dropped. We also show that in this two-player case, the requirement that all items must always be allocated is necessary – without it, there exist implementable social choice functions that are not almost affine maximizers (and do not satisfy IIA)[4].

[4]We note that the incentive compatible mechanism of [5] also does not always allocate all items.

The major open problem we leave is whether the IIA condition is necessary:

Main Open Problem: *Are there incentive compatible combinatorial auctions that are not "essentially" affine maximizers?*

The meaning of "essentially" in this open problem is soft, as we demonstrate that various "minor" variations from affine maximization are possible. The question is really whether anything *useful* is possible, e.g. can any non-trivial approximation be achieved.

We close the circle by returning to our original motivation, that of the existence of computationally efficient incentive compatible approximation mechanisms. We observe that essentially any affine maximizer is as computationally hard as exact social welfare maximization. This implies that if exact computation of the optimal allocation is computationally hard, then incentive compatible mechanisms that satisfy IIA are essentially powerless. For an exact statement of computational hardness we must first fix an input format, i.e. a "bidding language" [24] that is powerful enough to make the exact optimization problem computationally intractable[5]. We say that a combinatorial auction mechanism is *unanimity-respecting* if whenever every bidder values only a single bundle, and furthermore, these bundles compose a valid allocation, then this allocation is chosen[6]. This condition ensures that all allocations are possible outputs, ruling out "bundling" auctions[7].

Theorem: *(Assuming $P \neq NP$ and a sufficiently powerful bidding language) Any unanimity-respecting truthful polynomial-time combinatorial (or multi-unit) auction that satisfies IIA cannot obtain any polynomially-bounded approximation ratio.*

Note that without the incentive-compatibility requirement, an $O(\sqrt{k})$ and n approximations can be obtained for combinatorial auctions, and an FPAS for multi-unit auctions. Note also that for the two-player case where all items are always allocated, the IIA condition can be dropped.

An especially crisp result is obtained for the case of two-player multi-unit auctions. This case is computationally equivalent to a weighted knapsack problem and is thus NP-complete to solve exactly, but has a fully polynomial time approximation scheme. However, this approximation is not incentive compatible. Indeed, [16] who considered this problem were only able to show "almost incentive compatibility" [8]. Our results show that this is no accident. Exact incentive compatibility directly collides with an approximation scheme.

Corollary: *(Assuming $P \neq NP$ and a sufficiently powerful bidding language) No polynomial time incentive compatible mechanism for a multi-unit auction between two players that always allocates all units can achieve an approximation factor better than 2.*

1.4 Paper Organization

The rest of the paper is organized as follows (full details, proofs, and examples are given in the full paper [17]). In section 2 we describe our model. In section 3 we discuss the connection between monotonicity and truthfulness. Section 4 gives an exposition of the main theorem and its proof. Section 5 discusses the implications to computationally efficient combinatorial auctions. In the full paper, we also provide our alternative proof of Roberts' theorem.

2 Setting and Notations

2.1 Social choice functions on restricted domains

Social Choice Function. We study a general model of a social choice function $f : V_1 \times ... \times V_n \rightarrow A$. The interpretation is that f gets as its input a vector of *players' preferences* and chooses an alternative among a finite set of possible alternatives A. We denote $|A| = m$, and assume w.l.o.g that f is onto A.

The Domain (player types). Each player i ($1 \leq i \leq n$) assigns a real value $v_i(a)$ to each possible alternative from A. The vector $v_i \in R^m$ is called the player's *type* and is interpreted as specifying the player's preferences. The set $V_i \subseteq R^m$ is the set of possible valuations v_i. We denote $V = V_1 \times ... \times V_n$. We use the notation $v = (v_1, ..., v_n) \in R^{nm}$, and $v(a) = (v_1(a), ..., v_n(a)) \in R^n$. We also use the notation $v_{-i} = (v_1 ... v_{i-1}, v_{i+1} ... v_n) \in R^{n-1}$.

The main point in this paper is that V_i may be a proper subset of R^m. Here are some of the domains that we are concerned with in this paper:

- **Unrestricted Domains.** We say that the domain is *unrestricted* if $V_i = R^m$. In other words, the value of alternative a for player i does not place any restrictions upon i's values for the other alternatives.

[5] E.g.: for general combinatorial auctions any complete bidding language that can succinctly express single-minded bids is enough; if the number of players is a fixed constant, the language must allow OR-bids; for multi-unit (non-combinatorial) auctions, the bidding language must allow specifying the number of items in binary.

[6] This is essentially equivalent to the property of a "reasonable" auction of [25].

[7] E.g., where all items are sold as a single bundle in a simple auction – this clearly gives a factor $\min(n, k)$-approximation. Slightly better approximations in polynomial time are possible by partitioning the items into a constant number of bundles [15].

[8] A somewhat similar notion of "almost incentive compatibility" for an approximation scheme for a different problem was also obtained in [1].

- **Combinatorial Auctions (CA).** In a combinatorial auction, a set Ω of k items are auctioned between n bidders. The "alternatives" that the auction chooses among are allocations of items to bidders. That is, an alternative a is an allocation $a = (a_1...a_n)$, where $a_i \subseteq \Omega$ is the set of items allocated to player i, and $a_i \cap a_j = \emptyset$ for $i \neq j$ (each item can be allocated to at most one player). The valuations are assumed to satisfy three conditions:

  1. No externalities: v_i should only depend on i's allocated bundle a_i. I.e. $v_i(a) = v_i(a_i)$.
  2. Free disposal: v_i should be non-decreasing with the set of allocated items. I.e. for every $a_i \subseteq b_i$, we have that $v_i(a_i) \leq v_i(b_i)$.
  3. Normalization: $v_i(\emptyset) = 0$.

- **Multi Unit Auctions.** This is the special case of combinatorial auctions where the items are homogeneous. In this case an allocation $(a_1...a_n)$ is simply a vector of nonnegative integers, subject to the restriction that $\sum_i a_i \leq k$, and the valuation functions v_i can be represented as non-decreasing non-negative functions $v_i : \{1...k\} \rightarrow R_+$.

- **Order-Based Domains.** We will phrase our results in this paper in terms of a general family of domains termed "order-based", which contains all the previous classes, as well as others. These are domains where each V_i is defined by a (finite) family of inequalities and equalities of the form $v_i(a) \leq v_i(b)$, $v_i(a) < v_i(b)$, $v_i(a) = v_i(b)$ or $v_i(a) = 0$. Thus for example an unrestricted domain is defined by the empty family, while the domain of valuations for combinatorial auctions is defined by the following set of inequalities: for all $a, b \in A$ such that $a_i = b_i$: $v_i(a) = v_i(b)$ (no externalities); for all $a, b \in A$ such that $a_i \subseteq b_i$: $v_i(a) \leq v_i(b)$ (free disposal); for all $a \in A$ such that $a_i = \emptyset$: $v_i(a) = 0$.

- **Strict Order-Based Domains.** A subset of order-based domains for which we can prove strong statements is those defined only by strict inequalities $v_i(a) < v_i(b)$, as well as at most a single equality of the form $v_i(a) = 0$. Examples of strict order-based domains are two-players combinatorial auctions, or two-player multi-unit auctions, where all items must be allocated, i.e. $a_1 \cup a_2 = \Omega$ (this is discussed in details in section 5). Trivially, unrestricted domains are also strict order based.

2.2 Implementation and Truthfulness

Mechanisms. We assume that players' valuations are private information. Thus, a player might be motivated to declare a different type than his true type, in order to shift the social choice in some direction desirable for him. One solution is to construct a *mechanism*, which is allowed to charge payments ($p_i : V \rightarrow \mathcal{R}$) from the players, in addition to producing the chosen alternative. We assume that players are quasi-linear and rational in the sense of maximizing their total utility: $u_i = v_i(f(v)) - p_i(v)$. In *truthful* mechanisms, a player is motivated to be truthful and declare his true type, v_i, rather than a different type, u_i:

Definition 1 (Truthfulness) [9] *A mechanism* $(f, p_1...p_n)$, *where* $f : V \rightarrow A$ *and* $p_i : V \rightarrow \mathcal{R}$ *is called* truthful *if for any player* i, *any* $v_{-i} \in V_{-i}$, *and any* $v_i, u_i \in V_i$: $v_i(f(v)) - p_i(v) \geq v_i(f(u_i, v_{-i})) - p_i(u_i, v_{-i})$. *We say that such a mechanism* implements *the social choice function* f. *We say that the social choice function* f *is* implementable *or* truthful *if there exists some mechanism that implements it.*

The only known general class of truthful social choice functions over multi-dimensional domains are *affine maximizers*, which can be implemented using the VCG mechanism:

Affine maximization. A social choice function f is an *affine maximizer* if there exist constants $\omega_1, \ldots, \omega_n \geq 0$ and $\{\gamma_a\}_{a \in A}$ such that for any $v \in V$: $f(v) \in \text{argmax}_{a \in A}\{\sum_{i=1}^n \omega_i v_i(a) + \gamma_a\}$. It can be verified that, in this case, f is implemented by the payments $p_i = -\omega_i^{-1}(\sum_{j \neq i}^n \omega_j v_j(a) + \gamma_a)$.

3 Truthfulness and Monotonicity

It is well known that truthfulness is related to some notions of monotonicity. In this section we derive these relationships which serve as the embarking point towards our main characterization. All proofs are given in the full paper [17].

3.1 Weak monotonicity

In simple "one parameter" domains, monotonicity is usually the property of "still winning when raising my value". In general domains, we must examine value *differences*. Roberts [30] used a definition of monotonicity called PAD: f satisfies PAD if for every $v \in V$, $f(v) = a$, then $u_i(a) - v_i(a) > u_i(b) - v_i(b)$ for all $i = 1, \ldots, n$ and all $b \in A$ implies that $f(u) = a$. However, PAD has no real meaning for most restricted domains: Suppose there exists a player i and two alternatives a, b s.t. $v_i(a) = v_i(b)$ for all $v_i \in V_i$ (e.g. in CA, when i gets the same bundle in a and b). Then the condition of PAD is never satisfied if

[9] In this paper we only discuss direct revelation mechanisms with dominant strategy implementations in quasi-linear private value domains.

$f(v) = a$. One can make several attempts to "fix" this. In the full paper [17] we describe several natural "candidates" for a more general monotonicity condition, and demonstrate that they fail to be necessary for truthfulness. We identify the "correct" notion of monotonicity:

Definition 2 (Weak Monotonicity (W-MON)) *A social choice function f satisfies W-MON if for any $v \in V$, player i, and $u_i \in V_i$: $f(v) = a$ and $f(u_i, v_{-i}) = b$ imply that $u_i(b) - v_i(b) \geq u_i(a) - v_i(a)$.*

In other words, if player i caused the outcome of f to change from a to b by changing his valuation from v_i to u_i, then it must be that i's value for b has not increased less than i's value for a. W-MON implies PAD on every domain (see the full paper [17]), but makes sense also in domains where PAD does not. It is essentially equivalent to truthfulness:

Theorem 1 *Every implementable social choice function satisfies W-MON. For order based domains, W-MON is also a sufficient condition for truthfulness.*

The condition that the domain is order-based is needed (although it may be relaxed[10]) to ensure that W-MON is a sufficient condition, as the following example shows (this example is inspired by [32]):

Example 1 *Consider a single player with $A = \{a, b, c\}$ and a domain of three possible types v_a, v_b, v_c, as follows: $v_a = (0, 1, -2)$; $v_b = (-2, 0, 1)$; $v_c = (1, -2, 0)$, where the first coordinate in each type is a's value, the second is b's value, and the third c's value.*

The function f has $f(v_x) = x$, for every $x \in A$. f satisfies W-MON since $v_x(x) - v_y(x) > v_x(y) - v_y(y)$ for any $x, y \in A$.

Suppose by contradiction that there are truthful prices. Therefore: $-1 = v_c(c) - v_c(a) \geq p(c) - p(a)$. Similarly, $-1 = v_a(a) - v_a(b) \geq p(a) - p(b)$, and $-1 = v_b(b) - v_b(c) \geq p(b) - p(c)$. But the last two inequalities imply $p(c) - p(a) \geq 2$, a contradiction.

3.2 Strong monotonicity and IIA

So far we have seen that weak monotonicity is almost equivalent to truthfulness. We identify the following slightly stronger monotonicity condition, where the inequality in the definition is strict, as being of particular importance. We require this stronger condition for our main result.

Definition 3 (Strong Monotonicity (S-MON)) *A social choice function f satisfies S-MON if for any $v \in V$, player i, and $u_i \in V_i$: $f(v) = a$ and $f(u_i, v_{-i}) = b \neq a$ imply that $u_i(b) - v_i(b) > u_i(a) - v_i(a)$.*

[10] We conjecture that W-MON is sufficient for truthfulness on every *convex* domain. We are able to prove this if there are at most three alternatives.

In both definitions, we have the situation that i's valuation changed from v_i to u_i and this caused the outcome of f to change from a to b. S-MON asserts that this implies that i's valuation of b had to increase more than did the valuation of a. W-MON only requires that it did not increase less. While this seems like a slight change, it is in fact crucial. S-MON is not a necessary condition for truthfulness, and we give several counter examples in section 5, in the context of Combinatorial Auctions. The following definition, inspired by Arrow's notion for non-quasi-linear environments [4], essentially characterizes the difference between W-MON and S-MON:

Definition 4 (Independence of Irrelevant Alternatives (IIA)) *f satisfies IIA if for any $v, u \in V$, if $f(v) = a$ and $f(u) = b \neq a$ then there exists a player i such that $u_i(a) - u_i(b) \neq v_i(a) - v_i(b)$.*

In other words, if the social choice function on some valuations prefers a over b (i.e. a is chosen), and no player changes his preference of a with respect to b then it cannot be the case that the social choice function would prefer b over a.

Proposition 1 *W-MON + IIA $\Rightarrow$ S-MON. For order based domains, S-MON $\Rightarrow$ W-MON + IIA.*

Remark: We actually show that, for order based domains, S-MON implies the following "generalized S-MON": $f(v) = a$ and $f(u) = b \;\Rightarrow\; \exists i : \; u_i(b) - u_i(a) > v_i(b) - v_i(a)$. (Clearly this implies IIA). We also give some degenerate domain for which S-MON holds but IIA does not hold.

For certain domains, the notions of W-MON and S-MON are essentially equivalent. This equivalence is true for unrestricted domains and may be viewed as implicit in the proof of Roberts.

Theorem 2 *If V is strict order based then for every social choice function f there exists a function $\tilde{f} : V \to A$ such that, if f satisfies W-MON then $\tilde{f}$ satisfies S-MON, and then, if $\tilde{f}$ is an affine maximizer then f is an affine maximizer as well.*

4 Main Theorem

In what follows we give an exposition of our main theorem and its proof. A full description is given in the full paper [17].

Our main theorem shows that, under certain conditions, social choice functions that satisfy S-MON are "almost" affine maximizers. Let us first explain these conditions and qualifications:

- **The Domain:** The theorem holds for some family of domains which we call *order-based domains with conflicting preferences* – basically domains that satisfy various technical requirements in the proof. Combinatorial auctions, multi-unit auctions, and unrestricted domains all satisfy these requirements.

- **The Range:** The actual range of the social choice function must be *non-degenerate*. For combinatorial auctions or multi-unit auctions this means that there exists some player (w.l.o.g player 1) such that, for every other player i, the range includes an allocation a with $a_1 \neq \emptyset$ and $a_i \neq \emptyset$. For unrestricted domains this means that the range is of size at least 3. Without this condition, the problem may essentially be reduced to a single-dimensional setting, in which case many non affine maximizer truthful choice functions exist.

- **The Social Choice Function:** We require *player decisiveness*. For CAs and MUAs, this means that a player can always receive all goods if he bids high enough on them. [11]

- **Almost Affine Maximizer:** The theorem only shows that the social choice function must be an affine maximizer for large enough input valuations. I.e. there exists a threshold M s.t. the function is an affine maximizer if $v_i(a) \geq M$ for all a and i (except from inherently zero alternatives). We believe that this restriction is just a technical artifact of the current proof.

Theorem 3 *Every social choice function over an order-based domain with conflicting preferences and onto a non-degenerate range, that is player decisive and satisfies S-MON, must be almost affine maximizer.*

We now provide an intuitive outline of the proof. Full details appear in the full paper [17]. We start by attempting to infer some order that f induces on the domain. Specifically, if for some vector v of preferences the choice is $a = f(v)$ then we may say that the vector of values $v(a) = (v_1(a), \ldots, v_n(a))$ has more weight than the vector $v(b)$. This leads us to the following definition:

Definition 5 *("x at a" is larger than "y at b") For $a, b \in A$ and $x, y \in \mathcal{R}^n$ we say that $x@a > y@b$ if there exists $v \in V$ such that $v(a) = x,\ v(b) = y$, and $f(v) = a$.*

This notation certainly suggests that ">" is an order. In unrestricted domains this is indeed the case. However, in restricted domains, it is not generally so. The requirements of the theorem imply "just enough" of the properties of an order to proceed with the proof.

Once such a "near-order" is defined, we can compare every $x@a$ to multiples of some fixed reference $z@c$. This is inspired by the "min-function" model of Archer and Tardos [2]. We would expect that for small values of α we would have $x@a > (\alpha z)@c$, while for large values of α we would have $x@a < (\alpha z)@c$. The value of α where the change happens somehow summarizes the "weight" of $x@a$. To proceed we need to find such c and z where this holds for "enough" x and a. From now on, let such appropriate c and z be fixed.

Definition 6 *The "measure of x at a" is defined as:*

$$m(x@a) = inf\ \{\ \alpha \mid x@a < (\alpha \cdot z)@c\ \}$$

This measure captures the choice function, as the following property shows:

Claim: *Under the conditions of the theorem, if $m(v(a)@a) < m(v(b)@b)$ then $f(v) \neq a$.*

This claim basically shows that $f(v) \in \text{argmax}_a\{m(v(a)@a)\}$. What remains to show is that $m(x@a)$ is in fact a linear function (in x) on R^n! (And, that it does not depend on a, up to an additive constant.) To get this result let us, informally, consider the partial derivative $\partial m(x@a)/\partial x_i$. A key observation is that this partial derivative must be equal to $\partial m(y@b)/\partial y_i$ for any other "compatible" y and b. Let us see the intuition for this: consider some v such that $v(a) = x$ and $v(b) = y$. Since the S-MON requirement only looks at differences $v_i(a) - v_i(b)$ when "choosing between a and b", we would expect that adding a constant δ to both $x_i = v_i(a)$ and to $y_i = v_i(b)$ will also leave $m(x@a) - m(y@b)$ unchanged. This is indeed the case:

Claim: *Under the conditions of the theorem, for all (appropriate) a, b, x, y and δ we have that $m(x@a) - m(y@b) = m((x + \delta \cdot e_i)@a) - m((y + \delta \cdot e_i)@b)$. (Here e_i is the i'th unit vector).*

But now we claim that this means that $\partial m(x@a)/\partial x_i$ is independent of x! To see this, fix some y and denote $h_i(\delta) = m((y + \delta \cdot e_i)@b) - m(y@b)$. The previous claim states that $m((x + \delta \cdot e_i)@a) - m(x@a) = h_i(\delta)$. It is a simple exercise to verify that such a condition on $m(x@a)$ implies that it is linear in x. Specifically it must have the form $m(x@a) = \sum_i h_i(1) \cdot x_i + \gamma_a$, where γ_a is an arbitrary constant. This is (almost) the required result. (Delicate difficulties enter when we are unable to choose a single $y@b$ in an appropriate way for all $x@a$ – those are treated in the full proof.)

[11] We note the difference between this requirement and the decisiveness requirement of [20], where it is required that some player will be able to cause *any* alternative to be chosen, when declaring appropriately. For CAs, this is very strong – for example, it requires that player 1 will be able to decide whether player 2 or player 3 will receive all goods.

5 Computational Implications for Combinatorial Auctions

In this section we discuss the applicability of the main theorem to the main motivating problem: computationally efficient truthful mechanisms for approximating the optimal allocation in combinatorial auctions (CAs) and multi-unit auctions (MUAs). For this application most of the technical issues in the main theorem can be dropped. We start dealing with the general issues, proceed with those implied by approximation factors, and conclude with the computational ones.

5.1 General Issues

CAs and MUAs satisfy all the requirements on the domain of theorem 3. Thus any CA or MUA that satisfies S-MON and player decisiveness, onto a non-degenerate domain, must be almost affine maximizer. In fact, a non-degenerate domain captures even the case where each bidder is interested in only two, known in advance, bundles ("known double minded bidders"), where one of bundles is the set of all goods.

Let us now look at the different requirements of the theorem. First notice that if the range is degenerate, then as discussed above, the social choice function need not be an almost affine maximizer[12]. As for the strong monotonicity, in the full paper [17] we show that there are truthful CAs that do not satisfy S-MON, and indeed are *not* affine maximizers. (But are still "close" to affine maximizers).

When there are two players, and all the goods are always allocated, then S-MON is no longer a burden: in this case, for any distinct allocations a and b we have that $a_i \neq b_i$. Thus, V is very close to being strictly ordered, so we expect to be able to use Theorem 2 to reduce S-MON to W-MON. Specifically, define the *interior* of V to be $\overset{\circ}{V} = \{v \in V \mid v_i(a) < v_i(b)$ for all $a, b \in A$ s.t. $a_i \subsetneq b_i\}$ and define $\overset{\circ}{f} : \overset{\circ}{V} \to A$ by $\overset{\circ}{f}(v) = f(v)$.

Theorem 4 *Fix any truthful CA or MUA f for two players, that always allocates all the goods. Suppose that $\overset{\circ}{f}$ is player decisive and onto a non-degenerate range. Then:*

1. *f must be almost affine maximizer in the interior of V.*
2. *If the $\{\gamma_a\}_{a\in A}$'s are all zero then f is almost affine maximizer in all of V.*

[12] A simple class of examples is the "bundled auction" that allocates all items to the player with maximum value of $t_i(v_i(\Omega))$, where each t_i is an arbitrary monotone real function.

If we drop the assumption of always allocating all goods, then S-MON cannot be assumed w.l.o.g. In the full paper [17] we give an example of a specific CA that sometimes leaves unallocated goods, is player decisive and onto a full range, but does not satisfy S-MON, and is not almost affine maximizer.

5.2 Approximation

Since exact welfare optimization in CAs is computationally hard (see also below), we ask whether there exist truthful welfare approximations. A social choice function is a c-approximation of the optimal welfare if, for any type v, the alternative $f(v)$ has welfare of at least $1/c$ times the optimal welfare for v. For this class of functions, we are able to show that most of the qualifiers of the main theorem can be dropped.

Specifically, we define an auction to be *unanimity-respecting* (essentially equivalent to the notion of "reasonable" in [26]) if, whenever every player values only a single bundle a_i, and $a_i \cap a_j = \emptyset$ for all i, j, then f chooses the allocation $a = (a_1, \ldots, a_n)$. Using these, the "almost" qualifier and the player decisiveness property are dropped from the main theorem:

Lemma 1 *Any unanimity-respecting truthful CA or MUA that satisfies IIA and achieves a c-approximation must be an affine maximizer. Furthermore, the weights must satisfy $\gamma_a = 0$ for all alternatives a and and $(1/c) \leq (\omega_i/\omega_j) \leq c$ for all players i, j.*

For two players, where all the goods are always allocated, we can drop even the remaining qualifiers:

Lemma 2 *Any truthful CA or MUA for two players that always allocates all items and achieves an approximation factor of $c < 2$ must be an affine maximizer. Furthermore, it must have a full range, and the weights must satisfy $\gamma_a = 0$ far all a and $0.5 < (\omega_i/\omega_j) < 2$ for all i, j.*

5.3 Polynomial-Time Computation

All treatment of mechanisms so far assumed a fixed number of players n and a fixed number of items k. When formalizing the notion of computational running time we must let these parameters (or at least the number of items k) grow, and consider the running time as a function of them. A mechanism whose running time we wish to analyze would apply to all k and, if n is not fixed, for all n, i.e. would really be a uniform family of mechanisms. The characterization as affine maximizer above would then only apply to each mechanism in the family separately (with no explicit relationship across the different values of n and k.) However, to obtain computational hardness, we need that the constants of the different affine maximizers to be *polymoially bounded*, i.e. that for any fixed k, n, the relation between the affine constants of the mechanism will be bounded by $poly(k, n)$. Otherwise, for example, if the weight of player 1 is much larger than those of the other players, this depicts that player 1 will always receive all goods (for a polynomial input size).

In order to represent the mechanisms' running time as a function of its input size, we must fix an input representation for the valuations, i.e. a *bidding language* [24]. Our results apply to any such choice of a bidding language as long as it is complete (i.e. can represent all valuation) and sufficiently powerful. In fact, for claiming that affine maximization is as computationally hard as exact maximization, we only need the bidding language to have the following two elementary properties:

Definition 7 *A bidding language L is* **elementary** *if,*

1. *For any bid $b \in L$ that implicitly represents some valuation v, there exists a polynomial time procedure to construct a bid $b' \in L$ that represents the valuation $\alpha \cdot v$ (i.e. the value of each bundle is multiplied by α).*

2. *There exists a valid bid in which all bundles except Ω are valued as 0, and Ω is valued as α, for any $\alpha \geq 0$.*

For example, OR bids and XOR bids [24] are elementary: the first property is satisfied by just going over all the bid's blocks and multiplying their value by α.

We can now state formally that affine maximizers CAs and MUAs are as hard to compute as exact welfare maximizers:

Lemma 3 *Any affine maximizer CA or MUA with an elementary bid language, with polynomially bounded constants, and with the additive constants being equal to zero, is as computationally hard as the exact welfare maximization problem (with the same bidding language and the same range A).*

Our interest is in cases where the bidding language is sufficiently powerful as to make exact welfare maximization NP-complete. If the bid language forces the input to be long, e.g. the value of all possible bundles must be specified, then clearly we can construct an affine maximizer that will take linear time in the size of this input. Therefore, we need to allow short inputs. In particular, [19] show that as long as even single-minded bids are possible then the CA problem with n players is NP-complete (where n is not fixed). We observe that this is true for MUAs as well, as long as the number of desired items may be given in binary (rather than unary). When the number of players is fixed, then single-minded bids (as well as XOR-bids) may be handled in polynomial time, but we show that allowing OR bids results in an NP-complete optimization problem:

Claim 1 *Any welfare maximizing CA or MUA for n players (where n is not fixed), with full range, is NP-hard, even with single minded bids. If the number of players is fixed, then the above holds with OR bids as the bidding language.*

To integrate our main characterization with this computational hardness, we need a bidding language that will be rich enough to express all possible valuations, since the characterization does not assume any limitations on the possible valuation of the players. Notice that single minded bids and OR bids are not rich enough (OR bids can express only super-additive valuations).

Definition 8 *A bidding language L* **generalizes** *the bidding language L' if,*

1. *L contains all valid bids of L'.*

2. *L can express all possible player valuations.*

For example, XOR bids generalize single minded bids. And, OR bids with dummy items, and XOR of ORs, both generalize OR bids.

We can now integrate the above claims with our characterization of truthful welfare approximations:

Theorem 5 *Any Unanimity respecting truthful polynomial time combinatorial (or multi-unit) auction, with a bidding language that generalizes single minded bids, and that satisfies IIA, cannot obtain $poly(n, k)$ welfare approximation (unless $P = NP$).*

For the case of two-player auctions, we can omit the "unanimity-respecting" and "IIA" assumptions:

Corollary 1 *Any truthful polynomial-time multi-unit (or combinatorial) auction between two players, with a bidding language that generalizes OR bids, and that always allocates all goods, cannot obtain a welfare approximation better than 2 (unless $P = NP$).*

In contrast, for MUA without the truthfulness requirement there exists an FPAS [27]! Also notice that a truthful 2-approximation can be easily obtained using a simple auction of the bundle of all goods.

Acknowledgements

We thank Liad Blumrosen, Daniel Lehmann, Moritz Meyer-ter-Vehn, Benny Moldovanu, Dov Monderer, Motty Perry, and Amir Ronen for helpful discussions and comments.

References

[1] A. Archer, C. Papadimitriou, K. Talwar, and E. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *Proc. of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'03)*, 2003.

[2] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42st Annual Symposium on Foundations of Computer Science(FOCS'01)*, 2001.

[3] A. Archer and E. Tardos. Frugal path mechanisms. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, 2002.

[4] K. Arrow. *Social Choice and Individual Values*. Wiley, 1951.

[5] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi-unit combinatorial auctions. Working paper, 2002.

[6] L. Blumrosen and N. Nisan. Auctions with severely bounded communication. In *Proc. of the 43rd Annual Symposium on Foundations of Computer Science (FOCS'02)*, 2002.

[7] M. L. Breton and J. Weymark. Arrovian social choice theory on economic domains. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare: Volume 2*. North-Holland: Amsterdam, Forthcoming.

[8] E. Clarke. Multipart pricing of public goods. *Public Choice*, 8:17–33, 1971.

[9] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1), 2001.

[10] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *Proc. of the 34th ACM Symposium on Theory of Computing (STOC'02)*, 2002.

[11] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

[12] R. Gonen and D. Lehmann. Multi unit combinatorial auctions: Branch and bound heuristics. In *Proc. of the ACM Conference on Electronic Commerce (EC'00)*, 2000.

[13] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[14] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.

[15] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. Working paper, 2001.

[16] A. Kothari, D. Parkes, and S. Suri. Approximately-strategyproof and tractable multi-unit auctions. In *Proc. of the ACM Conference on Electronic Commerce (EC'03)*, 2003.

[17] R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. http://www.cs.huji.ac.il/ noam/partial.ps, 2003.

[18] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proc. of the ACM Conference on Electronic Commerce (EC'00)*, 2000.

[19] D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[20] M. Meyer-Ter-Vehn and B. Moldovanu. Ex-post implementation with interdependent valuations. Working paper, 2002.

[21] H. Moulin. *The Startegy of Social Choice*. North-Holland, 1983.

[22] A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *AAAI-02*, 2002.

[23] N. Nisan. Algorithms for selfish agents. In *Proc. of the 16th Symposium on Theoretical Aspects of Computer Science (STACS'99)*, 1999.

[24] N. Nisan. Bidding and allocation in combinatorial auctions. In *Proc. of the ACM Conference on Electronic Commerce (EC'00)*, 2000.

[25] N. Nisan and A. Ronen. Computationally feasible vcg mechanisms. In *Proc. of the ACM Conference on Electronic Commerce (EC'00)*, 2000.

[26] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001. Extended abstract appeared in STOC'99.

[27] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting lindahl prices. Working paper, 2002.

[28] N. Nisan and E. Zurel. An efficient approximate allocation algorithm for combinatorial auctions. In *Proc. of the ACM Conference on Electronic Commerce (EC'01)*, 2001.

[29] C. Papadimitriou. Algorithms, games, and the internet. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, 2001.

[30] K. Roberts. The characterization of implementable choise rules. In J.-J. Laffont, editor, *Aggregation and Revelation of Preferences. Papers presented at the first European Summer Workshop of the Econometric Society*, pages 321–349. North-Holland, 1979.

[31] A. Ronen and A. Saberi. Optimal auctions are hard. In *Proc. of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)*, 2002.

[32] I. Rozenshtrum. Dominant strategy implementation with quasi-linear preferences. Master's thesis, Dept. of Economics, The Hebrew University, Jerusalem, Israel, 1999.

[33] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. Cabob: A fast optimal algorithm for combinatorial auctions. In *Proc. of the 17th International Joint Conference on Artificial Intelligence (IJCAI'01)*, 2001.

[34] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[35] W. Vickrey. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.