False-Name Bidding and Economic Efficiency in Combinatorial Auctions

Colleen Alkalay-Houlihan and Adrian Vetta
School of Computer Science, McGill University

Abstract

Combinatorial auctions are multiple-item auctions in which bidders may place bids on any package (subset) of goods. This additional expressibility produces benefits that have led to combinatorial auctions becoming extremely important both in practice and in theory. In the computer science community, auction design has focused primarily on computational practicality and incentive compatibility. The latter concerns mechanisms that are resistant to bidders misrepresenting themselves via a single false identity; however, with modern forms of bid submission, such as electronic bidding, other types of cheating have become feasible. Prominent amongst them is *false-name bidding*; that is, bidding under pseudonyms. For example, the ubiquitous Vickrey-Clarke-Groves (VCG) mechanism is incentive compatible and produces optimal allocations, but it is not false-name-proof – bidders can increase their utility by submitting bids under multiple identifiers. Ergo, there has recently been much interest in the design and analysis of false-name-proof auction mechanisms. These false-name-proof mechanisms, however, have polynomially small efficiency guarantees: they can produce allocations with very low economic efficiency/social welfare. In contrast, we show that, provided the degree to which different goods are complementary is bounded (as is the case in many important, practical auctions), the VCG mechanism gives a constant efficiency guarantee. Constant efficiency guarantees hold even at equilibria where the agents bid in a manner that is not individually rational. Thus, while an individual bidder may personally benefit greatly from making false-name bids, this will have only a small detrimental effect on the objective of the auctioneer: maximizing economic efficiency. So, from the auctioneer's viewpoint the VCG mechanism remains preferable to false-name-proof mechanisms.

Introduction

In combinatorial auctions, bidders express a valuation for every possible subset of the entire set of goods. Combinatorial auctions allow bidders to express complementary and substitute preferences over the goods through their valuations of different packages. This can result in greater economic efficiency. The practical applications of combinatorial auctions are important and varied. They include industrial procurement, truck scheduling, and the allocation of bandwidth for wireless communication. See (Cramton, Shoham, and Steinberg 2006) for an in-depth guide to combinatorial auctions. Formally, we have a set $G = \{x_1, x_2, ..., x_m\}$ of goods and a collection $I = [n] = \{1, 2, ..., n\}$ of bidders. For each package of goods $S \subseteq G$, bidder $i \in I$ has a non-negative value $v_i(S)$. A *feasible allocation* of the goods to these bidders is a collection of pairwise-disjoint subsets, that is, an assignment $\mathcal{T} = \{T_1, T_2, \dots, T_n\}$ such that (i) $T_i \subseteq G, \forall i \in I$ and (ii) $T_i \cap T_j = \emptyset, \forall i \neq j$. In the standard sealed-bid auction, each bidder i submits a bid vector $\mathbf{b_i}$ consisting of a bid $b_i(S)$ for each package S; the auctioneer then uses these bids to determine a feasible allocation of the goods and the price each bidder must pay (this price may differ from the bid value).

The VCG Mechanism

For governmental auctions (such as bandwidth auctions), rather than maximize revenue, the stated objective is typically to maximize social welfare (economic efficiency); the social welfare of a feasible allocation $\mathcal{T}$ is $\omega(\mathcal{T}) = \sum_{i \in I} v_i(T_i)$. Clearly, this objective is achievable only if the bidders bid truthfully; that is, they declare $\mathbf{b_i} = \mathbf{v_i}$. Fortunately, there is a classical auction mechanism that ensures this. The Vickrey-Clarke-Groves (VCG) mechanism outputs the feasible allocation that maximizes welfare according to the declared valuations $\mathbf{b_i}$. Bidder i is then charged a price p_i that equates her declared utility[1] with her marginal contribution to social welfare; equivalently, p_i is the total loss of utility (welfare) incurred by the other players as a result of her participation. Under this mechanism, truthful reporting is a dominant strategy for each bidder and so social welfare can be maximized. Thus, VCG is a *truthful mechanism*.

The focus of this paper is such incentives for truthfulness. The practical problems of running sealed-bid VCG mechanisms, such as computational complexity, are not our concern here; see (Rothkopf 2007) for a discussion of such issues. We remark, however, that most of these problems have been addressed in the most modern combinatorial auctions.[2]

[1] Assuming quasi-linear utility functions, where the utility of a package equals its value minus its price.

[2] For example, currently the most important application of combinatorial auctions is bandwidth auctions. There, the state-of-the-art *combinatorial clock auction* is computational very fast on practical instances. This auction actually consists of two phases. The

False-Name-Proof Mechanisms

So VCG is a truthful mechanism, but there is a caveat. VCG is resistant against a bidder submitting a single false identity, that is, one false vector of bids instead of a truthful vector. What if a bidder assumes multiple false identities and places numerous bid vectors under false names? We define a **false-name bid** as a bid made under a pseudonym (Yokoo, Sakurai, and Matsubara 2004) In this case, it turns out that VCG mechanisms are no longer truthful (Sakurai, Yokoo, and Matsubara 1999). To illustrate this, we present two simple examples where bidders can increase their utility by submitting false-name bids under multiple identifiers. In the first example, false-name bidding allows the bidder to increase her utility by winning a better package. In the second example, false-name bidding allows the bidder to increase her utility by decreasing the price she pays for her package.

Example 1: Consider an auction of two broadband licenses:

Bidder	*License 1*	*License 2*	*Licenses 1 and 2*
Dodgers	$1bn	$1bn	$8bn
Horizon	$4bn	$4bn	$4bn

Under the VCG mechanism, Dodgers is assigned both licenses and pays $4 billion. But suppose Horizon uses two false names, Horizon-1 and Horizon-2, and bids as follows:

Bidder	*License 1*	*License 2*	*Licenses 1 and 2*
Dodgers	$1bn	$1bn	$8bn
Horizon-1	$8bn	$0	$8bn
Horizon-2	$0	$8bn	$8bn

The VCG mechanism will now allocate License 1 to Horizon-1 and License 2 to Horizon-2, while Dodgers receives no license at all. Horizon-1 and Horizon-2 both pay $1 billion. By using the false names Horizon-1 and Horizon-2, Horizon is allocated both licenses and pays $2 billion. It thus clearly benefits from having used false-name bids, since its utility is now $2 billion, as opposed to zero when it bid truthfully with a single identifier.

Example 2: Now consider two art galleries competing in an auction for two "Group of Seven" paintings.

Bidder	*Mount Lefroy*	*The Red Maple*	Both
Hermitage	$7m	$7m	$14m
McMichael	$4m	$4m	$12m

In this case, the Hermitage is allocated both Mount Lefroy and The Red Maple for the VCG price of $12 million. Now suppose the Hermitage uses two false identifiers, Hermitage-1 and Hermitage-2.

Bidder	*Mount Lefroy*	*The Red Maple*	Both
Hermitage-1	$7m	$0	$7m
Hermitage-2	$0	$7m	$7m
McMichael	$4m	$4m	$12m

The VCG mechanism will now allocate Mount Lefroy to Hermitage-1 and The Red Maple to Hermitage-2. It can then be seen that the VCG prices are $5 million each for Hermitage-1 and Hermitage-2. Thus, the Hermitage, which is using two false identifiers, receives both paintings for $10m, improving its utility by $2m.

These examples are particularly troubling, as false-name bidding is easier to implement and sustain than other forms of cheating, such as collusion, that require the participation of multiple bidders. Furthermore, for many modern auctions (e.g. Internet auctions), false-name bidding is hard to prohibit as it is difficult to verify all the participants' identities.

To rectify this problem, (Yokoo, Sakurai, and Matsubara 2001b) initiated the development of **false-name-proof (FNP) mechanisms**, where it is a dominant strategy for each bidder to declare her true valuation function using a single identifier (even if it is possible to use multiple identifiers). Such mechanisms have recently received much attention in combinatorial auctions (Yokoo, Sakurai, and Matsubara 2004; Todo et al. 2009; Yokoo, Sakurai, and Matsubara 2001a). False-name bidding has also been examined in the context of online auction mechanisms (Todo et al. 2012), auctions that allow for bid withdrawals (Guo and Conitzer 2010), double auction protocols (Yokoo, Sakurai, and Matsubara 2001c; Sakurai and Yokoo 2003), and combinatorial multi-attribute procurement auctions (Suyama and Yokoo 2004). (Conitzer 2007) also used the idea of confirming the identities of a subset of participants through an external process in order to induce false-name-proofness in the entire mechanism. We emphasize that the work of (Sakurai, Yokoo, and Matsubara 1999) shows that there is no single-round, sealed-bid auction mechanism that is false-name-proof *and* satisfies individual rationality, Pareto efficiency, and incentive compatibility. False-name-proof mechanisms have also been studied beyond the domain of auctions. In particular, they have been examined in voting games (Bachrach and Elkind 2008; Wagman and Conitzer 2008; Aziz and Paterson 2009), social choice mechanisms (Todo 2010), and social networks (Conitzer et al. 2010).

Submodularity: Complements, Substitutes and False-Name-Proof Mechanisms

In combinatorial auctions, bidders can bid for the packages they want at the prices they want. This means they do not suffer from two flaws inherent in alternate auction platforms: the exposure problem, and the specification problem. The *exposure problem* is where a bid for a package Q leaves the bidder liable to win a package $P \subset Q$. The *specification problem* arises when separate bids for packages P and Q leave the bidder liable to win a package $P \cup Q$. Such liabilities are very undesirable when the goods are complements (for the exposure problem) or substitutes (for the specification problem). We say that a collection of goods are **substitutes** if the demand for one is non-decreasing in the price of the others; that is, if the price of one such good increases, the demand for its substitutes must either increase or stay the same. Goods are **complements** if the demand for one is non-increasing in the price of the others. A simple example of a pair of substitute goods is two different car models, whereas cars and gasoline form a pair of complementary goods.

There is a close relationship between substitutability and submodular functions. A set function $f : 2^X \to \mathbb{R}$ is **submodular** if and only if for all $A \subseteq B \subseteq X$ and all $x \in X \setminus B$: $f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B)$. If goods are substitutes for a bidder $i \in I$ then bidder i's indi-

final phase is a sealed-bid VCG auction (or a close relative, such as the nearest-Vickrey, core-selecting auction). This auction is as described above, except that restrictions are imposed upon what combinatorial bids are allowed. These restrictions are derived from information obtained in the first "clock" phase, a learning phase, that incorporates bidding constraints to incentivize truth-telling.

rect utility function is submodular. More specifically, goods are substitutes for bidder $i \in I$ if and only if bidder i's indirect utility function is submodular; see (Cramton, Shoham, and Steinberg 2006) for details.

Furthermore, substitutability, and thus submodularity, has a close relationship with false-name-proof mechanisms. For example, the VCG mechanism is false-name-proof if goods are substitutes. Thus, in this case, we have a false-name-proof mechanism that produces optimal welfare. For general valuation functions, however, false-name-proof mechanisms can produce extremely poor welfare. Specifically, (Iwasaki et al. 2010) showed (assuming Independence of Irrelevant Goods) that for any false-name-proof mechanism, there are instances for which it gives social welfare of at most $\frac{2}{m+1} \cdot \text{OPT}$. Here m is the number of goods and OPT is the maximum social welfare over all feasible allocations.

The very negative welfare result of (Iwasaki et al. 2010) provides motivation for our work. Truthful bidding is only a means to an end. The auctioneer desires truthful bidding as it should allow it optimise its objective – in this case economic efficiency. Thus, if the incentives provided by a false-name-proof mechanism to ensure truthfulness themselves negatively impact this objective, then that mechanism will have little appeal to the auctioneer. So we ask if it is possible to design a mechanism that will achieve high economic efficiency even if the bidders can manipulate the mechanism by making false-name bids? The answer is yes and we quantify the extent to which the VCG mechanism has this property.

Our Results

Rather than a polynomially small welfare guarantee, we show that the VCG mechanism will provide a constant factor welfare guarantee even when a bidder uses pseudonyms, provided the degree of complementarity between the items is not particularly large. This characteristic is common in practice (Cramton, Shoham, and Steinberg 2006). In fact, in many cases, such as bandwidth auctions, marginal valuations do not increase much with set size; thus, if complementarities exist, "they are not too large and may be bounded in some way" (Lehmann, Lehmann, and Nisan 2006).

Thus we focus upon auctions where bidders' valuation functions are "near-submodular" to a certain specified degree. A natural measure of this *degree of submodularity* is:

$$\mathcal{D}(f) = \min_{x} \min_{A,B:A\subset B} \frac{f(A \cup x) - f(A)}{f(B \cup x) - f(B)}$$

Observe that f is submodular if and only if $\mathcal{D}(f) \geq 1$. Furthermore, for some $\alpha \geq 1$, we say that a bidder i's valuation function is α *near-submodular* if $\mathcal{D}(v_i) \geq 1/\alpha$.

To illustrate this, recall Example 1. There the valuation function of Dodgers is 7-near submodular: the marginal value of a license is \$1bn when added to the emptyset, but is \$7bn when added on top of the other license. Thus marginal values can increase there but they are bounded by a factor 7.

A similar concept to near-submodularity, called *bounded complementarity*, is introduced in (Lehmann, Lehmann, and Nisan 2006). (Abraham et al. 2012) also consider combinatorial auctions with restrictions on the "size" of complementarities; they use a more restrictive approach where valuation functions are defined using hypergraphs.

We may now state the first of our main results.

Theorem 1. *Given individually rational bidders with α near-submodular valuation and bidding functions, any Nash equilibrium S under VCG has welfare $\omega(\mathcal{S}) \geq \frac{1}{1+\alpha} \cdot \text{OPT}$.*

Here we make the standard auction assumption of *free disposal*: the valuation functions we consider are also monotone, that is, $v_i(A) \leq v_i(B)$ if $A \subset B$. The assumption that every agent is individually rational is also standard: bidders never bid in a way that could result in them receiving negative utility from the auction. Interestingly, though, our second main result shows that the VCG mechanism generates high welfare even if we allow for agents that are *not* individually rational.[3]

Theorem 2. *Given a combinatorial auction where each bidder has α near-submodular valuation, any Nash equilibrium S for the VCG mechanism obtained when one bidder makes false-name-bids has welfare $\omega(\mathcal{S}) \geq \frac{1}{1+\alpha} \cdot \text{OPT}$.*

Here the false-name bidder may submit bid vectors that are not individually rational and that are not α near-submodular. Furthermore, Theorem 2 is almost tight. There are instances where the VCG mechanism will output allocations with welfare at most $\frac{1}{\alpha} \cdot \text{OPT}$. Theorem 2 easily extends to the case where more than one bidder deviates from the norm behavior of submitting a single bid vector. Specifically, if k bidders use pseudonyms then the VCG mechanism will output an allocation of welfare at least $(\frac{1}{1+\alpha})^k \cdot \text{OPT}$, provided all but one of the deviants submit α near-submodular bid vectors.

We conclude this section by emphasizing two points. First, our results apply to all combinatorial auctions and to all valuation functions. No restriction need be imposed upon the bidding functions, but the welfare guarantee will vary with the degree of submodularity of the true valuation functions and of the bidding functions. Of course, our results imply that auctioneer may benefit by imposing restriction on the degree of submodularity allowed in the bidding functions.

Second, high welfare cannot be generated by any known false-name-proof mechanism, i.e. the Set Protocol (Sakurai, Yokoo, and Matsubara 1999), the Level Division Set Mechanism (Yokoo, Sakurai, and Matsubara 2001a) or the Minimal Bundle Mechanism (Yokoo 2003), in the setting of α near-submodular valuation functions. Indeed even when all bidders have submodular ($\alpha = 1$), or even additive, evaluation functions, it is easy to construct examples where all existing false-name-proof mechanisms give a social welfare of $\frac{1}{m} \cdot \text{OPT}$.

False-Name Bidding in VCG Auctions

Recall, as Example 1 shows, even when all bidders' valuation functions are α near-submodular, a bidder can lie and infinitely improve her utility. However, for the auctioneer, whose objective is to maximize social welfare, the corresponding VCG allocation is still reasonably good. The social welfare guarantee resulting from the false-name bidding is bounded by a constant rather than a polynomial as

[3]The nomenclature is unfortunate. In equilibrium analyses, we of course assume that each individual agent is *rational*. However, it is possible at equilibria for rational agents to bid in a manner that is not individually rational!

in FNP mechanisms. When agents bid in an individually rational manner, this is true regardless of how many agents use pseudonyms; when bidders need not be individually rational then it holds provided a constant number of bidders use pseudonyms.

Individually Rational Bidders

For individually rational agents, we prove that every Nash equilibrium has welfare at least $\frac{1}{1+\alpha} \cdot \text{OPT}$ even when *all* the agents use pseudonyms. To show this, we being with some notation. Let the truthful valuation vectors be $\{\mathbf{v_1}, \mathbf{v_2}, \ldots, \mathbf{v_n}\}$, and let the corresponding optimal allocation be $\mathcal{S}^* = \{S_1^*, S_2^*, \ldots, S_n^*\}$. Now assume that S is a Nash equilibrium at which bidder j uses k_j pseudonyms; let pseudonym j_r, $1 \leq r \leq j_{k_j}$ make the false-name bid $\mathbf{b_j^r}$. Given these bids, let the resulting VCG allocation be $\mathcal{S} = \{S_1, \ldots, S_n\}$, where $S_j = \cup_{r=1}^{k_j} S_j^r$ and S_j^r is the set allocated to pseudonym j_r.

Observe that since the VCG mechanism is not truthful under false-name bidding, the mechanism does not optimize the true welfare but instead optimizes declared welfare. So let's distinguish between the true welfare function ω and the declared welfare function $\hat{\omega}$.

Claim 3. *For individually rational bidders, we have $\hat{\omega}(\mathcal{S}) \leq \omega(\mathcal{S})$ for any allocation $\mathcal{S}$.*

Proof. Let $\mathcal{S}$ be an allocation (not necessarily an equilibrium) with false-name bidding and $\mathbf{p}$ the VCG prices. For each bidder j, by individual rationality, we must have

$$v_j(\bigcup_{r=1}^{k_j} S_j^r) \geq \sum_{r=1}^{k_j} b_j^r(S_j^r) \tag{1}$$

Therefore

$$\hat{\omega}(\mathcal{S}) = \sum_j \sum_{r=1}^{k_j} b_j^r(S_j^r) \leq \sum_j v_j(\bigcup_{r=1}^{k_j} S_j^r) = \omega(\mathcal{S}) \quad \square$$

Now assume that the degree of submodularity for each valuation function is $\mathcal{D}(v_i) \geq 1/\alpha$, for each bidder i. (We will assume all false-name bids have this property too.)

Claim 4. *Let S^* and $\mathcal{T}$ be the optimal allocations with and without bidder i, respectively. Then the utility of player i at any Nash equilibrium $\mathcal{S}$ is at least*

$$v_i(S_i^*) - \sum_{j:j\neq i} \sum_{r=1}^{k_j} \left(b_j^r(T_j^r) - b_j^r(T_j^r \setminus S_i^*)\right)$$

Proof. Take any Nash equilibrium $\mathcal{S}$ and suppose that bidder i attempts to deviate from this equilibrium. Without bidder i, let the VCG allocation be $\mathcal{T} = \{T_1, \ldots, T_{i-1}, \emptyset_i, T_{i+1}, \ldots, T_n\}$. Here $T_j = \cup_{r=1}^{k_j} T_j^r$. Thus, one possible deviation is for bidder i to bid truthfully and for the mechanism to select the following feasible allocation $\mathcal{X}$: bidder i is allocated S_i^* and, for $j \neq i$, the pseudonymous bidder j_r is allocated $T_j^r \setminus S_i^*$. Thus the allocation if bidder i bids truthfully has declared value at least

$$\hat{w}(\mathcal{X}) = v_i(S_i^*) + \sum_{j:j\neq i} \sum_{r=1}^{k_j} b_j^r(T_j^r \setminus S_i^*)$$

Without bidder i the declared welfare is $\hat{w}(\mathcal{T})$. So, when bidder i deviates to truthfulness, the VCG mechanism must give bidder i a declared welfare of at least

$$\begin{aligned}
&\hat{w}(\mathcal{X}) - \hat{w}(\mathcal{T}) \\
&= \left(v_i(S_i^*) + \sum_{j:j\neq i} \sum_{r=1}^{k_j} b_j^r(T_j^r \setminus S_i^*)\right) - \sum_{j:j\neq i} \sum_{r=1}^{k_j} b_j^r(T_j^r) \\
&= v_i(S_i^*) - \sum_{j:j\neq i} \sum_{r=1}^{k_j} \left(b_j^r(T_j^r) - b_j^r(T_j^r \setminus S_i^*)\right)
\end{aligned}$$

Since $\mathcal{S}$ is a Nash equilibrium, we have that, by rationality, the true welfare $u_i(\mathcal{S}, \mathbf{p})$ of bidder i is at least this. Thus

$$u_i(\mathcal{S}, \mathbf{p}) \geq v_i(S_i^*) - \sum_{j:j\neq i} \sum_{r=1}^{k_j} \left(b_j^r(T_j^r) - b_j^r(T_j^r \setminus S_i^*)\right)$$

Here $\mathbf{p}$ are the VCG prices for the allocation $\mathcal{S}$. $\square$

Theorem 1. *Given individually rational bidders with α near-submodular valuation and bidding functions, any Nash equilibrium S under VCG has welfare $\omega(\mathcal{S}) \geq \frac{1}{1+\alpha} \cdot \text{OPT}$.*

Proof. The true welfare of $\mathcal{S}$ is the sum of the true welfares of the bidders plus auction revenue. Thus

$$\omega(\mathcal{S}) \geq \sum_i u_i(\mathcal{S}, \mathbf{p}) \tag{2}$$

Applying Claim 4 and summing over all bidders, we have

$$\begin{aligned}
&\sum_i u_i(\mathcal{S}, \mathbf{p}) \\
&\geq \sum_i \left(v_i(S_i^*) - \sum_{j:j\neq i} \sum_{r=1}^{k_j} b_j^r(T_j^r) - b_j^r(T_j^r \setminus S_i^*)\right) \\
&= \sum_i v_i(S_i^*) - \sum_i \sum_{j:j\neq i} \sum_{r=1}^{k_j} \left(b_j^r(T_j^r) - b_j^r(T_j^r \setminus S_i^*)\right) \\
&= \omega(S^*) - \sum_i \sum_{j:j\neq i} \sum_{r=1}^{k_j} \left(b_j^r(T_j^r) - b_j^r(T_j^r \setminus S_i^*)\right) \quad (3)
\end{aligned}$$

To lower bound this, observe that

$$\begin{aligned}
&\sum_j \sum_{r=1}^{k_j} \sum_{i:i\neq j} \left(b_j^r(T_j^r) - b_j^r(T_j^r \setminus S_i^*)\right) \\
&\leq \sum_j \sum_{r=1}^{k_j} \sum_{i:i\neq j} \alpha \cdot \left(b_j^r(T_j^r \setminus \cup_{t=1}^{i-1} S_t^*) - b_j^r(T_j^r \setminus \cup_{t=1}^{i} S_t^*)\right) \\
&= \sum_j \sum_{r=1}^{k_j} \alpha \cdot \left(b_j^r(T_j^r) - b_j^r(T_j^r \setminus \cup_{i:i\neq j} S_i^*)\right) \\
&\leq \sum_j \sum_{r=1}^{k_j} \alpha \cdot b_j^r(T_j^r) \quad (4)
\end{aligned}$$

Here the first inequality follows by α-near submodularity. The equality is due to the telescoping sum. The second inequality follows by non-negativity of the bids. Furthermore

$$\begin{aligned}\sum_j \sum_{r=1}^{k_j} b_j^r(T_j^r) &\leq \sum_j \sum_{r=1}^{k_j} b_j^r(S_j^r) \leq \sum_j v_j(\cup_{r=1}^{k_j} S_j^r) \\ &= \sum_j v_j(S_j) = \omega(\mathcal{S}) \end{aligned} \tag{5}$$

Here, the first inequality follows from the optimality of $\mathcal{S}$ with respect to the bids $\mathbf{b}$. The second inequality follows by the individual rationality of the bidders, specifically from (1). The two equalities follow by definition.

Plugging (4) and (5) into (3) gives

$$\sum_i u_i(\mathcal{S}, \mathbf{p}) \geq \omega(S^*) - \alpha \cdot \omega(\mathcal{S}) \tag{6}$$

Thus (2) and (6) give $\omega(\mathcal{S}) \geq \frac{1}{1+\alpha}\, \omega(S^*)$. □

Bidders that are not Individually Rational

We now show that high welfare guarantees are provided by the VCG mechanism even when we discard the assumption that bidders are individually rational! To begin, we assume there is a unique bidder, whom we distinguish as bidder 0, who wishes to make multiple false-name bids. We show that if each non-false-name bidder's valuation function is α near-submodular then the social welfare of VCG is at least $\frac{1}{1+\alpha} \cdot$ OPT. This bound is almost tight; we give an example where the social welfare of VCG can be as low as $\frac{1}{\alpha} \cdot$ OPT in the presence of false-name bidding.

Denote the set of truthful bid vectors by $\{\mathbf{b_0}, \mathbf{b_1}, \ldots, \mathbf{b_n}\}$, and let the corresponding optimal allocation be $\mathcal{S}^* = \{S_0^*, S_1^*, \ldots, S_n^*\}$. Let the corresponding optimal welfare be $\text{OPT} = \omega(\mathcal{S}^*) = \sum_{i=0}^n v_i(S_i^*)$ and the VCG prices be $\mathbf{p}^*$. Then we may assume that all the players make (truthful) near-submodular bids, except for bidder 0 who makes the bids $\{\mathbf{b_0^1}, \mathbf{b_0^2}, \ldots, \mathbf{b_0^k}\}$. Given these bids, let the resulting VCG allocation be $\mathcal{S} = \{S_0 = \cup_{j=1}^k S_0^j, S_1, \ldots, S_n\}$, where S_0^j is the set allocated to pseudonym 0_j. Let the corresponding VCG price vector be $\mathbf{p}$. Then the total price paid by bidder 0 is $p_0 = \sum_{j=1}^k p_0^j$.

First we show that the price vector satisfies the inequality given below in Lemma 5. There $\mathcal{T} = \{\emptyset_0, T_1, \ldots, T_n\}$ is the optimal allocation when bidder 0 is allocated nothing.

Lemma 5. *For any collection of valuation functions,*

$$\sum_{i=0}^n v_i(S_i) \geq \text{OPT} - \sum_{i=1}^n v_i(T_i) + p_0 + \sum_{i=1}^n v_i(S_i)$$

Proof. Observe that the false-name bids $\{\mathbf{b_0^1}, \mathbf{b_0^2}, \ldots, \mathbf{b_0^k}\}$ are rational for bidder 0 only if

$$v_0(S_0^*) - p_0^* \leq v_0(S_0) - p_0 \tag{7}$$

In the truthful case, the VCG mechanism charges bidder 0

$$p_0^* = \sum_{i=1}^n v_i(T_i) - \sum_{i=1}^n v_i(S_i^*) \tag{8}$$

Combining (8) with the rationality constraint (7) produces

$$\begin{aligned} v_0(S_0) &\geq v_0(S_0^*) - p_0^* + p_0 \\ &= v_0(S_0^*) - \sum_{i=1}^n v_i(T_i) + \sum_{i=1}^n v_i(S_i^*) + p_0 \\ &= \text{OPT} - \sum_{i=1}^n v_i(T_i) + p_0 \end{aligned}$$

Therefore

$$\begin{aligned} \sum_{i=0}^n v_i(S_i) &= v_0(S_0) + \sum_{i=1}^n v_i(S_i) \\ &\geq \text{OPT} - \sum_{i=1}^n v_i(T_i) + p_0 + \sum_{i=1}^n v_i(S_i) \quad \square \end{aligned}$$

For a worst-case analysis, we may assume that each pseudonym bidder 0_j is single-minded and bids only on the set S_0^j, and that in the absence of false-name bidder 0_j the other false-name bidders still win in the optimal allocation. Otherwise the VCG price to bidder 0 will be larger, as it will incorporate charges for the damage the pseudonym bidders do to each other.

Lemma 6. *For any collection of α near-submodular valuation functions, we have*

$$p_0 \geq \frac{1}{\alpha} \cdot \sum_{i=1}^n (v_i(T_i \cap S_0) - v_i(S_i))$$

Proof. Observe that the price p_0^j that false-name bidder j pays is lower bounded by the value of directly reallocating S_0^j to the truthful bidders *given* that they have already won $\{S_1, \ldots, S_n\}$. Since the allocation $\{T_1, \ldots, T_n\}$ is still a feasible allocation in the false-name case, a feasible solution when false-name bidder j is omitted is for bidder i to win $(T_i \backslash S_0) \cup (T_i \cap S_0^j)$.

Furthermore, an alternative feasible solution when false-name bidder j is omitted is for bidder i to win $S_i \cup (T_i \cap S_0^j)$. Thus the damage the false-name bidders do to the truthful bidders is at least

$$\begin{aligned} &\sum_{j=1}^k \sum_{i=1}^n \Big(v_i(S_i \cup (T_i \cap S_0^j)) - v_i(S_i)\Big) \\ &= \sum_{i=1}^n \sum_{j=1}^k \Big(v_i(S_i \cup (T_i \cap S_0^j)) - v_i(S_i)\Big) \\ &\geq \sum_{i=1}^n \sum_{j=1}^k \frac{1}{\alpha}\Big(v_i(S_i \cup (T_i \cap \textstyle\bigcup_{r=1}^j S_0^r)) - v_i(S_i \cup (T_i \cap \bigcup_{r=1}^{j-1} S_0^r))\Big) \\ &= \sum_{i=1}^n \frac{1}{\alpha} \cdot \left(v_i(S_i \cup (T_i \cap \bigcup_{r=1}^k S_0^r)) - v_i(S_i)\right) \\ &= \frac{1}{\alpha} \cdot \sum_{i=1}^n (v_i(S_i \cup (T_i \cap S_0)) - v_i(S_i)) \end{aligned}$$

Here, the inequality above follows by near-submodularity, and the second equality follows as the sum telescopes. Thus

$$p_0 \quad \geq \quad \frac{1}{\alpha} \cdot \sum_{i=1}^{n} (v_i(T_i \cap S_0) - v_i(S_i)) \quad \Box$$

Recall, our aim is to show that $\omega(\mathcal{S})$ is comparable to the optimal welfare $\omega(\mathcal{S}^*)$. To do this, we essentially partition up the analysis by case depending upon how much damage bidder 0 does to the other players by participating *and* using pseudonyms. Specifically, we examine all possible settings of β such that

$$\sum_{i=1}^{n} v_i(S_i) \quad = \quad \beta \cdot \omega(\mathcal{T}) \; = \; \beta \cdot \sum_{i=1}^{n} v_i(T_i) \qquad (9)$$

and then calculate the worst-case welfare guarantee over all these possibilities. (Note that $\beta \leq 1$.)

To gain some intuition behind the proof, consider the following. Clearly, if both β and $\omega(\mathcal{T})$ are large then so is $\sum_{i=1}^{n} v_i(S_i)$. In this case, since $\omega(\mathcal{S}) \geq \sum_{i=1}^{n} v_i(S_i)$, the outcome is high social welfare. On the other hand, if $\omega(\mathcal{T})$ is small in comparison to $\omega(\mathcal{S}^*)$ then bidder 0 must do very well in $\mathcal{S}^*$ and, hence, in $\mathcal{S}$. Thus, again, $\mathcal{S}$ gives high social welfare. Consequently, the difficult case is when β is small. To analyse this, we require the following claim.

Lemma 7. *Let $\sum_{i=1}^{n} v_i(S_i) = \beta \cdot \sum_{i=1}^{n} v_i(T_i)$. Then*

$$\sum_{i=1}^{n} v_i(T_i \cap S_0) \geq (1 - \alpha \cdot \beta) \cdot \sum_{i=1}^{n} v_i(T_i)$$

Proof. By near-submodularity, we have

$$\begin{aligned} v_i(T_i) - v_i(T_i \cap S_0) \quad &= \quad v_i((T_i \cap S_0) \cup (T_i \backslash S_0)) - v_i(T_i \cap S_0) \\ &\leq \quad \alpha \cdot (v_i(T_i \backslash S_0) - v_i(\emptyset)) \\ &= \quad \alpha \cdot v_i(T_i \backslash S_0) \end{aligned}$$

Summing over the bidders gives:

$$\begin{aligned} \sum_{i=1}^{n}(v_i(T_i) - v_i(T_i \cap S_0)) \quad &= \quad \alpha \cdot \sum_{i=1}^{n} v_i(T_i \backslash S_0) \\ \leq \; \alpha \cdot \sum_{i=1}^{n} v_i(S_i) \quad &= \quad \alpha \cdot \beta \sum_{i=1}^{n} v_i(T_i) \end{aligned}$$

Here the inequality follows from the optimality of the mechanism. Specifically, if S_0 is allocated to (copies of) bidder 0 then $\{S_1, \ldots, S_n\}$ is a better allocation to the other players than $\{T_1 \backslash S_0, \ldots, S_n \backslash S_0\}$. Rearranging, we obtain

$$\sum_{i=1}^{n} v_i(T_i \cap S_0) \quad \geq \quad (1 - \alpha \cdot \beta) \cdot \sum_{i=1}^{n} v_i(T_i) \quad \Box$$

We are now ready to prove the second main theorem.

Theorem 2. *For any collection of α near-submodular valuation functions, the VCG mechanism outputs an allocation with $\omega(\mathcal{S}) \geq \frac{1}{1+\alpha} \cdot \text{OPT}$.*

Proof. Applying Lemma 5 and (9) we have

$$\begin{aligned} \omega(\mathcal{S}) \quad &= \quad \sum_{i=0}^{n} v_i(S_i) \\ &\geq \quad \text{OPT} - \sum_{i=1}^{n} v_i(T_i) + p_0 + \sum_{i=1}^{n} v_i(S_i) \\ &\geq \quad \text{OPT} - \sum_{i=1}^{n} v_i(T_i) + p_0 + \beta \cdot \sum_{i=1}^{n} v_i(T_i) \\ &= \quad \text{OPT} - (1 - \beta) \cdot \sum_{i=1}^{n} v_i(T_i) + p_0 \qquad (10) \end{aligned}$$

Therefore

$$\begin{aligned} p_0 \quad &\geq \quad \frac{1}{\alpha} \cdot \sum_{i=1}^{n} (v_i(T_i \cap S_0) - v_i(S_i)) \\ &\geq \quad \frac{1}{\alpha} \cdot \sum_{i=1}^{n} v_i(T_i) \cdot (1 - \alpha\beta - \beta) \\ &= \quad \frac{1 - (\alpha + 1) \cdot \beta}{\alpha} \cdot \sum_{i=1}^{n} v_i(T_i) \qquad (11) \end{aligned}$$

Here the first inequality is Lemma 6. The second inequality follows by Lemma 7 and (9). Substituting (11) into (10) gives

$$\begin{aligned} \omega(\mathcal{S}) \quad &\geq \quad \text{OPT} - \left(1 - \beta - \frac{1 - (\alpha + 1) \cdot \beta}{\alpha}\right) \cdot \sum_{i=1}^{n} v_i(T_i) \\ &= \quad \text{OPT} - \left(1 - \frac{1}{\alpha} \cdot (1 - \beta)\right) \cdot \sum_{i=1}^{n} v_i(T_i) \qquad (12) \\ &\geq \quad \text{OPT} - \left(1 - \frac{1}{\alpha} \cdot (1 - \beta)\right) \cdot \text{OPT} \\ &\geq \quad \left(\frac{1}{\alpha} \cdot (1 - \beta)\right) \cdot \text{OPT} \qquad (13) \end{aligned}$$

So we obtain the lower bound (13) on social welfare. When β is large this guarantee is weak. However, by (9), we know

$$\omega(\mathcal{S}) \quad \geq \quad \sum_{i=1}^{n} v_i(S_i) = \beta \cdot \sum_{i=1}^{n} v_i(T_i) \qquad (14)$$

Combining inequality (12) with (14) we obtain

$$\left(1 - \frac{1}{\alpha} \cdot (1 - \beta) + \beta\right) \cdot \omega(\mathcal{S}) \quad \geq \quad \beta \cdot \text{OPT}$$

In turn, this rearranges to

$$\omega(\mathcal{S}) \quad \geq \quad \frac{\beta}{1 - \frac{1}{\alpha} + \beta(1 + \frac{1}{\alpha})} \cdot \text{OPT} \qquad (15)$$

Finally, together the lower bounds (15) and (13) give

$$\begin{aligned} \omega(\mathcal{S}) \quad &\geq \quad \max\left[\frac{1}{\alpha} \cdot (1 - \beta) \, , \, \frac{\beta}{1 - \frac{1}{\alpha} + \beta(1 + \frac{1}{\alpha})}\right] \cdot \text{OPT} \\ &\geq \quad \frac{1}{1 + \alpha} \cdot \text{OPT} \end{aligned}$$

Here the second inequality holds since the maximum is minimized by setting $\beta = \frac{1}{\alpha+1}$. $\Box$

Theorem 2 is almost tight; see appendix for an example.

References

Abraham, I.; Babaioff, M.; Dughmi, S.; and Roughgarden, T. 2012. Combinatorial auctions with restricted complements. In *Proc. of EC*, 3–16.

Aziz, H., and Paterson, M. 2009. False name manipulations in weighted voting games: Splitting, merging and annexation. In *Proc. of AAMAS*, 409–416.

Bachrach, Y., and Elkind, E. 2008. Divide and conquer: False-name manipulations in weighted voting games. In *Proc. of AAMAS*, 975–982.

Conitzer, V.; Immorlica, N.; Letchford, J.; Munagala, K.; and Wagman, L. 2010. False-name-proofness in social networks. In *Proc. of WINE*, 209–221.

Conitzer, V. 2007. Limited verification of identities to induce false-name-proofness. In *Proc. of TARK*, 102–111.

Cramton, P.; Shoham, Y.; and Steinberg, R. 2006. *Combinatorial Auctions*. MIT Press.

Guo, M., and Conitzer, V. 2010. False-name-proofness with bid withdrawl. In *Proc. of AAMAS*, 1475–1476.

Iwasaki, A.; Conitzer, V.; Omori, Y.; Sakurai, Y.; Todo, T.; Guo, M.; and Yokoo, M. 2010. Worst-case efficiency ratio in false-name-proof combinatorial auction mechanisms. In *Proc. of AAMAS*, 633–640.

Lehmann, B.; Lehmann, D.; and Nisan, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55(3):270–296.

Rothkopf, M. 2007. Thirteen reasons why the vickrey-clarke-groves process is not practical. *Operations Research* 55(3):191–197.

Sakurai, Y., and Yokoo, M. 2003. A false-name-proof double auction protocol for arbitrary evaluation values. In *Proc. of AAMAS*, 329–336.

Sakurai, Y.; Yokoo, M.; and Matsubara, S. 1999. A limitation of the generalized vickrey auction in electronic commerce: Robustness against false-name bids. In *Proc. of AAAI*, 86–92.

Suyama, T., and Yokoo, M. 2004. Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction. In *Proc. of AAMAS*, 160–167.

Todo, T.; Iwasaki, A.; Yokoo, M.; and Sakurai, Y. 2009. Characterizing false-name-proof allocation rules in combinatorial auctions. In *Proc. of AAMAS*, 265–272.

Todo, T.; Mouri, T.; Iwasaki, A.; and Yokoo, M. 2012. False-name-proofness in online mechanisms. In *Proc. of AAMAS*, 753–762.

Todo, T. 2010. Characterization of false-name-proof social choice mechanisms. In *Proc. of AAMAS*, 1675–1676.

Wagman, L., and Conitzer, V. 2008. Optimal false-name-proof voting rules with costly voting. In *Proc. of AAAi*, 190–195.

Yokoo, M.; Sakurai, Y.; and Matsubara, S. 2001a. Bundle design in robust combinatorial auction protocol against false-name bids. In *Proc. of IJCAI*, 1054–1101.

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

Yokoo, M.; Sakurai, Y.; and Matsubara, S. 2001c. Robust double auction protocol against false-name bids. In *Proc. of ICDCS*, 137–145.

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

Yokoo, M. 2003. Characterization of strategyfalse-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In *Proc. of IJCAI*, 733–739.

Appendix

Here is an almost tight example for Theorem 2.

Theorem 8. *There are combinatorial auctions with α near-submodular valuation functions where the VCG mechanism outputs an allocation with welfare* $\frac{1}{\alpha} \cdot \text{OPT}$.

Proof. To see this, consider the following example. There are two bidders and k items. Bidder 0 has an additive valuation function with value $1+\delta$ for each item. Bidder 1 has value $1+(r-1)\cdot\alpha$ for any subset containing exactly $r \geq 1$ items. This function is α-near submodular. The optimal allocation $\mathcal{S}^*$ is to assign all the items to bidder 1 for a total welfare of $1+(k-1)\cdot\alpha$.

Now suppose bidder 0 creates k false names. The jth false-name bidder 0_j bids α for item j. The VCG allocation $\mathcal{S}$ then assigns all the items to the false-name bidders. The mechanism thinks false-name bidder 0_j contributes value $\alpha - 1$ so it is charged 1. Overall, therefore, bidder 0 wins all the items for a price k. The true value of this allocation to bidder 0 is $k(1+\delta)$, so this false set of bids is beneficial to Bidder 0. The welfare of this allocation is

$$\omega(\mathcal{S}) \quad = \quad \frac{k(1+\delta)}{1+\alpha(k-1)} \cdot \text{OPT} \qquad (16)$$

For small δ, the fraction in (16) tends to $\frac{1}{\alpha}$ as k gets big. $\square$