Revenue Monotonicity in Deterministic, Dominant-Strategy Combinatorial Auctions

Baharak Rastegari*,[a], Anne Condon[a], Kevin Leyton-Brown[a]

[a] *Department of Computer Science, University of British Columbia*
2366 Main Mall, Vancouver, B.C., Canada, V6T 1Z4

Abstract

In combinatorial auctions using VCG, a seller can sometimes increase revenue by *dropping* bidders. In this paper we investigate the extent to which this counter-intuitive phenomenon can also occur under other deterministic dominant-strategy combinatorial auction mechanisms. Our main result is that such failures of "revenue monotonicity" can occur under *any* such mechanism that is *weakly maximal*—meaning roughly that it chooses allocations that cannot be augmented to cause a losing bidder to win without hurting winning bidders—and that allows bidders to express arbitrary single-minded preferences. We also give a set of other impossibility results as corollaries, concerning revenue when the set of goods changes, false-name-proofness, and the core.

JEL classification: D44

Key words: Mechanism design, Combinatorial auctions, Revenue

1. Introduction

In combinatorial auctions, multiple goods are sold simultaneously and bidders are allowed to place bids on bundles, rather than just on individual goods. These auctions are interesting in settings where bidders have non-additive—and particularly, superadditive—values for goods. (For an introduction, see Cramton et al. (2006).) As with other applications of mechanism design, the design of combinatorial auctions has tended to focus on the theoretical properties that a given design can guarantee. We begin with a discussion of such properties, specifically considering dominant strategy truthfulness, allocative efficiency, and revenue.

1.1. Dominant Strategy Implementation

One useful property for an auction mechanism is that it offers bidders the dominant strategy of truthfully revealing their private information to the mechanism. (By the revelation principle, the assumption that bidders declare *truthfully* is without loss of

*Corresponding author

Email addresses: `baharak@cs.ubc.ca` (Baharak Rastegari), `condon@cs.ubc.ca` (Anne Condon), `kevinlb@cs.ubc.ca` (Kevin Leyton-Brown)

generality; however, not all mechanisms offer dominant strategies.) Considerable research has characterized the space of social choice functions that can be implemented in dominant strategies. A classic result of Roberts (1979) showed that for bidders with unrestricted quasilinear valuations, affine maximizers are the only dominant-strategy-implementable social choice functions. Subsequent work has focused mainly on restricted classes of preferences (Rochet, 1987; Lavi et al., 2003; Bikhchandani et al., 2006; Saks and Yu, 2005; Constantin and Parkes, 2005).

The VCG mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973) has gained substantial attention in mechanism design literature because of its strong theoretical properties. In particular, it offers dominant strategies and achieves efficiency. Indeed, no substantially different (technically, no non-Groves) mechanism can guarantee these properties for agents with general quasilinear valuations (Green and Laffont, 1977). VCG is computationally intractable[1], thus, there have been many attempts to design feasible dominant strategy truthful mechanisms, even if for restricted classes of valuations. Archer and Tardos (2001), Andelman and Mansour (2006) and Mu'alem and Nisan (2002) studied the design of truthful mechanisms for combinatorial settings with single-parameter agents: agents whose private information can be encoded in a single positive real number. Babaioff et al. (2005, 2006) studied CA design in single-value domains under the further assumption that each agent has the same value for all desired outcomes. Yokoo et al. (2001, 2004) studied the design of truthful mechanisms for settings in which bidders may submit multiple bids using pseudonyms.

1.2. Allocative Efficiency

Much of the literature on combinatorial auctions has focused on the issue of achieving efficient allocations. Computing the efficient outcome in combinatorial auctions is an NP-complete problem. Therefore, finding the efficient allocation is not tractable in many combinatorial settings, and so approximation schemes can be useful. Sandholm (2002) showed that approximating efficiency in combinatorial auctions to within a factor of $n^{1-\varepsilon}$ for any fixed positive ε is NP-complete, where n is the number of bidders. The proof is based on a result of Håstad (1999). Lehmann et al. (2002) introduced a dominant strategy combinatorial auction mechanism for a restricted class called single-minded bidders. This mechanism runs in polynomial time and approximates efficiency by a factor of $\sqrt{m}$, where m is the number of goods for sale. Applying the result of Sandholm (2002) to the single-minded case, approximating efficiency to within a factor of $m^{1/2-\varepsilon}$ is NP-complete even if bidders are all single-minded (see, e.g., Nisan (2007)). Bartal et al. (2003) introduced approximately efficient dominant strategy truthful CA mechanisms for general valuations for both online (i.e., bidders arrive one at a time) and off-line scenarios that run in polynomial time. Lehmann et al. (2001) introduced a simple greedy 2-approximation algorithm for the class of valuations with decreasing marginal utilities. Later, Dobzinski and Schapira (2006) provided an approximation algorithm with the ratio of $\frac{1}{1-1/e}$ for this class of valuations and Khot et al. (2005) proved that $\frac{1}{1-1/e}$ is the lower bound for any polynomial-time approximation algorithm for this class.

[1]Indeed, VCG has a host of other drawbacks too; see, e.g., Rothkopf (2007).

Ascending auctions are another widely-studied family of CA mechanisms, primarily because they can reduce communication as compared to direct mechanisms. (For work on communication complexity of combinatorial auctions see, e.g., Segal (2006) and Nisan and Segal (2006).) Demange et al. (1986), Gul and Stacchetti (1999, 2000), Milgrom (2002), Ausubel (2006) and Bikhchandani et al. (2001) studied ascending item-price combinatorial auctions. Ascending bundle-price auctions were first suggested by Parkes and Ungar (2000) and Ausubel and Milgrom (2002). Hybrid designs appear in Kelly and Steinberg (2000) and Ausubel et al. (2006a).

1.3. Revenue

Besides the design of efficient allocation rules, the other main concern for auctioneers is an auction's revenue. VCG has some good revenue properties, at least among efficient mechanisms. Specifically, in settings where VCG is *ex post* individually rational, VCG collects at least as much revenue as any other efficient, *ex interim* individually rational mechanism, including those that do not offer dominant strategies (Krishna and Perry, 1998). However, it is possible to attain yet more revenue by relaxing efficiency. In the context of single-good auctions, Myerson (1981) and Riley and Samuelson (1981) characterized the *optimal auctions*, those mechanisms that in equilibrium maximize expected revenue under the bidders' valuation distribution. Other important early work includes Haris and Raviv (1981) and Maskin and Riley (1980). Recent work has studied the revenue properties of multigood auctions from experimental (see, e.g., Ledyard et al. (1997) and Englmaier et al. (2006)) and theoretical points of view (much of the latter from computer science). Goldberg and Hartline (2003) studied the design of an auction that achieves a constant fraction of the optimal revenue even on worst-case inputs in the unlimited supply (digital good) setting. More recent work links this prior-free approach to the more standard Bayesian setting (Hartline and Roughgarden, 2008). Ronen (2001) designed a unit-demand multi-unit auction that runs in polynomial time and is approximately optimal on expectation. Monderer and Tennenholtz (2005) gave an upper bound on the expected revenue from multi-object auctions with risk-averse bidders, and showed that under some additional assumptions VCG is asymptotically optimal as the number of bidders grows. Likhodedov and Sandholm (2005a,b) gave algorithmic methods for finding approximately optimal combinatorial auctions from the (VCG-like) family of affine maximizers. Balcan and Blum (2006) presented an approximation algorithm for optimal item pricing for single-minded bidders in the unlimited supply setting.

In this paper we focus on a particular revenue-related property: that a seller's revenue from an auction is guaranteed weakly to increase as the number of bidders grows. Ausubel and Milgrom (2002) dubbed this property *bidder monotonicity*. In order to emphasize that we are concerned with monotonicity of revenue—as compared to some other auction property—we prefer the term *bidder revenue monotonicity*. This can be contrasted with e.g., *good revenue monotonicity*, the property that a seller's revenue from an auction is guaranteed to weakly increase as the number of goods at auction grows. We are primarily interested in the former property; thus, as a shorthand we abbreviate bidder revenue monotonicity simply as *revenue monotonicity*.

It is easy to see that even VCG is not (bidder) revenue monotonic. Following an example due to Ausubel and Milgrom (2006), consider an auction with three bidders and two goods for sale. Suppose that bidder 2 values both goods at \$2 billion whereas bidder 1 and bidder 3 value the first and the second goods at \$2 billion respectively. The VCG mechanism awards the goods to bidders 1 and 3 for the price of zero, yielding the seller zero revenue. However, in the absence of either bidder 1 or bidder 3, the auction would generate \$2 billion in revenue.

Different approaches have been proposed to understanding the extent of revenue non-monotonicity problems. One approach has considered VCG's performance under restricted valuation classes. Say that *the combined valuation of bidders satisfies bidder submodularity* if and only if for any bidder i and any two sets of bidders S and S' with $S \subseteq S'$, it is the case that $V^*_{S\cup\{i\}} - V^*_S \geq V^*_{S'\cup\{i\}} - V^*_{S'}$, where V^*_S is the maximum social welfare achievable under S. Ausubel and Milgrom (2002) showed that if the combined valuation of bidders satisfies bidder submodularity then VCG is guaranteed to be revenue monotonic. Bidder submodularity is implied by the *goods are substitutes* condition (see, e.g., Ausubel and Milgrom (2002) for a definition). However, in many application domains for which combinatorial auctions have been proposed, goods are not substitutes and bidders' valuations exhibit complementarity. We therefore wish to investigate revenue monotonicity in domains where arbitrary complementarity may exist. The simplest such domain is that of single-minded bidders. Note that VCG is not revenue monotonic is this domain, as demonstrated in the above example.

Day and Milgrom (2007) showed that auctions that always select an outcome that is in the core with respect to declared valuations (so-called core-selecting auctions) are revenue monotonic when they select a core outcome that minimizes the seller's revenue. (A preliminary version of this result also appeared in Ausubel and Milgrom (2002).) Thus, the ascending proxy auction proposed by Ausubel and Milgrom (2002) and the clock-proxy auction proposed by Ausubel et al. (2006b) are both revenue monotonic but do not offer dominant strategies. Other mechanisms that have been proposed for use in practice similarly lack dominant strategies (see, e.g., Bernheim and Whinston (1986), Rassenti et al. (1982) and Porter et al. (2003)). We are not aware of any result in the literature that shows whether or not these mechanisms are revenue monotonic.

While revenue monotonicity *is* a feature of some auction mechanisms that have been deployed in practice, dominant strategies *are* (perhaps surprisingly) uncommon. This fact underscores the practical importance of revenue properties like revenue monotonicity, while pointing out that auctioneers are willing to sacrifice the strategic simplicity of dominant strategies.

1.4. Overview of Our Work

In our work, we ask whether there exists a combinatorial auction mechanism that allows bidders to express arbitrary single-minded preferences and that is both dominant strategy truthful and revenue monotonic.

If dominant strategy truthfulness and revenue monotonicity are the only conditions we require, it is easy to answer the above question in the affirmative. Specifically, we can offer all goods as one indivisible bundle using a second-price sealed-bid auction. However, this mechanism is unappealing, because it is combinatorial only in a degenerate sense. If we want to require the mechanism to allocate the goods more sensibly than through a static prebundling, we must rely on a further property something like efficiency. In this work[2], we exchange efficiency for the much more inclusive notion of *weak maximality*. While efficiency requires the mechanism to choose an allocation that maximizes social welfare, weak maximality requires that the mechanism choose an allocation that cannot be augmented to make some bidder better off, while making none worse off.

Our main contribution roughly states that, when bidders are allowed to express arbitrary single-minded preferences, no deterministic, dominant-strategy combinatorial auction mechanism is revenue monotonic, under some standard assumptions and the further assumption of weak maximality. As noted above, none of the auctions in practical use offer dominant strategies, while we know that at least some are revenue monotonic. Our impossibility result helps to explain this phenomenon: if revenue monotonicity is an important property in practice, deployed deterministic mechanisms will be unable to offer dominant strategies.

In Section 2 we define terminology for discussing combinatorial auction mechanisms and their properties. In Section 3 we define a restricted family of bidder valuations, define some properties of mechanisms for such bidders, and show that an existing, inefficient combinatorial auction mechanism for this setting (Lehmann et al., 2002) fails revenue monotonicity. We present our main impossibility result in Section 4. As corollaries, in Section 5 we prove similar impossibility results concerning the existence of mechanisms that yield weakly increasing revenue as the set of goods (rather than bidders) increases, that are false-name-proof (i.e., that offer truthful dominant strategies when agents are able to submit multiple bids under different identities), and that choose outcomes guaranteed to belong to the core.

2. Preliminaries

In this section we define terminology for discussing combinatorial auction mechanisms. In particular, we want to reason about changing the setting to include or exclude bidders/goods, which is difficult using traditional notations. Thus, we provide a general definition of mechanisms in which the allocation and payment rules may depend on which bidders participate, which goods are for sale, *ex-ante* knowledge the mechanism has about bidders' valuations (e.g., a single-minded bundle of interest), as well as on bidders' declared preferences. This detailed setting is necessary for the full formality of our claims and proofs; nevertheless, a reader who only skims Section 2.1 will be able to understand most of the details in what follows.

2.1. Bidders and Combinatorial Auction Mechanisms

Let $\mathbb{N} = \{1, 2, \ldots, n\}$ be the universal set of n bidders—all the potential bidders who exist in the world. Let $N \subseteq \mathbb{N}$ denote the set of bidders participating in a particular auction. Let $\mathbb{G}$ be the finite universe of goods for sale. Let $G \subseteq \mathbb{G}$ denote the set of goods for sale in a particular auction.

[2]We published a six-page preliminary version of our main result at a computer science conference (Rastegari et al., 2007). This work considered a very limited version of our weak maximality condition that can be understood as requiring Pareto efficiency with respect to bidder valuations (i.e., ignoring payments).

A *valuation function* describes the values that a bidder holds for subsets of the set of goods in G. Let valuation function $v_{\mathbb{G},i}$ for bidder $i \in \mathbb{N}$ map $2^{\mathbb{G}}$ to the nonnegative reals. For every $G \subset \mathbb{G}$ let valuation function $v_{G,i}$ be the projection of $v_{\mathbb{G},i}$ into G.

Whenever G is understood, we drop it from the subscript. We assume that bidders have quasilinear utility functions; that is, bidder i's utility for bundle a_i is $v_i(\mathrm{a}_i) - p_i$, where v_i is her valuation and p_i is any payment she is required to make.

A *valuation profile* is an n-tuple $v = (v_1, \ldots, v_n)$, where, for every participating bidder i, v_i is a valuation function. Let $\mathbb{V}$ denote the universal set of all possible valuation profiles. Observe that valuation profiles always have one entry for every potential bidder, regardless of the number of bidders who participate in the auction. We use the symbol $\varnothing$ in such tuples as a placeholder for each non-participating bidder (i.e., each bidder $i \notin N$). When v is an n-dimensional tuple, then $(v_1, \ldots, v_{i-1}, \varnothing, v_{i+1}, \ldots, v_n)$ is denoted by v_{-i}. Note that if $i \notin N$, then $v = v_{-i}$. Let $\mathbb{V}_{N,G}$ denote the set of all valuation profiles given a set of participating bidders N and a set of goods for sale G; that is, the set of all valuation profiles v_G for which $v_i = \varnothing$ if and only if $i \notin N$.

If asked to reveal her valuation, a bidder may not tell the truth. Denote the declared valuation function of a (participating) bidder i as $\widehat{v}_i$. Let $\widehat{v}$ be the declared valuation profile. Use the same notation to describe declared valuation profiles as valuation profiles (e.g., all declared valuation profiles are n-tuples), and furthermore write $(v_i, \widehat{v}_{-i})$ to denote $(\widehat{v}_1, \ldots, \widehat{v}_{i-1}, v_i, \widehat{v}_{i+1}, \ldots, \widehat{v}_n)$.

In a particular auction, bidders' valuation functions may be drawn from some restricted set. Let $V_{N,G} \subseteq \mathbb{V}_{N,G}$ denote a subspace of the universal set of valuation profiles for the set of participating bidders N and the set of goods for sale G. Let $\mathcal{V}_{\mathbb{N},\mathbb{G}}$ denote the universal set of valuation profile subspaces, that is $\mathcal{V}_{\mathbb{N},\mathbb{G}} = \{V_{N,G} \mid N \subseteq \mathbb{N}, G \subseteq \mathbb{G}, V_{N,G} \subseteq \mathbb{V}_{N,G}\}$. Let $\mathcal{V}$ denote a set of valuation profile subspaces with at least one member corresponding any $N \subseteq \mathbb{N}$ and $G \subseteq \mathbb{G}$. That is, $\mathcal{V} \subseteq \mathcal{V}_{\mathbb{N},\mathbb{G}}$ and $\exists V_{N,G} \in \mathcal{V}, \forall N \subseteq \mathbb{N}, G \subseteq \mathbb{G}$. Note that there could be more than one subspace corresponding to a fixed N and a fixed G in $\mathcal{V}$.[3]

We are now ready to define a combinatorial auction mechanism. Observe that our definition requires a mechanism to define allocations and payments for all possible sets of bidders, all possible sets of goods, and all corresponding valuation profiles belonging to a given, possibly restricted set. Also, note the implicit assumption that the auction setting—i.e., N, G and $V_{N,G}$—is common knowledge among all bidders and the auctioneer.

Definition 1 (CA Mechanism). *Let the set of valuation profile subspaces $\mathcal{V}$ be given. A* deterministic direct Combinatorial Auction (CA) mechanism *M (CA mechanism) maps each $V_{N,G} \in \mathcal{V}$, $N \subseteq \mathbb{N}$ and $G \subseteq \mathbb{G}$, to a pair (a, p) where*

- *a, the allocation scheme, maps each $\widehat{v} \in V_{N,G}$ to an allocation tuple $a =$*

[3] The reader might wonder why all this machinery is useful. We use it, for example, to model the case of "known single-minded" bidders (Section 3), in which each bidder values one bundle—that is known to the mechanism—and all its supersets at some amount v_i, and values all other bundles at zero. We can use $V_{N,G}$ to represent all valuations consistent with each bidder having a single-minded interest in one known bundle. $\mathcal{V}$ can describe subspaces corresponding to all the possible sets of known bundles for different bidders.

$(a_1(\widehat{v}), \ldots, a_n(\widehat{v}))$ of goods, where $\cup_i a_i(\widehat{v}) \subseteq G$, $a_i(\widehat{v}) \cap a_j(\widehat{v}) = \emptyset$ if $i \neq j$, and $a_i(\widehat{v}) = \emptyset$ if $\widehat{v}_i = \varnothing$.

- *p, the payment scheme, maps each $\widehat{v} \in V_{N,G}$ to a payment tuple $p = (p_1(\widehat{v}), \ldots, p_n(\widehat{v}))$, where $p_i(\widehat{v})$ is the payment from bidder i to the auctioneer such that $p_i(\widehat{v}) = 0$ if $\widehat{v}_i = \varnothing$.*

We say that CA mechanism M is defined for $\mathcal{V}$. We refer to (a, p) as the outcome of the CA mechanism. We refer to a_i and p_i as bidder i's allocation and payment functions respectively. Whenever $\widehat{v}$ can be understood from the context, we refer to $a_i(\widehat{v})$ and $p_i(\widehat{v})$ by a_i and p_i, respectively. If $\widehat{v}_i(a_i) > 0$, we say that bidder i "wins". We denote by $\mathbb{A}_{N,G}$ the set of all possible partitions of G into $|N| + 1$ partitions; i.e., the set of all possible ways of distributing goods among participating bidders and the auctioneer—note that the auction may allocate no good to any of the bidders. For any given allocation $\mathrm{a} \in \mathbb{A}_{N,G}$, we denote by a_i the set of goods that are allocated to bidder i under a.

2.2. Desirable Mechanism Properties

Now we survey properties that we would like to require of combinatorial auction mechanisms.

2.2.1. Dominant strategy truthfulness

In mechanism design, it is especially desirable for a mechanism to give rise to dominant strategies, as then there is no need for bidders to reason about each others' behavior in order to maximize their utilities.

A direct CA mechanism M is said to be *truthful* if bidders declare their true valuations to the mechanism in equilibrium. M is said to be *DS truthful* if it is a dominant strategy for every bidder to reveal her true preferences.

Definition 2 (Dominant-strategy truthfulness). *A CA mechanism M is* dominant strategy truthful *(or* DS truthful*) if and only if for all fixed set of participating bidders, it is a best response for each participating bidder to declare her true valuation regardless of the declarations of the other participating bidders. That is, for all $N \subseteq \mathbb{N}, G \subseteq \mathbb{G}$, $V_{N,G} \in \mathcal{V}$, $\widehat{v} \in V_{N,G}$ and for every bidder i we have that*

$$v_i(a_i(v_i, \widehat{v}_{-i})) - p_i(v_i, \widehat{v}_{-i}) \geq v_i(a_i(\widehat{v})) - p_i(\widehat{v}).$$

Observe that the revelation principle tells us that any social choice function that can be implemented in dominant strategies can also be implemented truthfully in dominant strategies. This means that our conflation of dominant strategies with truth-telling is without loss of generality. In fact, the revelation principle applies not only to implementation in dominant strategies but also to implementation in any equilibrium. That is, adding truthfulness does not change the space of implementable social choice functions.

2.2.2. Participation

It is natural to require that no bidder be made to make any payment unless she wins.

Definition 3 (Participation). *A truthful CA mechanism M satisfies* participation *if and only if for all $N \subseteq \mathbb{N}, G \subseteq \mathbb{G}$, $V_{N,G} \in \mathcal{V}$, and $v \in V_{N,G}$, $p_i(v) = 0$ for all bidder i for whom $v_i(a_i) = 0$ (i.e., who does not win).*

Unlike the property of individual rationality (IR), which requires roughly that no bidder has to make a payment more than her value for the bundle she gets, participation does not constrain the payments of bidders who win. Participation is therefore a weaker condition than IR.

2.2.3. Efficiency

As discussed earlier, one of the most commonly desired properties for an auction mechanism is efficiency. A CA mechanism is said to be efficient if it always chooses an allocation that maximizes the social welfare.

Definition 4 (Efficiency). *A CA mechanism M is efficient if its chosen allocation in equilibrium, a^*, maximizes the social welfare; that is, for all $N \subseteq \mathbb{N}, G \subseteq \mathbb{G}$, $V_{N,G} \in \mathcal{V}$, and $v \in V_{N,G}$,*

$$\mathrm{a}^* \in \arg\max_{\mathrm{a} \in \mathbb{A}_{N,G}} \sum_i v_i(\mathrm{a}_i).$$

2.2.4. Revenue Monotonicity

The revenue of an auction mechanism is the sum over all the payments made by the bidders to the auctioneer. Informally, an auction mechanism is revenue monotonic if, when a bidder drops out, the auctioneer never collects more money as a result.

Definition 5 (Revenue monotonicity). *A truthful CA mechanism M is bidder revenue monotonic (or* revenue monotonic *) if and only if for all $N \subseteq \mathbb{N}, G \subseteq \mathbb{G}$, $V_{N,G} \in \mathcal{V}$, $v \in V_{N,G}$ and for all bidders j,*

$$\sum_{i \in \mathbb{N}} p_i(v) \geq \sum_{i \in \mathbb{N} \setminus \{j\}} p_i(v_{-j}).$$

Our goal in this paper is to investigate whether broad families of dominant-strategy truthful CA mechanisms satisfy revenue monotonicity. As mentioned above, the only dominant-strategy truthful and efficient CA mechanisms are Groves mechanisms (Green and Laffont, 1977). We have already seen that VCG fails revenue monotonicity; therefore, efficiency is perhaps a strong condition to require. The following example, however, shows that revenue monotonicity is unsatisfyingly easy to achieve if we simply drop efficiency.

Consider the *set protocol*, a simple mechanism that offers all goods as one indivisible bundle and uses the second price sealed-bid auction to determine the winner and the payment. It is trivial to show that the set protocol is dominant-strategy truthful and satisfies participation. This mechanism is also revenue monotonic since dropping a bidder cannot cause the second-price bid to increase. However, the set protocol is a combinatorial auction only in a degenerate sense: it pre-bundles all goods and treats them as a single indivisible good. If we want to insist on doing otherwise, we need to require a property that is like, but weaker than, efficiency.

2.2.5. Maximality

We propose weakening our requirement of efficiency by instead requiring *maximality*. This property requires that whenever bidder i values any subset of goods s sufficiently highly, the mechanism never chooses allocations that could be augmented to satisfy i.

Definition 6 (Maximality). *A truthful CA mechanism M is* maximal with respect to bidder i *if and only if for all $N \subseteq \mathbb{N}$ where $i \in N$, for all $G \subseteq \mathbb{G}$, and for all $V_{N,G} \in \mathcal{V}$, there exists a set of nonnegative finite constants $\{\alpha_{N,G,i,s} \mid s \subseteq G\}$ such that the following holds. For all $v \in V_{N,G}$, M always chooses an allocation* a *where either:*

1. $v_i(\mathrm{a}_i) > 0$*; or*
2. *for all allocation a' with $v_i(\mathrm{a}'_i) > \alpha_{N,G,i,\mathrm{a}'_i}$, and $\mathrm{a}'_j = \mathrm{a}_j \setminus \mathrm{a}'_i$ for all $j \neq i$, it must be the case that for some j, $v_j(\mathrm{a}'_j) < v_j(\mathrm{a}_j)$.*

Intuitively, maximality ensures that the mechanism does not withhold any subset of goods, or give the goods away to the bidders who do not value them, when they are sufficiently valued by a losing bidder. By a losing bidder we mean a bidder who does not win. The quantities $\{\alpha_{N,G,i,s} \mid s \subseteq G\}$ can be thought of as bidder- and bundle-specific reserve prices. Note that the set protocol does not satisfy maximality with respect to any bidder, because the winning bidder may be given goods that she does not value, even if there exists another bidder who values these goods and bid more than an arbitrary constant amount.

Many interesting mechanisms are maximal. First, it is straightforward to show that efficiency implies maximality. Second, we show here that a broad class of affine maximizing mechanisms are maximal.

Affine maximizers generalize the idea behind the VCG mechanism's allocation rule (which aims to maximize the social welfare) by allowing the mechanism to restrict the set of possible allocations, to assign different non-negative weights ω_i to different players, and to assign different additive weights γ_{a} to different allocations.

Definition 7 (Affine maximizer). *A CA mechanism is an affine maximizer if for some $\mathbb{A}'_{N,G} \subseteq \mathbb{A}_{N,G}$, nonnegative $\{\omega_i\}_{i\in\mathbb{N}}$ and $\{\gamma_{\mathrm{a}}\}_{\mathrm{a}\in\mathbb{A}'_{N,G}}$, for all $N \subseteq \mathbb{N}, G \subseteq \mathbb{G}, V_{N,G} \in \mathcal{V}, v \in V_{N,G}$, its chosen allocation in equilibrium,* a^* *satisfies the following:*

$$\mathrm{a}^* \in \arg\max_{\mathrm{a}\in\mathbb{A}'_{N,G}} \left(\sum_i \omega_i v_i(\mathrm{a}_i) + \gamma_{\mathrm{a}}\right).$$

We call $\{\omega_i\}_{i\in\mathbb{N}}$ and $\{\gamma_{\mathrm{a}}\}_{\mathrm{a}\in\mathbb{A}'_{N,G}}$ the allocation parameters of affine maximizer M.

Theorem 8. *Let M be an affine maximizing truthful CA mechanism with finite allocation parameters $\{\omega_i\}_{i\in\mathbb{N}}$ and $\{\gamma_{\mathrm{a}}\}_{\mathrm{a}\in\mathbb{A}_{N,G}}$. Suppose that for some $i \in \mathbb{N}$, $\omega_i > 0$. Then M is maximal with respect to bidder i.*

Proof Let $\alpha_{N,G,i,s} = \frac{\max_{\mathrm{a}}\{\gamma_{\mathrm{a}}\}}{\omega_i}$, $\forall N \subseteq \mathbb{N}$ where $i \in N$, $\forall G \subseteq \mathbb{G}$ and $\forall s \subseteq G$. We prove that M is maximal with respect to bidder i. Assume for contradiction that M is not maximal with respect to i. Then, for some v, M's allocation scheme maps v to an allocation a that satisfies the following properties: (i) $v_i(\mathrm{a}_i) = 0$, and (ii) $\exists s \subseteq G$, $\exists \mathrm{a}' \in \mathbb{A}_{N,G}$: $\mathrm{a}'_i = s$ and $\forall j \neq i$, $\mathrm{a}'_j = \mathrm{a}_j \setminus s$ such that $v_i(\mathrm{a}'_i) > \alpha_{N,G,i,s}$ and $v_j(\mathrm{a}'_j) \geq v_j(\mathrm{a}_j)$.

From construction and (i) and (ii) we have that

$$\sum_{k\in N} \omega_k v_k(\mathrm{a}'_k) + \gamma_{\mathrm{a}'} > \sum_{k\in N\setminus\{i\}} \omega_k v_k(\mathrm{a}'_k) + \max\{\gamma\} + \gamma_{\mathrm{a}'} \geq \sum_{k\in N} \omega_k v_k(\mathrm{a}_k) + \gamma_{\mathrm{a}}.$$

Since M is affine maximizing it would not choose allocation a, giving us our contradiction. $\square$

Finally, all mechanisms that are *strongly Pareto efficient with respect to bidders' valuations* are also maximal. We define that term as follows.

Definition 9 (Strong Pareto efficiency with respect to bidders' valuations). *A mechanism is* strongly Pareto efficient with respect to bidders' valuations *if it chooses allocations that would be strongly Pareto efficient if monetary transfers were disallowed.*

Note that these allocations are a superset of those that are strongly Pareto efficient when transfers are permitted. Thus, a broader class of mechanisms achieves strong Pareto efficiency with respect to bidders' valuations than achieve strong Pareto efficiency.

2.2.6. Weak maximality

Now we define a weakened version of maximality. This version will be sufficient for our purposes in what follows. Since it is a weaker constraint, using it will make our impossibility result stronger.

Definition 10 (Weak Maximality). *A truthful CA mechanism M is* weakly maximal with respect to bidder i *if and only if for all $N \subseteq \mathbb{N}$ where $i \in N$, for all $G \subseteq \mathbb{G}$, and for all $V_{N,G} \in \mathcal{V}$, there exists a set of nonnegative finite constants $\{\alpha_{N,G,i,g} \mid g \in G\}$ such that the following holds. For all $v \in V_{N,G}$, M always chooses an allocation a where either:*

1. $v_i(\mathrm{a}_i) > 0$*; or*
2. *for all allocation a' with $v_i(\mathrm{a}'_i) > \alpha_{N,G,i,\mathrm{a}'_i}$, $|\mathrm{a}'_i| = 1$, and $\mathrm{a}'_j = \mathrm{a}_j \setminus \mathrm{a}'_i$ for all $j \neq i$, it must be the case that for some j, $v_j(\mathrm{a}'_j) < v_j(\mathrm{a}_j)$.*

Observe that the above definition is simply derived from the definition of maximality by restricting a'_i to be of size 1.

Our weak maximality property is conceptually related to the *reasonableness* condition of Nisan and Ronen (2000), which says that whenever an item is desired by a single agent only, that agent must receive the item. It is easy to see that if $\alpha_{N,G,i,s}$'s are all set to zero, then weak maximality implies reasonableness. However, reasonableness does not imply weak maximality. Consider the case where there are exactly two bidders who desire an item. Reasonableness still holds even if that item is never allocated to either of the agents, regardless of their declarations. However, such an allocation rule would violate weak maximality.

2.2.7. Consumer sovereignty

Roughly speaking, a mechanism satisfies consumer sovereignty if any bidder can win any bundle that according to her valuation space she may value above zero, as long as she bids high enough. In what follows let $(V_{N,G})_{-i}$ denote the set $\{v_{-i} | v \in V_{N,G}\}$.

Definition 11 (Consumer sovereignty[4]). *A CA mechanism M satisfies consumer sovereignty if and only if for all $N \subseteq \mathbb{N}$, $G \subseteq \mathbb{G}$, and $V_{N,G} \in \mathcal{V}$, $\forall i \in N$ and $\forall s \subseteq G$, for all $\widehat{v}_{-i} \in (V_{N,G})_{-i}$, there exists some finite amount $\in \mathbb{R}$, $k_i^s > 0$ such that if i reports that she values s and supersets of s at amount at least k_i^s and values any bundle that does not contain s at zero, then i is allocated at least s.*

It is useful at this point to contrast consumer sovereignty with maximality. Consumer sovereignty implies that by bidding at or above the critical value k_i^s, bidder i surely wins bundle s_i. In contrast, maximality does not imply that i necessarily wins s_i if she values it at or above the bidder-specific reserve price α_{N,G,i,s_i}.

3. Known Single-Minded Bidders

We now define a restricted class of valuation spaces of which we will make use in the proof of our main theorem. Then we define some properties of mechanisms designed for such bidders.

The class of *unknown single-minded bidders*, or simply *single-minded bidders*, was first introduced by Lehmann et al. (2002). Informally, a participating bidder i is single-minded if there exists a particular bundle b_i such that bidder i only has a nonzero valuation for bundles that contain b_i, and values all these bundles equally. *Known single-minded bidders*, an even more restricted bidder model, was first introduced by Mu'alem and Nisan (2002). A bidder i is known single-minded in a mechanism if she is single-minded and the mechanism knows her bundle of interest b_i.

Definition 12 (Single-minded bidder). *Bidder i is* single-minded *if her valuation function is defined as*

$$v_i(b'_i) = \begin{cases} v_i & b'_i \supseteq b_i \\ 0 & \textit{otherwise,} \end{cases}$$

where $v_i > 0$ and $b_i \subseteq \mathbb{G}$.

[4] Our definition follows that of Feigenbaum et al. (2002).

Note that a single-minded bidder i's valuation can be characterized by two parameters: $\langle b_i, v_i\rangle$. Therefore, we use $\langle b_i, v_i\rangle$ and v_i interchangeably when a bidder is single-minded. We let $\langle \widehat{b}_i, \widehat{v}_i\rangle$ denote the *declared valuation* of single-minded bidder i. When b_i is known to the mechanism—called the known single-minded bidder case—the valuation of bidder i can be characterized by the single parameter v_i, representing i's valuation for any superset of bundle b_i. Thus in this case we use v to denote a valuation profile for a group of single-minded bidders, $\widehat{v}_i$ to denote the declared valuation of a participating bidder i, and $\widehat{v}$ to denote a tuple consisting of declared valuations for participating bidders and $\varnothing$ symbols for non-participating bidders.

Let $N \subseteq \mathbb{N}$ and $G \subseteq \mathbb{G}$ be fixed. Let $b = (b_1, b_2, \dots, b_n) \in (2^G)^n$. If i is a participating bidder, let $V^{(b)}_{N,G,i}$ be the set of all possible single-minded valuation functions, taken over all possible choices of v_i, and otherwise let $V^{(b)}_{N,G,i} = \varnothing$. Let $V^{(b)}_{N,G} = V^{(b)}_{N,G,1} \times \dots \times V^{(b)}_{N,G,n}$. Then, $V^{(b)}_{N,G}$ is simply the space of valuation profiles in which participating bidders are all single-minded and each participating bidder i values bundle b_i.

Definition 13 (Set of valuation profile subspaces for known single-minded bidders). *Let $\mathcal{V}^{(ksm)}$ denote the set of valuation profile subspaces for known single-minded bidders,*

$$\mathcal{V}^{(ksm)} = \{V^{(b)}_{N,G} \mid N \subseteq \mathbb{N}, G \subseteq \mathbb{G}, b \in (2^G)^n\}.$$

A CA mechanism is then *defined for known single-minded bidders* if its set of valuation profile subspaces is $\mathcal{V}^{(ksm)}$. From the definition of CA mechanism (Definition 1), it follows that the allocation and payment functions depend on the set $V^{(b)}_{N,G} \in \mathcal{V}^{(ksm)}$ from which bidders' valuation profiles are drawn. Informally, b is *known*, since the allocation and payments depend on b. Observe that our definition requires that the mechanism be defined for *all* possible known single-minded valuations, not just for the set of bundles that a given set of bidders might value.

A set of valuation subspaces $\mathcal{V}$ *subsumes* another set of valuation subspaces $\mathcal{V}'$ if and only if for all $V'_{N,G} \in \mathcal{V}'$, there exists $V_{N,G} \in \mathcal{V}$ such that $V'_{N,G} \subseteq V_{N,G}$.

We can say that the class of mechanisms defined for $\mathcal{V}$ is a subset of the class of mechanisms defined for known single-minded bidders when $\mathcal{V}$ subsumes known single-minded valuations. The preceding claim is in fact true in a general sense, that is even if we replace known single-minded valuations with any $\mathcal{V}'$. The following lemma states it formally. On some level this result is obvious; however, we were not able to find any formal discussion of it in the literature and so present it here for completeness.

Lemma 14. *Every social choice function that can be implemented by a mechanism defined for $\mathcal{V}$ can also be implemented by a mechanism defined for $\mathcal{V}'$, when $\mathcal{V}$ subsumes $\mathcal{V}'$.*

Proof Without loss of generality (see Section 2.2.1), we can restrict the proof to truthful mechanisms. The allocation function in a truthful mechanism is precisely the social choice function. Let $M^{(\mathcal{V})}$ be a truthful mechanism defined for $\mathcal{V}$. Modify $M^{(\mathcal{V})}$ such that given declared valuation profile $\widehat{v} \in V'_{N,G}$, $V'_{N,G} \in \mathcal{V}'$, runs the same allocation and payment functions as $M^{(\mathcal{V})}$ would run on $\widehat{v} \in V_{N,G}$, $V_{N,G} \in \mathcal{V}$, where

$V'_{N,G} \subseteq V_{N,G}$. As $\mathcal{V}$ subsumes $\mathcal{V}'$ such $V_{N,G}$ exists. Let $M^{(\mathcal{V}')}$ be this new mechanism that is defined for $\mathcal{V}'$. $M^{(\mathcal{V}')}$ is clearly truthful as is $M^{(\mathcal{V})}$. □

Informally speaking, for each mechanism M defined for $\mathcal{V}$, there is a corresponding mechanism M' defined for $\mathcal{V}'$. We will use the above claim in Section 4 to state our result for general CA mechanisms.

3.1. Criticality and Consumer Sovereignty for Mechanisms Defined For Known Single-Minded Bidders

Consider a mechanism defined for known single-minded bidders. We say that the mechanism *offers critical values to bidder* i if two properties hold. First, bidder i wins whenever she bids more than some critical value that depends only on the other bidders' declarations, and loses whenever she bids less. Second, bidder i's payment is equal to the aforementioned critical value if she wins, and is zero otherwise. A mechanism defined for known single-minded bidders satisfies criticality if it offers critical values to all bidders.

In what follows let $(V^{(b)}_{N,G})_{-i}$ denote the set $\{v_{-i} | v \in V^{(b)}_{N,G}\}$.

Definition 15 (Criticality). *A CA mechanism M defined for known single-minded bidders satisfies* criticality *if and only if for all $N \subseteq \mathbb{N}$, $G \subseteq \mathbb{G}$, and $V^{(b)}_{N,G} \in \mathcal{V}^{(ksm)}$, for all $i \in N$ and $\widehat{v}_{-i} \in \left(V^{(b)}_{N,G}\right)_{-i}$, there exists a* critical value $cv_i(\widehat{v}_{-i}) \in \mathbb{R}$ *where:*

- *if $\widehat{v}_i > cv_i(\widehat{v}_{-i})$, i wins and pays $cv_i(\widehat{v}_{-i})$;*
- *if $\widehat{v}_i < cv_i(\widehat{v}_{-i})$, i loses and pays 0.*

From necessary and sufficient conditions for dominant-strategy truthfulness (see, e.g., Mu'alem and Nisan (2002) and Nisan (2007)), it is straightforward to show that dominant-strategy truthful combinatorial auction mechanisms defined for known single-minded bidders that satisfy participation must also satisfy criticality.

Theorem 16 (Following Lehmann et al. (2002) and Mu'alem and Nisan (2002)). *Any CA mechanism defined for known single-minded bidders that satisfies dominant-strategy truthfulness and participation also satisfies criticality.*

The following corollary, which immediately follows from Definition 11 and Theorem 16, is used in the proof of our main theorem.

Corollary 17. *Any CA mechanism defined for known single-minded bidders that satisfies dominant-strategy truthfulness, participation and consumer sovereignty offers finite critical values to all bidders.*

4. Impossibility of Revenue Monotonicity

In this section we turn to our main claim, that no CA mechanism can be revenue monotonic if it satisfies our desired properties of dominant-strategy truthfulness, participation, consumer sovereignty and weak maximality with respect to at least two bidders. We begin by giving an example of how an existing inefficient mechanism fails revenue monotonicity, and then prove the general result.

4.1. An Example with an Inefficient Mechanism

In the introduction we gave a well-known example showing that VCG does not satisfy revenue monotonicity. Now we show—we believe, for the first time—that another widely-studied mechanism also fails revenue monotonicity, even though it does not have an efficient allocation rule. This example is an application of our impossibility theorem and so is not of independent interest; however, it offers intuition for what follows.

Lehmann et al. (2002) introduced an inefficient, dominant-strategy truthful, direct CA mechanism for single-minded bidders. Naming it after its authors, we call the mechanism LOS. Like VCG, LOS satisfies participation and consumer sovereignty. LOS is also strongly Pareto efficient with respect to bidders' valuations (see Definition 9), and so satisfies maximality (and hence, also weak maximality) with respect to all bidders.

Let $ppg_i = v_i/|b_i|$, bidder i's declared price per good. LOS ranks bids in a list L in decreasing order of ppg, and then greedily allocates bids starting from the top of L. Thus, each bidder i's bid is granted if b_i does not conflict with any previously allocated bids. If i's bid is allocated she is made to pay $|b_i| * v_{inext}/|b_{inext}|$ where $inext$ is the first bidder following i in L whose bid was denied but would have been allocated if i's bid were not present. Bidder i pays zero if she does not win or if there is no bidder $inext$.

Consider three bidders $\{1, 2, 3\}$ and two goods $\{g_1, g_2\}$. Let the true valuations of bidder 1, 2 and 3 be $\langle\{g_1\}, v_1\rangle$, $\langle\{g_1, g_2\}, v_2\rangle$ and $\langle\{g_2\}, v_3\rangle$, respectively. Now consider the following conditions on the bidders' valuations: (1) $v_1 > v_3 > v_2/2$; (2) $v_2 > 0$. It is possible to assign values to the v_i's in a way that satisfies both conditions: e.g., $v_1 = 5$, $v_2 = 4$ and $v_3 = 3$.

We will demonstrate that the auctioneer's revenue under LOS can be increased by dropping a bidder, whenever the bidders and their valuations are as described above. From Condition 1, $ppg_1 > ppg_3 > ppg_2$ and therefore bidders 1 and 3 win. Each pays zero, so the total revenue is zero. To see this, note that the next bidder in the list after bidder 1 whose bid conflicts with b_1 is bidder 2. However, bidder 2 would not win even if bidder 1 were not present, since b_2 also conflicts with b_3. Therefore bidder 1 pays zero. The same is true for bidder 3, and thus she also pays zero. If bidder 1 is dropped, bidder 3 wins and must pay $ppg_2 = v_2/2$. Since $v_2 > 0$ (Condition 2), this payment is more than zero and so revenue monotonicity fails.

4.2. Impossibility Theorem

We first prove a strong form of the theorem, for mechanisms defined for known single-minded bidders. Then we state a weaker form of the theorem for general CA mechanisms, which follows directly from the strong form.

Theorem 18. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism defined for known single-minded bidders that offers dominant strategies to the bidders and satisfies participation,[5] consumer sovereignty, and weak maximality with respect to at least two bidders. Then M is not revenue monotonic.*

[5] In the case of single-parameter domains—which includes known single-minded bidders—one can get

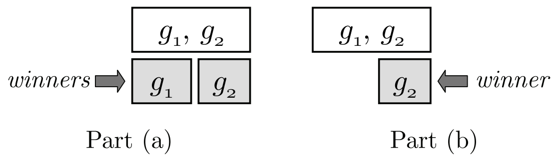

Figure 1: A high-level illustration of Theorem 19: Given $\langle\{g_1\}, v_1\rangle$, $\langle\{g_1, g_2\}, v_2\rangle$ and $\langle\{g_3\}, v_3\rangle$—$v_i$'s as constructed in the proof of the theorem—(a) Bidders 1 and 3 win bundle $\{g_1\}$ and bundle $\{g_2\}$ respectively and each pay more than a predefined constant amount, (b) bidder 3 wins bundle $\{g_2\}$ and pays more than the sum of the payments in part (a).

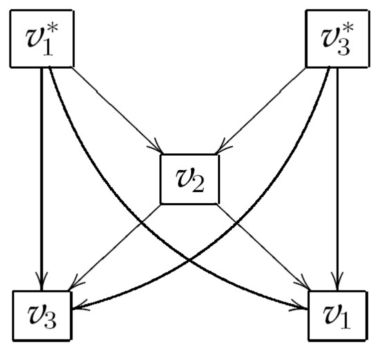

Figure 2: Illustration of dependencies between the constructed values in the proof of Theorem 19

Proof Without loss of generality (by the revelation principle) assume that M is dominant-strategy truthful. We will further assume that bidders follow their dominant strategies and bid truthfully. Since $|\mathbb{N}| \geq 3$, there are at least three bidders; let us name the first three 1, 2 and 3. (For notational simplicity in what follows, we will write the proof as though $|\mathbb{N}| = 3$. If in fact $|\mathbb{N}| > 3$ our argument does not change, but all valuation profiles must include extra $\varnothing$ entries.) Assume without loss of generality that M is weakly maximal with respect to bidders 1 and 3. Since $|\mathbb{G}| \geq 2$, there are at least two goods; let us name the first two g_1 and g_2. Let the bundles valued by bidders 1, 2, and 3 be $b_1 = \{g_1\}$, $b_2 = \{g_1, g_2\}$ and $b_3 = \{g_2\}$ respectively. Throughout we fix $G \subseteq \mathbb{G}$, subject to $g_1, g_2 \in G$

We now show how to construct valuations for the three bidders. First pick an arbitrary positive constant k, and then define $v_1^* = \alpha_{\{1,2,3\},G,1,g_1} + k$ and $v_3^* =$

participation for free. That is, the space of social choice functions that are implementable in dominant strategies is the same with or without adding a participation constraint. This is mainly because each bidder has to pay either of the two specific amounts: one if she wins and one if she loses. If we "normalize" the payment function and unconditionally pay each bidder the losing amount—which could be negative—then we achieve a dominant-strategy mechanism that satisfies participation. However, there are revenue implications to these unconditional payments that vary as the number of bidders in the auction varies. Therefore, we nevertheless state the participation condition explicitly.

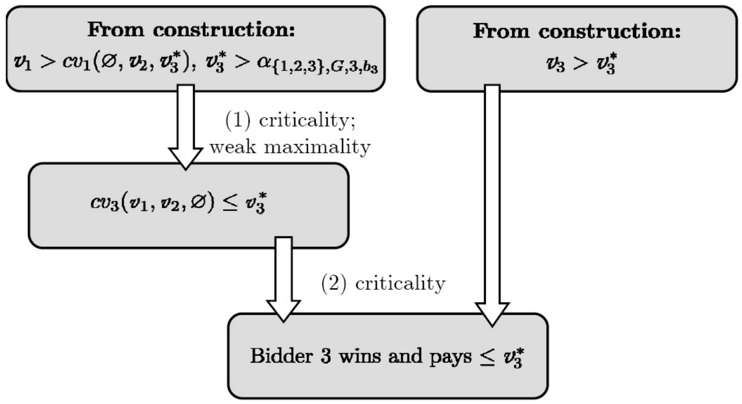

Figure 3: Illustration of the proof of Theorem 19: Part 1

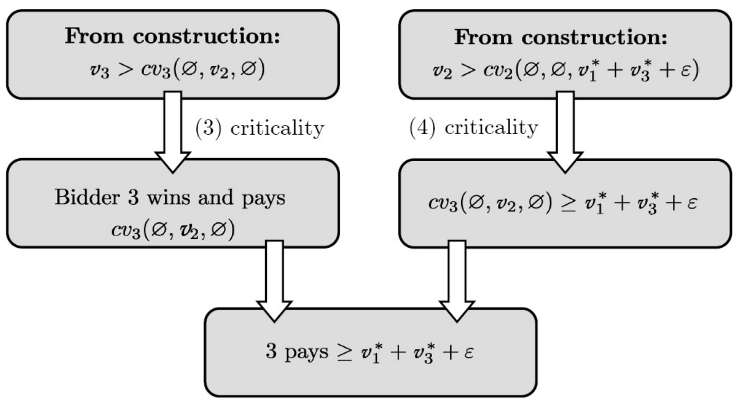

Figure 4: Illustration of the proof of Theorem 19: Part 2

$\alpha_{\{1,2,3\},G,3,g_2} + k$. Next pick an arbitrary positive constant ε, and then pick an arbitrary value for v_2 that satisfies

$$v_2 > cv_2(\varnothing, \varnothing, v_1^* + v_3^* + \varepsilon).$$

Finally, pick values for v_1 and v_3 that satisfy

$$v_1 > \max\{cv_1(\varnothing, v_2, v_3^*), cv_1(\varnothing, v_2, \varnothing), v_1^*\}, \text{ and}$$

$$v_3 > \max\{cv_3(v_1^*, v_2, \varnothing), cv_3(\varnothing, v_2, \varnothing), v_3^*\}.$$

By Corollary 17 the above critical values are all finite. Dependencies between v_1^*, v_3^*, v_2, v_1, and v_3 are shown in Figure 2, illustrating the fact that it is possible to pick values for these variables that satisfy all our constraints by following the ordering given.

The rest of the proof is divided into two parts. In Part 1 we consider $N = \{1, 2, 3\}$ and construct an expression for the auction's revenue. In Part 2 we consider $N = \{2, 3\}$ and show that more revenue is obtained than in Part 1. Sketches of the arguments in each of these parts are given in Figures 3 and 4 respectively.

Part 1: Since $v_1 > cv_1(\varnothing, v_2, v_3^*)$ (by construction), if bidder 3 were to bid v_3^* then bidder 1 would win (by criticality). By construction, bidder 3 is the only bidder whose bundle does not overlap with b_1 and $v_3^* > \alpha_{\{1,2,3\},G,3,b_3}$; thus, by weak maximality bidder 3 would also win and, by criticality,

$$cv_3(v_1, v_2, \varnothing) \leq v_3^* \text{ (see (1) in Figure 3)} \tag{1}$$

Symmetrically, from $v_3 > cv_3(v_1^*, v_2, \varnothing)$ we can also conclude that

$$cv_1(\varnothing, v_2, v_3) \leq v_1^* \tag{2}$$

By construction, $v_1 > v_1^*$ and $v_3 > v_3^*$. Then, using Inequalities (1) and (2) and by criticality, bidders 1 and 3 win (see (2) in Figure 3). By participation, since bidder 2 loses she must pay zero. Therefore the revenue of the auction, by criticality, is $R = cv_1(\varnothing, v_2, v_3) + cv_3(v_1, v_2, \varnothing) \leq v_1^* + v_3^*$.

Part 2: If bidder 1 is not present, then only bidders 2 and 3 compete. Since $v_3 > cv_3(\varnothing, v_2, \varnothing)$, by criticality, bidder 3 wins and pays $cv_3(\varnothing, v_2, \varnothing)$ (see (3) in Figure 4). Since b_2 and b_3 overlap, bidder 2 loses and by participation pays zero. The revenue of the auction is therefore $R_{-1} = cv_3(\varnothing, v_2, \varnothing)$. By construction, $v_2 > cv_2(\varnothing, \varnothing, v_1^* + v_3^* + \varepsilon)$. Thus if bidder 3 were to bid $v_1^* + v_3^* + \varepsilon$ then bidder 2 would win (by criticality) and so bidder 3 would lose. This tells us (again by criticality) that $cv_3(\varnothing, v_2, \varnothing) \geq v_1^* + v_3^* + \varepsilon$ (see (4) in Figure 4). Therefore, $R_{-1} = cv_3(\varnothing, v_2, \varnothing) \geq v_1^* + v_3^* + \varepsilon > v_1^* + v_3^* \geq R$. Thus, M is not revenue monotonic. □

Finally, as stated in the following theorem, the result holds for any mechanism for which $\mathcal{V}$, the set of valuation subspaces, *subsumes* known single-minded bidders valuation subspace.

Theorem 19. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism whose set of valuation subspaces $\mathcal{V}$ subsumes known single-minded bidders, that offers dominant strategies to the bidders and satisfies participation, consumer sovereignty, and weak maximality with respect to at least two bidders. Then M is not revenue monotonic.*

Proof The proof directly follows Lemma 14 and Theorem 18. Following Lemma 14, if there is a mechanism M defined for $\mathcal{V}$—i.e., M's set of valuation subspaces is $\mathcal{V}$—that offers dominant strategies to the bidders and satisfies participation, consumer sovereignty, weak maximality with respect to at least two bidders and revenue monotonicity, then there exists a mechanism $M^{(ksm)}$ defined for known single-minded bidders that has all the above properties. By Theorem 18, such a mechanism $M^{(ksm)}$ does not exist, and thus nor does such a mechanism M. □

One might have imagined that maximality would increase an auction mechanism's revenue by not "leaving money on the table," augmenting allocations to award available goods to the bidders who value them. Instead, we have shown above that any dominant-strategy truthful combinatorial auction mechanism that satisfies weak maximality with respect to at least two bidders—along with some other, very standard conditions—can sometimes collect no more than predefined constant amounts despite competition between bidders. Specifically, given the constructed valuations, bidder 2's losing bid has no effect on the prices paid by winning bidders 1 and 3, who also offer each other no competition as they bid on separate bundles. Thus bidders 1 and 3 each pay an amount arbitrarily close to a predefined constant. On the other hand, when bidder 1 is dropped then bidders 2 and 3 do compete. Although bidder 3 still wins, she pays more than before and, given the constructed valuations, more than the sum of the payments in the three-bidder case.

Observe that, given the constructed valuations, the mechanism can gain arbitrarily higher revenue in the two-bidder case than in the three-bidder case, since ε and k can be set to be arbitrarily large and arbitrarily small, respectively. In the three-bidder case the mechanism may generate almost the lowest possible revenue (the sum of the predefined constant amounts) as k can be set arbitrarily close to zero.

5. Related Impossibility Results

Our main result from Theorem 19 straightforwardly implies several other impossibility results. Here we demonstrate, considering the same family of mechanisms as before, that it is impossible to achieve monotonicity in the set of goods rather than the set of bidders, false-name-proofness, and outcomes belonging to the core.

5.1. Monotonicity in the set of goods

First, we show that we can also obtain the same impossibility results as in Theorem 19 when we define revenue monotonicity over the set of goods instead of over the set of bidders. This result may be more intuitive than our first result, as it relies on the fact that adding goods to an auction can reduce the level of competition between the bidders.

Introduce the notation $p_i^G(v)$ to denote bidder i's payment to a truthful CA mechanism when all bidders' valuations are v, and where the set of goods at auction is G. Then we can give the following definition.

Definition 20 (Good revenue monotonicity). *A truthful CA mechanism M is* good revenue monotonic *if and only if for all $N \subseteq \mathbb{N}, G \subseteq \mathbb{G}$, $V_{N,G} \in \mathcal{V}$, $v \in V_{N,G}$ and for all goods $g \in G$,*

$$\sum_{i \in \mathbb{N}} p_i^G(v) \geq \sum_{i \in \mathbb{N}} p_i^{G \setminus \{g\}}(v).$$

Corollary 21. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism whose set of of valuation subspaces subsumes known single-minded bidders, that offers dominant strategies to the bidders and satisfies participation, consumer sovereignty, and weak maximality with respect to at least two bidders. Then M is not good revenue monotonic.*

Proof The claim follows directly from the proof of Theorem 18 and Theorem 19 with the following modifications: (i) add an extra good g_3 to bidder 1's bundle b_1, and (ii) instead of dropping bidder 1 in Part 2, drop g_3—then bidder 1's valuation for all available bundles will be 0. □

5.2. False-name-proofness

False-name (pseudonymous) bidding has been studied extensively (e.g., Yokoo (2006) and Yokoo et al. (2001, 2004)). This work is concerned with auctions in which a bidder may submit multiple bids using pseudonyms. An auction mechanism is said to be false-name-proof if truth-telling without using false-name bids is a dominant strategy for each bidder. Yokoo et al. (2001) proved that there does not exist any combinatorial auction mechanism that is false-name-proof and efficient. Observe that this is a somewhat narrow result, because—as discussed earlier—only Groves mechanisms are both dominant-strategy truthful and efficient (Green and Laffont, 1977).

There is a connection between false-name-proofness and revenue monotonicity. From the seller's perspective, false-name bidding is the same as having more bidders in the auction. If an auction is not revenue monotonic, more bidders can mean less revenue. Our results are therefore relevant to research on false-name bidding. For technical reasons, we have to make minor changes to our formal model to capture false-name bidding (e.g., we have assumed that mechanisms know bidders' identities.) We can then prove the following corollary which generalizes the result of Yokoo et al. (2001) by replacing their requirement of efficiency with the much weaker criterion of weak maximality. Recall that all efficient mechanisms are maximal and therefore weakly maximal with respect to all bidders, but there exist other mechanisms that are inefficient and still maximal.

Corollary 22. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism whose set of of valuation subspaces subsumes known single-minded bidders, that offers dominant strategies to the bidders, and that satisfies participation, consumer sovereignty, and weak maximality with respect to at least two bidders. Then M is not false-name-proof.*

Proof Given the valuations constructed in the proof of Theorem 18, bidder 3 gains by pseudonymously bidding also as bidder 1, and so truth telling is not a dominant strategy for bidder 3. □

5.3. Outcomes in the Core

It is relatively standard (see, e.g., Ausubel and Milgrom (2002) and Day and Milgrom (2007)) to describe efficient auction mechanisms as coalitional games. Coalitional game theory focuses on groups of players and the utility they can achieve together. Thus, this theory can be useful for discussing what happens to an auction's revenue when bidders are added or removed.

While the application of coalitional game theory to modeling efficient auction mechanisms is unproblematic (see Section 5.3.1), it is less clear whether it is appropriate to model inefficient mechanisms as coalitional games, and if so how to do so. We discuss this concern in Section 5.3.2. We then present several alternate formulations, and for each consider whether a combinatorial auction can be guaranteed to select outcomes that belong to the core.

5.3.1. Modeling Efficient Mechanisms as Coalitional Games

A *transferable utility (TU)* coalitional game is defined by a set of *players* N_p and a *characteristic function* w that maps each coalition of players S to the coalition's value, $w(S)$. The grand coalition is the coalition of all players. An imputation is a payoff profile in which each player receives a non-negative payoff and the sum of the payoffs does not exceed the grand coalition's value. An efficient combinatorial auction naturally defines a TU coalitional game. Let N_p be the set of participating bidders, N, plus the seller whom we denote by 0. An efficient auction game is then defined as follows.

Definition 23 (Efficient auction game). *For any coalition $S \subseteq N_p$, define the coalition's value as*

$$w(S) = \begin{cases} \max_{\mathrm{a} \in \mathbb{A}_{S,G}} \sum_{i \in S} v_i(\mathrm{a}_i) & 0 \in S, \\ 0 & 0 \notin S. \end{cases}$$

Intuitively, in an efficient auction game, the value of a coalition consisting of any set of players S including the seller is the maximum social welfare achievable under S. When the seller does not belong to a coalition, the coalition's value is zero.

In an auction, the mechanism picks a specific imputation by imposing the chosen allocation and payments. We call this *this auction's imputation*. In an auction game, define the payoff of the seller under the auction's imputation as the auction's revenue, $\pi_0 = R = \sum_{i \in N} p_i$. Define bidder i's payoff as her utility from the auction, $\pi_i = u_i = v_i - p_i$. Observe that in an efficient auction game $\sum_{i \in N_p} \pi_i = w(N_p)$.

Definition 24 (Core in TU coalitional game). *An imputation π is in the* core *of a TU coalitional game if and only if no subset of players can achieve higher payoff:*

$$\forall S \subseteq N_p, \sum_{j \in S} \pi_j \geq w(S).$$

If an auction's imputation is in the core, no coalition has an incentive to deviate from it. Note that we consider the possibility that the grand coalition (in addition to smaller coalitions) would make such a deviation. We say that the outcome of an auction mechanism is in the core if the auction's imputation is in the core. Note that any efficient mechanism is maximal with respect to all bidders. Our impossibility result then implies the following corollary.

Corollary 25. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism whose set of of valuation subspaces subsumes known single-minded bidders, that offers dominant strategies to the bidders, and that satisfies participation, consumer sovereignty, and efficiency. Then, there exists a valuation profile for which the auction's imputation does not belong to the core.*

This result follows as a special case of Corollary 28, so we omit the proof.

5.3.2. Modeling Inefficient Mechanisms as Coalitional Games

The literature on modeling auctions as coalitional games focuses on efficient mechanisms. This makes sense under the assumption that any deviating coalition can achieve a social welfare maximizing outcome. Recall that in an auction game, the payoff of the seller is the auction's revenue and the payoff of each bidder is her utility. If one attempts to describe an inefficient auction mechanism as a TU game following Definition 23, the outcome of the auction is not guaranteed to be in the core. This is because the sum of the payoffs may not add up to the grand coalition's value. In other words, if the auction mechanism chooses an inefficient outcome then the grand coalition has an incentive to deviate to an efficient outcome. However, if a seller elects to use an inefficient mechanism, it is inconsistent to then imagine all bidders and the seller jointly deviating to an efficient allocation. The use of an inefficient mechanism can nevertheless make sense, e.g., because regulatory or computational constraints may limit the set of outcomes that can be achieved. Therefore, here we aim to model inefficient mechanisms as coalitional games. Specifically, we discuss three alternate coalitional game models of the auction game, none of which obviously dominates the others.

In the first alternative, which makes minimal changes to Definition 23, we assume that players can reach the efficient allocation for all but the grand coalition. (That is, we assume that the coalition's value for all coalitions except the grand coalition is as stated in Definition 23.)

Definition 26 (Inefficient auction game (first alternative)). *For any coalition $S \subseteq N_p$, define the coalition's value as*

$$w(S) = \begin{cases} \sum_{i \in S} v_i(a_i(v)) & S = N_p, \\ \max_{\mathrm{a} \in \mathbb{A}_{S,G}} \sum_{i \in S} v_i(\mathrm{a}_i) & 0 \in S \text{ and } S \neq N_p, \\ 0 & 0 \notin S. \end{cases}$$

In the second alternative, we assume that players have to obey the mechanism's allocation choice under all coalitions, rather than only under the grand coalition.

Definition 27 (Inefficient auction game (second alternative)). *For any coalition $S \subseteq N_p$, define the coalition's value as*

$$w(S) = \begin{cases} \sum_{i \in S} v_i(a_i(v)) & 0 \in S, \\ 0 & 0 \notin S. \end{cases}$$

The second alternative may seem more plausible than the first one. We do not need to choose between them, however, as both lead to the following impossibility result.

Corollary 28. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism whose set of of valuation subspaces $\mathcal{V}$ subsumes known single-minded bidders, that offers dominant strategies to the bidders and satisfies participation, consumer sovereignty, and weak maximality with respect to at least two bidders. Define the auction game as in Definition 26 or Definition 27. Then, there exists a valuation profile for which the auction's imputation does not belong to the core.*

Proof The proof can be derived from the proof of Theorem 18, by slight modifications, and Lemma 14, by making a similar argument as in the proof of Theorem 19. First, construct valuations as in the proof of Theorem 18, but now choose v_2 to satisfy the constraint $v_2 > \max(cv_2(\varnothing, \varnothing, v_1^* + v_3^* + \varepsilon), cv_2(\varnothing, \varnothing, \varnothing), v_1^* + v_3^*)$. Then, notice that in the auction game defined as in either of Definitions 26 or 27, the coalition of the seller and bidder 2 has an incentive to deviate from the grand coalition since $w(\{2, 0\}) = v_2 > v_1^* + v_3^* \geq u_2 + R$. In other words, the seller can sell the bundle to bidder 2 for the price of p_2', $v_1^* + v_3^* < p_2' < v_2$, making both herself and bidder 2 better off. □

In the coalitional game formulations that we have considered so far, the mechanism only dictates its choice of allocation to some or all of the coalitions—specifically, to the grand coalition in Definition 26 and to all coalitions in Definition 27. We may want to assume that the mechanism imposes not only its choice of allocation, but also its choice of payments. This motivates our third coalitional game model, which describes an inefficient auction game as a coalitional game with *nontransferable utility (NTU)*.

Formally, a NTU coalitional game is defined by a set of players N_p and a characteristic function w that maps each coalition of players S to a set of real-valued vectors describing different sets of payoffs achievable by the members.

Definition 29 (Core in an NTU coalitional game). *A payoff vector $\pi \in w(N_p)$ is in the* core *of a NTU coalitional game if and only if $\forall S \subseteq N_p$, $\neg\exists x \in w(S)$ such that $\forall i \in S, \pi_i \leq x_i$ and $\exists j \in S, \pi_j < x_j$.*

Definition 30 (Inefficient auction game (third alternative)). *Let the characteristic function w map each coalition $S \subseteq N_p$ to a single real-valued vector in which each player's payoff is exactly her utility under the mechanism's chosen allocation and taking into account her payment to the mechanism, when the set of participating bidders is $S \setminus \{0\}$.*

For known single-minded bidders, all mechanisms that involve only a single bidder i, that satisfy participation, and that offer dominant strategies can be understood as offering i her desired bundle at a fixed price, $cv_i(\varnothing, \ldots, \varnothing)$. The following result can be understood as showing that any mechanism satisfying our conditions either already sets $cv_i(\varnothing, \ldots, \varnothing)$ in such a way that both the seller and i can gain when all other bidders are excluded from the mechanism, or can be modified to do so. Intuitively, our counterexample cannot be used to show that a given (unmodified) mechanism *always* suffers from this problem because, while i is always better off when the other bidders are dropped, the seller could be worse off if $cv_i(\varnothing, \ldots, \varnothing)$ is set too low.

Corollary 31. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism whose set of valuation subspaces $\mathcal{V}$ subsumes known single-minded bidders, that offers dominant strategies to the bidders and satisfies participation, consumer sovereignty, and weak maximality with respect to at least two bidders. Then there exists a CA mechanism M' whose set of valuation subspaces is $\mathcal{V}$, and that*

1. *has the same allocation and payment functions as M, except that it may have a different $cv_i(\varnothing, \ldots, \varnothing)$ for some (single) $i \in \mathbb{N}$;*

2. *satisfies participation, consumer sovereignty, and weak maximality with respect to at least two bidders; and*
3. *chooses an outcome that is not guaranteed to belong to the core.*

Proof The result follows from the proof of Theorem 18, by slight modification, and from Lemma 14, by making a similar argument as in the proof of Theorem 19. Consider the three-bidder two-good setting in the proof of Theorem 18. To emphasize that M may already choose outcomes that do not belong to the core, our proof considers two cases.

Case 1: $cv_2(\varnothing, \varnothing, \varnothing) > \alpha_{\{1,2,3\},G,1,g_1} + \alpha_{\{1,2,3\},G,3,g_3}$. Pick an arbitrary positive $k < \frac{1}{2}(cv_2(\varnothing, \varnothing, \varnothing) - \alpha_{\{1,2,3\},G,1,g_1} - \alpha_{\{1,2,3\},G,3,g_3})$. Construct valuations as in the proof of Theorem 18, given the chosen k, but now choose v_2 to satisfy the constraint $v_2 > \max(cv_2(\varnothing, \varnothing, v_1^* + v_3^* + \varepsilon), cv_2(\varnothing, \varnothing, \varnothing), v_1^* + v_3^*)$. By Corollary 17, M's revenue when only bidder 2 participates is $R_2 = cv_2(\varnothing, \varnothing, \varnothing) > v_1^* + v_3^* \geq R$. The utility of bidder 2 in this case—i.e., when only bidder 2 participates—is $u_2 = v_2 - cv_2(\varnothing, \varnothing, \varnothing) > 0$, which is strictly greater than bidder 2's utility when all three bidders participate. Thus, the outcome chosen by M does not belong to the core; let $M' = M$.

Case 2: $cv_2(\varnothing, \varnothing, \varnothing) \leq \alpha_{\{1,2,3\},G,1,g_1} + \alpha_{\{1,2,3\},G,3,g_3}$. Construct M' to be the same as M, except choose $cv_2(\varnothing, \varnothing, \varnothing) > \alpha_{\{1,2,3\},G,1,g_1} + \alpha_{\{1,2,3\},G,3,g_3}$. Observe that this change preserves dominant strategies (this property is unaffected by the specific value taken by $cv_2(\varnothing, \varnothing, \varnothing)$), participation (bidder 2 pays nothing if she loses), consumer sovereignty ($cv_2(\varnothing, \varnothing, \varnothing)$ is finite), and weak maximality with respect to bidders 1 and 3 (nothing changes for these bidders). Then, proof follows from the argument in Case 1. □

Earlier, when we modeled inefficient auctions as TU games, we assumed that bidder 2 and the seller could divide gains between them, meaning that the pair were always better off forming a coalition. Under the NTU model, that division must be described explicitly through the auction's payment rule. The proof of Theorem 31 shows that such a division can always be accomplished by an appropriate choice of $cv_i(\varnothing, \dots, \varnothing)$.

6. Conclusions and Future Work

In this work, we investigated whether there exists any dominant-strategy truthful CA mechanism that satisfies participation, consumer sovereignty and weak maximality with respect to at least two bidders and is revenue monotonic. We showed that no such mechanism exists; as corollaries, we were able to show similar results concerning mechanisms that yield weakly decreasing revenue when goods are dropped and false-name-proof mechanisms. Also, we investigated the relationship between a mechanism being revenue monotonic and the mechanism yielding an outcome that belongs to the core. More specifically, we showed that for any mechanism that satisfies our desired properties, the outcome of the mechanism is not guaranteed to belong to the core.

In future work, we are interested in investigating the probability that such revenue monotonicity failures occur in practical auctions. In a similar vein, it is also interesting to ask what dominant-strategy truthful CA mechanism has all the properties we demanded before and has the minimum probability of violating revenue monotonicity.

References

Andelman, N., Mansour, Y., 2006. A sufficient condition for truthfulness with single parameter agents. In: Proceedings of the ACM Conference on Electronic Commerce (EC). pp. 8–17.

Archer, A., Tardos, E., 2001. Truthful mechanisms for one-parameter agents. In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS). pp. 482–491.

Ausubel, L. M., 2006. An efficient dynamic auction for heterogeneous commodities. American Economic Review 96, 602–629.

Ausubel, L. M., Cramton, P., Milgrom, P., 2006a. The clock-proxy auction: A practical combinatorial auction design. In: (Cramton et al., 2006), Ch. 5.

Ausubel, L. M., Cramton, P., Milgrom, P., 2006b. The clock-proxy auction: A practical combinatorial auction design. In: (Cramton et al., 2006), Ch. 5, pp. 115–138.

Ausubel, L. M., Milgrom, P., 2002. Ascending auctions with package bidding. Frontiers of Theoretical Economics 1.

Ausubel, L. M., Milgrom, P., 2006. The lovely but lonely Vickrey auction. In: (Cramton et al., 2006), Ch. 1, pp. 17–40.

Babaioff, M., Lavi, R., Pavlov, E., 2005. Mechanism design for single-value domains. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 241–247.

Babaioff, M., Lavi, R., Pavlov, E., 2006. Single-value combinatorial auctions and implementation in undominated strategies. In: Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SOCA). pp. 1054–1063.

Balcan, M.-F., Blum, A., 2006. Approximation algorithms and online mechanisms for item pricing. In: Proceedings of the ACM Conference on Electronic Commerce (EC). pp. 29–35.

Bartal, Y., Gonen, R., Nisan, N., 2003. Incentive compatible multi-unit combinatorial auctions. In: Proceedings of the ACM Conference on Theoretical Aspects of Rationality and Knowledge (TARK). pp. 72–87.

Bernheim, D. B., Whinston, M., 1986. Menu auctions, resource allocation and economic influence. Quarterly Journal of Economics 101, 1–31.

Bikhchandani, S., Chatterji, S., Lavi, R., Mu'alem, A., Nisan, N., Sen, A., 2006. Weak monotonicity characterizes deterministic dominant-strategy implementation. Econometrica 74 (4), 1109–1132.

Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R. V., 2001. Linear programming and Vickrey auctions. IMA Volumes in Mathematics and its Applications 127, 75–116.

Clarke, E. H., 1971. Multipart pricing of public goods. Public Choice 8, 17–33.

Constantin, F., Parkes, D. C., 2005. Preference-based characterizations of truthfulness and the limited expressiveness of order-based domains. Proceeding of the Workshop on Preference HandlingPosition Paper.

Cramton, P., Shoham, Y., Steinberg, R. (Eds.), 2006. Combinatorial Auctions. MIT Press, Cambridge, Massachusetts.

Day, R., Milgrom, P., 2007. Core-selecting package auctions. International Journal of Game Theory, 393–407.

Demange, G., Gale, D., Sotomayor, M., 1986. Multi-item auctions. Journal of Political Economy 94, 863–872.

Dobzinski, S., Schapira, M., 2006. An improved approximation algorithm for combinatorial auctions with submodular bidders. In: Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SOCA). pp. 1064–1073.

Englmaier, F., Guillen, P., Llorente, L., Onderstal, S., Sausgruber, R., 2006. The chopstick auction: a study of the exposure problem in multi-unit auctions. CESIFO Working Paper 1782.

Feigenbaum, J., Krishnamurthy, A., Sami, R., Shenkar, S., 2002. Hardness results for multicast cost sharing. Foundations of Software Technology and Theoretical Computer Science 2556, 133–144.

Goldberg, A. V., Hartline, J. D., 2003. Envy-free auctions for digital goods. In: Proceedings of the ACM Conference on Electronic Commerce (EC). pp. 29–35.

Green, J., Laffont, J.-J., 1977. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. Econometrica 45 (2), 427–438.

Groves, T., 1973. Incentives in teams. Econometrica 41 (4), 617–631.

Gul, F., Stacchetti, E., 1999. Walrasian equilibrium with gross substitutes. Journal of Economics Theory 87, 95–124.

Gul, F., Stacchetti, E., 2000. The English auction with differentiated commodities. Journal of Economics Theory 92, 66–95.

Haris, M., Raviv, A., 1981. Allocation mechanisms and the design of auctions. Econometrica 49 (6), 1477–1499.

Hartline, J. D., Roughgarden, T., 2008. Optimal mechanism design and money burning. In: Proceedings of the Annual ACM Symposium on Theory of Computing (STOC). pp. 75–84.

Håstad, J., 1999. Clique is hard to approximate to within $n^{1-\epsilon}$. Acta Mathematica 182, 627–636.

Kelly, F., Steinberg, R., 2000. A combinatorial auction with multiple winners for universal service. Management Science 46, 586–596.

Khot, S., Lipton, R. J., Markakis, E., Mehta, A., 2005. Inapproximability results for combinatorial auctions with submodular utility functions. In: Proceedings of the International Workshop on Internet and Network Economics (WINE). pp. 92–101.

Krishna, V., Perry, M., 1998. Efficient mechanism design. Tech. rep., Pennsylvania State University.

Lavi, R., Mu'alem, A., Nisan, N., 2003. Towards a characterization of truthful combinatorial auctions. In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS). pp. 574–583.

Ledyard, J. O., Porter, D., Rangel, A., 1997. Experiments testing multi-object allocation mechanisms. Journal of Economics and Management Strategy 6 (3), 639–675.

Lehmann, B., Lehmann, D., Nisan, N., 2001. Combinatorial auctions with decreasing marginal utilities. In: Proceedings of the ACM Conference on Electronic Commerce (EC). pp. 18–28.

Lehmann, D., O'Callaghan, L. I., Shoham, Y., 2002. Truth revelation in approximately efficient combinatorial auctions. Journal of the ACM 49 (5), 1–26.

Likhodedov, A., Sandholm, T., 2005a. Approximating revenue-maximizing combinatorial auctions. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 267–274.

Likhodedov, A., Sandholm, T., 2005b. Methods for boosting revenue in combinatorial auctions. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 232–237.

Maskin, E. S., Riley, J. G., 1980. Auctioning an indivisible object. Tech. rep., Kennedy School, Harvard University.

Milgrom, P., 2002. Putting auction theory to work: the simultaneous ascending auction. Journal of Political Economy 018, 245–272.

Monderer, D., Tennenholtz, M., 2005. Asymptotically optimal multi-object auctions, http://ie.technion.ac.il/ dov/cauc28.pdf.

Mu'alem, A., Nisan, N., 2002. Truthful approximation mechanisms for restricted combinatorial auctions: Extended abstract. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 379–384.

Myerson, R. B., 1981. Optimal auction design. Mathematics of Operations Research 6 (1), 58–73.

Nisan, N., 2007. Introduction to mechanism design. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. V. (Eds.), Algorithmic Game Theory. Cambridge University Press, Cambridge, UK, Ch. 9.

Nisan, N., Ronen, A., 2000. Computationally feasible VCG mechanisms. In: Proceedings of the ACM Conference on Electronic Commerce (EC). pp. 242–252.

Nisan, N., Segal, I., 2006. The communication requirements of efficient allocations and supporting prices. Journal of Economic Theory 129, 192–224.

Parkes, D., Ungar, L., 2000. Iterative combinatorial auctions: Theory and practice. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 74–81.

Porter, D., Rassenti, S., Roopnarine, A., Smith, V., 2003. Combinatorial auction design. Proceedings of the National Academy of Sciences of the United States of America, PNAS 100, 11153–11157.

Rassenti, S., Smith, V. L., Bulfin, R. L., 1982. A combinatorial auction mechanism for airport time slot allocation. The Bell Journal of Economics 13, 402–417.

Rastegari, B., Condon, A., Leyton-Brown, K., 2007. Revenue monotonicity in combinatorial auctions. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 122–127.

Riley, J. G., Samuelson, W. F., 1981. Optimal auctions. American Economic Review 71 (3), 381–392.

Roberts, K., 1979. The characterization of implementable choice rules. Aggregation and Revelation of Preferences, 321–348.

Rochet, J.-C., 1987. A necessary and sufficient condition for rationalizability in a quasi-linear context. Journal of Mathematical Economics 16, 191–200.

Ronen, A., 2001. On approximating optimal auctions. In: Proceedings of the ACM Conference on Electronic Commerce (EC). pp. 11–17.

Rothkopf, M. H., 2007. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. Operations Research 55 (2), 191–197.

Saks, M., Yu, L., 2005. Weak monotonicity suffices for truthfulness on convex domains. In: Proceedings of the ACM Conference on Electronic Commerce (EC). pp. 286–293.

Sandholm, T., 2002. An algorithm for optimal winner determination in combinatorial auctions. Artificial Intelligence 135, 1–54.

Segal, I., 2006. The communication requirements of combinatorial allocation problems. In: (Cramton et al., 2006), Ch. 11.

Vickrey, W., 1961. Counterspeculations, auctions, and competitive sealed tenders. Journal of Finance 16, 8–37.

Yokoo, M., 2006. Pseudonymous bidding in combinatorial auctions. In: (Cramton et al., 2006), Ch. 7, pp. 161–187.

Yokoo, M., Sakurai, Y., Matsubara, S., 2001. Robust combinatorial auction protocol against false-name bids. Artificial Intelligence 130, 167–181.

Yokoo, M., Sakurai, Y., Matsubara, S., 2004. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. Games and Economic Behavior 46 (1), 174–188.