

Optimal Shill Bidding in the VCG Mechanism

Itai Sher*

University of Minnesota

This Version: February 2009

(First Version: December 29, 2008)

Abstract

This paper studies shill bidding in the VCG mechanism applied to combinatorial auctions. Shill bidding is a strategy whereby a single decision-maker enters the auction under the guise of multiple identities (Sakurai, Yokoo, and Matsubara 1999). I formulate the problem of *optimal* shill bidding for a bidder who knows the aggregate bid of her opponents. A key to the analysis is a subproblem—the *cost minimization problem*—which searches for the cheapest way to win a given package using shills. This formulation leads to an exact characterization of the aggregate bids b such that some bidder would have an incentive to shill bid against b . It is well known that when goods are substitutes, there is no incentive to shill bid. In contrast, I show that when goods are *pure* complements, the incentive to shill takes a simple form: there is an incentive to disintegrate and bid for each item using a different identity. With a mix of substitutes and complements, I show that the *winner determination problem* (for single minded bidders)—the problem of finding an efficient allocation in a combinatorial auction—can be embedded into the optimal shill bidding problem. Shill bidding is closely related to collusion. Setting aside the ordinary incentive to suppress competition, the disincentive to disintegrate using shills when facing a substitutes valuation translates into an incentive to merge for a coalition facing the same valuation. Only when valuations are additive can the incentives to shill and merge simultaneously disappear. The paper also shows that there does not exist a dominant strategy in the VCG mechanism when shill bidding is possible. I find a large class of shill bidding strategies which sometimes outperform truthful bidding, but also show that no shill bidding strategy dominates truthful bidding.

JEL Classification: C72, D44

Keywords: shill bidding, VCG mechanism, combinatorial auctions, winner determination problem, collusion

*I am grateful to Vincent Conitzer and David Rahman for useful discussions. I am also grateful to seminar participants at The University of Texas at Austin, The University of Wisconsin and Madison, and at the Midwest Economic Theory Meetings at Ohio State University. All errors are my own.

Contents

1	Introduction	1
2	Preliminaries	3
2.1	Combinatorial Auctions	3
2.2	The VCG Mechanism	4
2.3	Shill Bidding	5
2.4	Partitional Bid Profiles	6
2.5	Worst Case Analysis vs. Equilibrium Analysis	7
3	Optimal Shill Bidding	7
3.1	Shill Bidding Cost Minimization Problem	8
3.2	Shill Prices vs. VCG Prices	9
3.3	Optimal Shill Bidding Problem	10
3.4	Proof and Discussion of Theorems 1 and 2	11
4	Characterization of the Incentive to Shill	13
4.1	Comparison to a Result by Lehmann, Lehmann, and Nisan (2006)	14
4.2	Comparison to a Result by Ausubel and Milgrom (2002)	16
5	Pure Complements	17
6	Mixture of Substitutes and Complements:	
	Computational Complexity	21
6.1	The Winner Determination Problem	21
6.2	Reduction from The Winner Determination Problem	22
6.3	Relation to a Results by Sanghvi and Parkes (2004) and Conitzer and Sandholm (2006)	24
6.4	Interim Summary: Substitutes, Complements, and a Mixture of the Two	25
7	Shill Bidding and Dominance	25
8	The Riskiness of Shill Bidding	27
9	Collusion	28
9.1	VCG prices vs. Merged Prices	29
9.2	Collusion in The Second Price Auction For a Single Item	29
9.3	Competition and Integration Effects	30
9.4	Merging is the Inverse of Shilling	32
9.5	The Optimal Collusive Strategy	34
9.6	Efficiency	35
10	Conclusion	36
11	Appendix	37

1 Introduction

The Vickrey-Clarke-Groves (VCG) mechanism serves as an important benchmark for combinatorial auctions—or in other words, auctions in which many items are sold simultaneously and bidders may submit bids on packages of items. Assuming transferable utility, the VCG mechanism is essentially the unique mechanism that implements efficient outcomes in dominant strategies on smoothly connected domains (Green and Laffont 1977, Holmstrom 1979). On the other hand, the VCG mechanism is known to suffer from a host of problems (Milgrom 2004, Ausubel and Milgrom 2006, Rothkopf 2007). The purpose of this paper is to contribute to the assessment of one such problem: namely, the problem of *skill bidding*, which was introduced by Sakurai, Yokoo, and Matsubara (1999). Consider a VCG auction conducted over the internet in which the seller observes screen names belonging to the different bidders participating in the auction. Suppose, however, that the seller cannot verify that these different screen names actually correspond to different bidders. In this case, when the items on auction may be complements, a bidder may have an incentive to enter the auction under the guise of multiple identities, and manipulate the mechanism. This problem does not require the internet—but may occur in a VCG auction whenever there are multiple bidders who appear independent but who actually bid on behalf of a single decision-maker.¹

The main contribution of this paper is to study *optimal* skill bidding. Let Ann be a bidder who is considering the use of skills. This paper studies the following optimization problem: if Ann knows or guesses the profile of bids submitted by her opponents, what is her optimal use of skills? Theorems 1 and 2 formulate Ann’s optimization problem. This formulation decomposes Ann’s problem into two parts: (i) for any package B , what is the cheapest way that Ann could win B using skills? (ii) which package does Ann want to win given that she wins it in the cheapest way? (i) is referred to as the *skill bid cost minimization problem* (CMP). The theorems also show that the optimal skill bidding strategy will always have the characteristic that the different skills which are sponsored by Ann do not compete with one another.

Theorem 3 uses the formulation of Ann’s problem to characterize the circumstances in which some bidder may have an incentive to use skills. The characterization is related to similar results by Lehmann, Lehmann, and Nisan (2006) and Ausubel and Milgrom (2002).²

¹Rastegari, Condon, and Leyton-Brown (2007) and Day and Milgrom (2008) show that skill bidding can be a problem not just in the VCG mechanism, but for a larger class of auctions. Several papers have proposed solutions to the skill bidding problem, either in the form of limited verification of identities (Conitzer 2007) or in the form of proposed auctions which do not suffer from the skill bidding problem (Ausubel and Milgrom 2002, Matsuo, Takayuki, Day, and Shintani 2006, Yokoo 2006, Yokoo and Iwasaki 2007, Day and Milgrom 2008). Of course, as these papers point out, verification of identities may sometimes be infeasible, and the auctions which avoid the skill bidding problem do not have all of the attractive features of the VCG mechanism.

²Another characterization is due to Yokoo, Sakurai, and Matsubara (2004); this characterization is slightly

In general, shill bidding is not worthwhile if goods are substitutes. At the other extreme, when goods are pure complements (by which I mean that Ann faces an aggregate bid which is supermodular), optimal shill bidding takes on a simple form; there is an incentive for Ann to disintegrate and bid for each item that she would like to win, using a different identity. The intermediate case, in which Ann faces a mix of complements and substitutes, is complex: Ann may want to break up into different identities to some degree, but each shill bidder may bid for a package and not just an item. I show that there is a close relationship between the cost minimization part of the shill bidding problem—part (i)—and the *winner determination problem* (WDP), which is the problem of finding an efficient allocation in a combinatorial auction. This relationship provides an economically interesting interpretation of the optimal shill bidding problem. WDP is known to be an NP-hard problem; the NP-hardness of optimal shill bidding also follows from Sanghvi and Parkes (2004) (See Section 6.3).³

The analysis of shill bidding is a *worst-case* analysis. The optimal shill bidding strategy is analyzed under the assumption that Ann knows or can correctly guess the aggregate bid of her opponents. This is a highly unrealistic assumption, but it is meant to show that in the worst case, Ann may have an incentive to attempt to manipulate the mechanism. Whether this will actually happen in equilibrium is another matter. What is clear is that the dominant strategy incentive compatibility of the VCG mechanism does not survive the introduction of shill bidding.⁴ I go on to show that not only is truthful bidding no longer a dominant strategy in the VCG mechanism with shill bidding, but there does not exist any dominant strategy in the VCG mechanism with shill bidding. Moreover, I find a large class of strategies which are sometimes better than truthful bidding when shill bidding is possible. I also examine the worst case from Ann's perspective. In some cases, the strategies which solve Ann's shill bidding optimization problem are risky in the sense that if Ann is wrong about the aggregate bid of her opponent, Ann may receive a negative utility, which is not possible under truthful bidding. However, I show that sometimes, Ann has an optimal shill bidding strategy which gives Ann a zero utility in the worst case, matching the performance of truthful bidding along this dimension. This is possible even in cases when optimal shill bidding leads to an inefficient allocation.

Finally in Section 9, I show that the analysis in this paper has relevance to more general forms of collusion. In the VCG mechanism, there is almost always an incentive to collude in order to reduce the prices. This is because losing bids for packages determine the opportunity cost, and hence the VCG price of that package. However, given a collection of

further removed from the current characterization than those mentioned above.

³Conitzer and Sandholm (2006) have a related—but distinct—result on the complexity of collusive strategies in the VCG mechanism. The relationship is also discussed in Section 6.3.

⁴Some authors that study this problem refer to the VCG mechanism as dominant strategy incentive compatible, but say that it is not shill proof. I prefer to say that the VCG mechanism is dominant strategy incentive compatible if shill bidding is impossible, but not if shill bidding is possible.

bidders J and a collection of packages $(B_j : j \in J)$ that they would win at the efficient allocation, one can ask: what is the minimum payment possible for the members of J subject to the constraint that each member j of J wins B_j , holding fixed the bids of bidders outside of J ? Given that bidders suppress competition in this way, the question remains as to whether they would be better off merging, and bidding under a single identity. From this perspective, merging is simply the inverse of shill bidding. Whereas complementarities in the aggregate bid of one's opponents make shill bidding attractive, complementarities make merging unattractive. In contrast, when the opposing aggregate bid is a substitutes valuation, shill bidding is unattractive, but merging is attractive.⁵ In general—whether considering shill bidding or collusion—substitutes valuations create an incentive for integration, and complements valuations create an incentive for disintegration.

2 Preliminaries

In this section, I introduce several concepts and issues which will be necessary for the development of the main ideas of the paper.

2.1 Combinatorial Auctions

A combinatorial auction is an auction in which bidders may bid for packages of goods. Formally, assume a finite collection N of goods, and a finite collection I of bidders. Each bidder $i \in I$ has a valuation $v_i : 2^N \rightarrow \mathbb{R}_+$, which assigns to each package the maximum amount that bidder i would be willing to pay to receive exactly that package. If i receives package B for a price p , his utility is $v_i(B) - p$. We assume that $v_i(\emptyset) = 0$ and:

$$A \subseteq B \Rightarrow v_i(A) \leq v_i(B)$$

The latter assumption is called **monotonicity**, or alternatively **free disposal**.

Actual or potential uses of combinatorial actions of actual or potential uses of combinatorial auctions include auctions for provision of transportation services (Caplice and Sheffi 2006, Cantillon and Pesendorfer 2006) industrial procurement (Bicheler, Davenport, Hohner, and Kalagnanam 2006), arrival and departure times at airports (Rassenti, Smith, and Bulfin 1982), and use of radio spectrum (Milgrom 2000). Complementarities are often important in combinatorial auctions. For example, in a spectrum auction, a firm may need to win licenses for several regions in order to have a viable business, and therefore the firm may assign a value to the package of licences for this region which is higher than the sum of the values of the stand alone licenses.

⁵Milgrom (2004) presents an example which illustrates the incentive to merge in the VCG mechanism when goods are substitutes. In this paper, I provide an exact and general characterization of such incentives.

2.2 The VCG Mechanism

The VCG Mechanism is an example of a combinatorial auction. In the VCG Mechanism a bid is a valuation. The VCG mechanism then implements an efficient allocation taking the bids at face value. In other words, the goods are divided among the bidders in such a way as to maximize the sum of their reported valuations. Formally, let

$$\mathcal{X} := \{(X_i : i \in I) : \forall i, j \in I, X_i \subseteq N, X_i \cap X_j = \emptyset\}$$

be the set of allocations, where X_i is the package received by i . Similarly let \mathcal{X}_{-i} be the set of allocations such that i does not receive any items, or equivalently such that $X_i = \emptyset$. Let $(v_i : i \in I)$ be the profile of bids, and let X_i^* be the package which is assigned to i at the efficient allocation. In other words $(X_i^* : i \in I)$ solves:

$$\max\left\{\sum_{i \in I} v_i(X_i) : (X_i : i \in I) \in \mathcal{X}\right\}$$

If there are multiple efficient allocations, then the VCG mechanism selects one such allocation according to some tie-breaking rule. The VCG payment of bidder i is:

$$p_i = \underbrace{\max\left\{\sum_{j \in I \setminus i} v_j(X_j) : (X_j : j \in I \setminus i) \in \mathcal{X}_{-i}\right\}}_{(*)} - \underbrace{\sum_{j \in I \setminus i} v_j(X_j^*)}_{(**)}$$

Term (*) refers to the value of the efficient allocation in the marginal economy excluding bidder i , and term (**) refers to the total value to bidders other than i at the efficient allocation in the economy including i . The difference between (*) and (**) can be interpreted as the externality imposed by bidder i on other bidders, or alternatively as the opportunity cost of the items won by i . Truthful bidding—or in other words, submitting a bid equal to one's value—is a dominant strategy in the VCG mechanism. The VCG mechanism is essentially unique efficient auction in which truthful bidding is a dominant strategy (Green and Laffont 1977, Holmstrom 1979). By essentially unique, I mean that for each bidder i the payments may be altered by a constant depending on the reports of other bidders. The VCG mechanism is known to suffer from a host of problems (Milgrom 2004, Ausubel and Milgrom 2006, Rothkopf 2007) and it is not commonly in use for selling multiple heterogeneous objects. The motivation for studying this mechanism is that—as explained above—it is the unique auction with certain desirable properties, and we would like to understand how these properties inevitably lead to certain defects.

2.3 Shill Bidding

Imagine that a VCG auction is conducted over the internet, and that a single bidder may enter the auction under multiple screen names. Suppose that it is impossible to verify that different screen names correspond to different bidders. The incentive to undertake such a manipulative strategy in the VCG mechanism is related other forms of collusion and has been studied in a series of papers (Sakurai, Yokoo, and Matsubara 1999, Yokoo, Sakurai, and Matsubara 2000, Yokoo, Sakurai, and Matsubara 2004, Lehmann, Lehmann, and Nisan 2006, Ausubel and Milgrom 2006, Conitzer and Sandholm 2006). The following example shows how shill bidding may be effective.

Example 1 *Suppose that there are two bidders, Ann and Bob, and two goods, good 1 and good 2. Bob values an individual good at \$1, but the package containing both at \$4. Ann values each good at \$1.50 and the package at \$3. It is efficient to give the package to Bob. Under truthful bidding, the VCG mechanism gives package to Bob and charges Bob \$3. Now suppose that, rather than bidding truthfully, Ann enters auction under two identities i and j . i claims only to value good 1 at \$3.50, and j claims only to value good 2 at \$3.50. Taking these bids at face value, it would be efficient to give bidder i good 1 and bidder j good 2. Thus ultimately Ann would receive both items. Bidder i 's VCG payment would be $\$1 = \$4.50 - \$3.50$. \$4.50 is the total value to all bidders in the marginal economy excluding i : j would still receive good 2, and Bob would receive good 1. \$3.50 is the combined value to j and Bob in the economy including i . Similar reasoning shows that j 's VCG payment would also be \$1. Since Ann is responsible for the payments of all of her shill bidders, Ann's total payment is \$2. Since Ann receives both goods, her utility is $3 - 2 = 1$, which is higher than her utility to truthful bidding. One can show that the strategy described here is Ann's optimal use of skills in this example.*

The main problem is that under shill bidding, a decision-maker no longer corresponds to a single bidder. The idea behind the VCG mechanism is that each *decision-maker* is charged his externality. However, with shill bidding, it is no longer possible to identify the decision-maker.

This paper will focus on two characters, Ann and Bob. Ann is the character who shill bids. Therefore, rather than submitting a single bids, Ann may submit a profile of bids ($v_j : j \in J$). Ann selects not only the values v_j but also the (finite) set J . When Ann submits ($v_j : j \in J$), the VCG auction is run as if these were distinct legitimate bids. Bob is not interpreted as Ann's only opponent, but rather as Ann's *aggregate* opponent. In other words, Ann is assumed to face a set I of opponents, who submit bids ($v_i : i \in I$). Bob's bid is then defined by:

$$v_b(A) := \max \left\{ \sum_{i \in I} v_i(X_i) : \bigcup_{i \in I} X_i = A, i \neq j \Rightarrow X_i \cap X_j = \emptyset \right\}$$

Ann’s payment and allocation—whether or not she uses skills—depends only on the aggregate bid of her opponent.⁶ Therefore, the analysis that follows applies to situations in which Ann faces an arbitrary (finite) number of opponents.

2.4 Partitional Bid Profiles

For any $B \subseteq N$, let $\Pi(B)$ be the set of partitions on B . Formally, a partition $\mathcal{P} \in \Pi(B)$ is a collection of nonempty subsets of B satisfying:

$$\begin{aligned} \forall P, P' \in \mathcal{P} \quad P \neq P' &\Rightarrow P \cap P' = \emptyset \\ \bigcup \mathcal{P} &= B \end{aligned}$$

In other words, the elements of \mathcal{P} are pairwise disjoint and the union of all elements in \mathcal{P} is equal to B . An element P of the partition \mathcal{P} is called a **cell** of \mathcal{P} . For example, $\mathcal{P} = \{\{1, 2\}, \{3\}\}$ is a partition of $B = \{1, 2, 3\}$, and so belongs to $\Pi(B)$. $P = \{1, 2\}$ is a cell of \mathcal{P} .

For any $A, B \subseteq N$ and $r \in \mathbb{R}_+$, define:

$$v^{B,r}(A) := \begin{cases} r, & \text{if } B \supseteq A; \\ 0, & \text{otherwise.} \end{cases}$$

Valuations of the form $v^{B,r}$ are **single-minded**. A bidder with a single-minded valuation $v^{B,r}$ assigns a value of r to package B but does not care about anything else. If the bidder fails to receive even one item in B , his utility is 0, and his marginal utility for additional items is 0 once he has received B .

A **partitional bid profile** is a profile of valuations of the form $(v^{P,r_P} : P \in \mathcal{P})$ for some partition \mathcal{P} of some package B , where for all $P \in \mathcal{P}$, $r_P \in \mathbb{R}_+$. In other words, a partitional bid profile is a profile of pairwise disjoint single-minded bids. For example consider the bid profile $(v^{\{1,2\},7}, v^{\{3\},10})$ is a partitional profile corresponding to the partition $\mathcal{P} = \{\{1, 2\}, \{3\}\}$ of $B = \{1, 2, 3\}$. One bidder bids single-mindedly for the package $\{1, 2\}$ at a value of 7, and the other bidder bids single-mindedly for the package $\{3\}$ at a value of 10. Note that the bidders using a partitional bid profile do not compete with one another in the sense that they bid on mutually disjoint collections of items. If Ann shill bids using a partitional bid profile $(v^{P,r_P} : P \in \mathcal{P})$, then we can identify each of her shill bidders with

⁶The one possible exception is when there are multiple efficient allocations and the VCG mechanism selects one. In that case, in principle, depending on how the VCG mechanism breaks ties, different configurations of bids by the members of I which lead to the same aggregate valuation v_b may lead to different allocations and payments for Ann because they effect the outcome of the tie-breaking rule. If Ann bids truthfully, then Ann is indifferent among all of these allocation-payment combinations, but if Ann uses shills she might not be indifferent. In the main body of the text, I will ignore this knife-edge case. None of the theorems would be altered by taking this possibility seriously. However, the exposition would definitely be lengthened if I were explore this possibility.

the package P that he bids for; in other words, bidder P bids for package P .

2.5 Worst Case Analysis vs. Equilibrium Analysis

Before preceding to the analysis, it is important to discuss the interpretation of the results presented here. The analysis to be presented below should be considered as a *worst-case analysis*. For example, Theorem 3 gives a necessary and sufficient condition for the VCG mechanism to be vulnerable to shills if Ann knows or guesses the valuation of her aggregate opponent. We do not actually assume that Ann has such knowledge; the situation in which Ann has such knowledge is the worst case. This is to be contrasted with the traditional analysis of the VCG mechanism in which shills are implicitly assumed to be unavailable. In this case, truthful bidding is a dominant strategy, so even in the extreme case in which Ann knows the aggregate bid submitted by her opponents, she has no incentive to do anything other than bid truthfully. Traditionally, the term *strategy-proof* is used in the context in which it is assumed that a single agent cannot enter a mechanism under the guise of multiple identities, and the terms *shill-proof* or *false-name-proof* have been introduced to account for manipulation when entry under multiple identities is possible; however, I think that a better interpretation is that the VCG mechanism ceases to be strategy-proof given this new possibility. Even when I study *optimal shill bidding*—this should be considered as a worst case analysis; how would Ann optimally manipulate the VCG mechanism in the extreme case in which she knows the aggregate valuation of her opponent? The analysis of optimal shill bidding is closely related to the analysis of the exact conditions under which there is an incentive to use shills.

A further question concerns *equilibrium analysis*. Just because Ann has the incentive to use shills under extreme assumptions about her knowledge, does not mean that Ann will actually use shills given more reasonable assumptions. For various assumptions about the distribution of valuations, one may ask whether Ann would actually use shills in equilibrium. This appears to be a difficult question, and while I have little to say about it, I will discuss it briefly in the conclusion.

3 Optimal Shill Bidding

The main contribution of this paper is to study the problem of *optimal* shill bidding. Assume that Ann *knows* or *correctly guesses* the aggregate bid that she is facing. In this case, what is Ann's optimal use of shills? As explained above, the assumption that Ann knows the aggregate bid that she faces is extreme. In reality, Ann would typically have to decide on her shill bidding strategy in the face of uncertainty about the bids of her opponents. Nevertheless, studying Ann's optimization problem in this case sheds insight on the general problem of shill bidding.

In what follows, I decompose the problem of optimal skill bidding into two subproblems. The first subproblem concerns the lowest price at which Ann can buy any package using skills. The second subproblem concerns which package Ann would like to buy given the prices established by the first problem. I will now explore these two subproblems in turn.

3.1 Skill Bidding Cost Minimization Problem

The **Skill Bidding Cost Minimization Problem (CMP)** is as follows:

Input A valuation v_b for Bob and a subset B of N .

Output A skill bid profile $(v_j : j \in J)$ that wins B but makes the smallest possible payment among all skill bid profiles that win B .

So the CMP deals with the following scenario: Ann knows v_b , and she would like to win package B . What is the cheapest way for Ann to win B using skills? The following theorem presents an optimization problem whose solution provides a solution to the CMP.

Theorem 1 *Fix bid v_b for Bob. Suppose \mathcal{P}^* solves:*

$$\min\left\{\sum_{P \in \mathcal{P}} [v_b((N \setminus B) \cup P) - v_b(N \setminus B)] : \mathcal{P} \in \Pi(B)\right\} \quad (1)$$

Then $(v^{P, r_P} : P \in \mathcal{P}^)$ is the cheapest way Ann can win B using skills provided that the numbers $(r_P : P \in \mathcal{P}^*)$ are chosen large enough so that:*

1. *Bidder P wins package P (in original economy).*
2. *Bidder P would still win package P in marginal economy excluding any other skill bidder P' .⁷*

Ann's payment when she uses skills optimally to purchase B is given by the value of (1).

A proof of this theorem as well as a discussion occur in Section 1. Here I discuss a few points. The optimization problem 1 has as its feasible set, the set of all partitions \mathcal{P} of B . The objective is to find such a partition which minimizes a certain objective function.

⁷Formally conditions 1 may be written as follows:

$$\forall \mathcal{Q} \subseteq \mathcal{P}^*, \quad \sum_{P \in \mathcal{Q}} r_P > v_b([\bigcup \mathcal{Q}] \cup [N \setminus B]) - v_b(N \setminus B)$$

Condition 2 may be written as:

$$\forall P' \in \mathcal{P}^*, \forall \mathcal{Q} \subseteq \mathcal{P}^* \setminus P', \quad \sum_{P \in \mathcal{Q}} r_P > v_b([\bigcup \mathcal{Q}] \cup P' \cup [N \setminus B]) - v_b(P' \cup [N \setminus B])$$

Whether the strict inequalities can be changed into weak inequalities depends on the tie-breaking rule.

The objective function is a sum of terms such that there is a term for each cell P of the partition \mathcal{P} . Each of these terms give the marginal value to Bob of partition cell P given that Bob already has all items outside of B . We would like to find the partition of B which minimizes the sum of these marginal values.

Once we find a partition \mathcal{P}^* solving (1), then there is an optimal shill bid profile for Ann which is partitional, and such that the shill bidders correspond to the cells of \mathcal{P}^* . For each cell P in \mathcal{P}^* , there is a shill bidder P who bids single-mindedly for P . Shill bidder P must bid high enough for P so that he wins P (Condition 1), but also high enough so that he would still win P if some other shill bidder P' were excluded from the auction (Condition 2).

Condition 2 does not follow from Condition 1. In particular, imagine that there are two disjoint packages P and P' which are complements for Bob. Bob's marginal value for P is zero unless Bob also wins P' . Suppose that Ann uses two shill bidders, one to bid on P and another to bid on P' . Suppose that Ann wins both P and P' through her shills. The shill bidder P' submits a very high bid for P' and shill bidder P submits a low bid for P . This low bid is sufficient to win P because without P' , Bob does not value P . However, in the marginal economy excluding P' , bidder P 's bid might not be high enough to win P because once Bob receives P' , he has a higher marginal value for P .

Notice that a sufficient condition for the profile $(r_P : P \in \mathcal{P})$ to be above all of the lower bounds given by Conditions 1 and 2 in Theorem 1 is for $r_P \geq v_b(N)$ for all $P \in \mathcal{P}$. An exact mathematical formulation for these conditions is given in footnote 7. Notice that in general, Ann *may* want to make her bids as small as possible subject to conditions 1 and 2 because Ann may be bidding above her value for certain packages, and although we are assuming that Ann knows the aggregate bid v_b that she is facing, if there is a small chance that Ann is mistaken, she may want to minimize the chances that she has to overpay for certain packages. Of course, depending on Ann's valuation and beliefs, Ann may not want to minimize her bids subject to conditions 1 and 2 taking into account a small probability of a mistake.

3.2 Shill Prices vs. VCG Prices

One way of viewing Ann's problem in the VCG mechanism—without shills—is as of the problem of choosing the optimal bundle given VCG prices, where the VCG price for a package B is given by:

$$p_B^{VCG} := v_b(N) - v_b(N \setminus B) \tag{2}$$

$v_b(N)$ is Bob's utility if Ann were absent from the economy, since in this case, it would be efficient to reward Bob all goods. $v_b(N \setminus B)$ would be Bob's utility if Ann bid sufficiently high on B to win B (and won nothing else). Thus, if Ann knows v_b , then Ann effectively

faces price p_B^{VCG} for every package B . If Ann chose to win B her utility would be:

$$v_a(B) - p_B^{VCG} = v_a(B) + v_b(N \setminus B) - v_b(N)$$

Since the term $v_b(N)$ is subtracted in Ann's utility regardless of which package she purchases, Ann would like to select the package B which maximizes:

$$v_a(B) + v_b(N \setminus B)$$

In other words, if Ann knew, v_b then Ann would be motivated to choose the package which maximizes social welfare. Of course, we usually assume that Ann does not know Bob's valuation, but by truthfully revealing her value, Ann is effectively asking the mechanism to purchase the package which maximizes her utility at the VCG prices on her behalf. This is one way of understanding the dominant strategy incentive compatibility and efficiency of the VCG mechanism.

In contrast, it follows from Theorem 1 that the price that Ann faces for a package when she may use skills is:

$$p_B^{\text{Skill}} := \min\left\{\sum_{P \in \mathcal{P}} [v_b((N \setminus B) \cup P) - v_b(N \setminus B)] : \mathcal{P} \in \Pi(B)\right\} \quad (3)$$

This is the cheapest way that Ann may win B using skills. Observing that $\{B\}$ is a partition in $\Pi(B)$, and for the unique cell $P = B$ in $\{B\}$, $(N \setminus B) \cup P = N$, it follows that value of the objective in (3) at the feasible argument $\{B\}$ is simply equal to the VCG price (2). It follows that:

$$p_B^{\text{Skill}} \leq p_B^{VCG} \quad (4)$$

Another—essentially equivalent—way of seeing that (4) must be true is by observing that with skills, Ann simply has more strategies than without, and therefore the price she pays for any given package must be weakly less than the VCG price. It will be useful to think in terms of the skill prices p_B^{Skill} below.

3.3 Optimal Skill Bidding Problem

We are now in a position to discuss Ann's overall optimization problem. The **Optimal Skill Bidding Problem** is as follows:

Input A valuation v_b for Bob and a valuation v_a for Ann.

Output A skill bid profile $(v_j : j \in J)$ that maximizes Ann's utility in the VCG mechanism.

So we assume that Ann knows Bob's valuation v_b , and Ann has valuation v_a . What is Ann's optimal skill bidding strategy? The following theorem gives a program which answers this

question:

Theorem 2 Fix a bid v_b for Bob. For some $B^* \subseteq N$ and $\mathcal{P}^* \in \Pi(B^*)$, Ann has optimal skill bid profile: $(v^{P, r_P} : P \in \mathcal{P}^*)$ that is partitional.

- B^* solves:

$$\max\{v_a(B) - p_B^{\text{Skill}} : B \subseteq N\} \quad (5)$$

- \mathcal{P}^* solves the CMP (i.e., program (1)) for B^* .
- The profile $(r_P : P \in \mathcal{P}^*)$'s must be chosen large enough to satisfy Conditions 1 and 2 in Theorem 1 for B^* .

Some properties of optimal skill bidding are notable. In the canonical solution to the skill bidding problem identified by Theorem 2:

1. Each skill bidder bids single-mindedly for only one package.
2. Different skill bidders do not compete with one another, in the sense that they bid on disjoint collections of items.

3.4 Proof and Discussion of Theorems 1 and 2

Theorem 2 about optimal skill bidding is an immediate consequence of Theorem 1 about the skill bid cost minimization problem. Therefore, to prove both theorems, it is sufficient to prove Theorem 1. Let us think for a moment about why Theorem 1 is not obvious. If we do not know that there is a cost minimizing skill bid profile that is partitional, then it is possible that the payment of one skill bidder depends on the bids of other skill bidders. Even if we know that the cost minimizing skill bid profile is partitional, there may be a subtle dependence of the payment of one skill bidder on the bid of another, and moreover, the payment of skill bidder P may be different from $v_b((N \setminus B) \cup P) - v_b(N \setminus B)$. In particular, if Condition 2 is not satisfied, then in the marginal economy excluding some skill bidder P , Bob will receive package P . If P and P' are complements and skill bidder P' did not submit a sufficiently high bid for P' , then Bob may receive P' as well, in which case skill bidder P 's payment may be:

$$v_b((N \setminus B) \cup P) + r_{P'} - v_b((N \setminus B) \cup P \cup P') \quad (6)$$

However, if we know that Condition 2 is satisfied, then bidder P 's payment will be $v_b((N \setminus B) \cup P) - v_b(N \setminus B)$. Of course Condition 1 must also be satisfied since bidder P must actually win package P . It follows that to prove Theorem 1, it is sufficient to prove that there is always a cost minimizing skill bid profile which is partitional and which satisfies Conditions

1 and 2. In order to do this, we start with an arbitrary skill bid profile $(v_j : j \in J)$ winning B , and construct a partitional profile satisfying Conditions 1 and 2 that wins B but makes a smaller payment.

So let us consider a skill profile $(v_j : j \in J)$ that wins B against v_b . Then there must exist a partition $(P_j : j \in J) \in \Pi(B)$ such that bidder j wins package P_j .⁸ Let us tentatively set $r_{P_j} := v_j(P_j)$ and suppose that Ann submitted the partitional profile $(v^{P_j, r_{P_j}} : j \in J)$ instead of $(v_j : j \in J)$. Notice that for all packages A , $v^{P_j, r_{P_j}}(A) \leq v_j(A)$, which implies that the value of every allocation is lower under $(v^{P_j, r_{P_j}} : j \in J)$ than under $(v_j : j \in J)$. However, the value of the efficient allocation X selected by the VCG mechanism when Ann submits $(v_j : j \in J)$ is unchanged. Therefore X is still efficient when Ann submits $(v^{P_j, r_{P_j}} : j \in J)$. I will ignore the possibility that there are multiple efficient allocations as this case complicates the proof slightly but does not present any real problems.⁹ Given that X is the unique efficient allocation, the VCG mechanism will still select X , when Ann submits $(v^{P_j, r_{P_j}} : j \in J)$, and so skill bidder j will still win P_j . Let \hat{v}_j be either v_j or $v^{P_j, r_{P_j}}$. Then j 's VCG payment—both when Ann submits $(v_j : j \in J)$ and when Ann submits $(v^{P_j, r_{P_j}} : j \in J)$ —takes the form:

$$p_j = \underbrace{\max\{v_b(X_b) + \sum_{\ell \in I \setminus j} \hat{v}_\ell(X_\ell) : (X_\ell : \ell \in (J \cup b) \setminus j) \in \mathcal{X}_{-j}\}}_{(*)} - \underbrace{[v_b(N \setminus B) + \sum_{\ell \in J \setminus j} \hat{v}_\ell(P_\ell)]}_{(**)},$$

where \mathcal{X}_{-j} is the set of allocations in the marginal economy excluding j . Notice that term $(**)$ has the same value when Ann submits $(v_j : j \in J)$ as when Ann submits $(v^{P_j, r_{P_j}} : j \in J)$. However term $(*)$ is (weakly) lower when Ann submits $(v^{P_j, r_{P_j}} : j \in J)$ than when Ann submits $(v_j : j \in J)$ because as explained above, the value of every allocation is weakly lower when Ann submits $(v^{P_j, r_{P_j}} : j \in J)$. Therefore, p_j is lower when Ann submits $(v^{P_j, r_{P_j}} : j \in J)$ for all skills j , and thus Ann's total payment is lower. Now suppose that for some skill bidder k , r_{P_k} is not large enough to win P_k in marginal economies excluding skills j exactly in some set $H \subseteq J \setminus k$.¹⁰ Suppose that Ann raises r_{P_k} slightly. Of course this will not alter the VCG allocation, nor does it alter k 's payment, as in the VCG mechanism, a bidder's payment is independent of his bid conditional on the allocation. Let us consider the effect on the payments of the other skill bidders. There are two cases to consider. First suppose that $j \in H$. Raising r_{P_k} by a sufficiently small amount raises $(**)$, but does

⁸Observe that we may assume that each member j of J wins at least one item, since otherwise we could eliminate the members of J who do not win any items. This would simply eliminate possible allocations and therefore reduce the value of the optimal allocation in the marginal economy excluding any skill bidder j who actually wins some items, and thus reduce j 's payment.

⁹If there are multiple efficient allocations, then Ann may instead submit the profile $(v^{P_j, r_{P_j} + \epsilon} : j \in J)$ for small $\epsilon > 0$ so that X becomes the unique efficient allocation and the proof would proceed similarly.

¹⁰Here I mean that H is the set of skill bidders j such that r_{P_k} is not large enough for k to win P_k in *any* efficient allocation in the marginal economy excluding j .

not alter (*), because if r_{P_k} is sufficiently small, it will still be inefficient to allocate P_k to k in the marginal economy excluding k , and so the efficient allocation in this marginal economy is unchanged. It follows that in this case p_j is lowered. Next consider skill bidders $j \in J \setminus (H \cup k)$. Then raising r_{P_k} raises (*) and (**) by the same amount and therefore leaves p_j unaltered. It follows that if the components of the profile $(r_{P_j} : j \in J)$ are not initially large enough to satisfy Conditions 1 and 2, we may raise them until they do, and Ann's payment will be lowered in the process. This completes the proof.

4 Characterization of the Incentive to Shill

The optimization problem presented in Theorem 1 immediately generates an necessary and sufficient condition on the aggregate bid v_b for a collection of bidders I such that there exists a potential bidder outside of I who would have an incentive to use skills against v_b .

Definition 1 *Ann has a **profitable shill bid against** v_b if Ann has some shill bidding strategy that outperforms truthful bidding against v_b .*

Theorem 3 *The following conditions are equivalent:*

1. (**Submodularity at the Top (SubTop)**): For all $B \subseteq N$ and $\mathcal{P} \in \Pi(B)$:

$$v_b(N) - v_b(N \setminus B) \leq \sum_{P \in \mathcal{P}} [v_b((N \setminus B) \cup P) - v_b(N \setminus B)] \quad (7)$$

2. For all valuations for Ann, there is no profitable shill bid for Ann against v_b in the VCG auction for N .

Proof. SubTop is equivalent to:

$$v_b(N) - v_b(N \setminus B) \leq \min \left\{ \sum_{P \in \mathcal{P}} [v_b((N \setminus B) \cup P) - v_b(N \setminus B)] : \mathcal{P} \in \Pi(B) \right\}, \quad \forall B \subseteq N$$

which in turn, is equivalent to:

$$p_B^{VCG} \leq p_B^{\text{Shill}}, \quad \forall B \subseteq N$$

This means that the price that Ann has to pay for any package without skills is no more than the price that she has to pay with skills. (In fact, the two prices must be equal by (3)). So SubTop implies that Ann does not have a profitable shill bid.

Going in the other direction, suppose that SubTop fails. This means that there exists $B \subseteq N$ such that $p_B^{\text{Shill}} < p_B^{VCG}$. Then if Ann values B sufficiently, and values nothing outside of B , she will have an incentive to use skills. \square

The condition SubTop as well as the proof of the theorem come directly out of the cost minimization problem (1) in Theorem 1.

4.1 Comparison to a Result by Lehmann, Lehmann, and Nisan (2006)

I will now discuss the relation of Theorem 3 to the analysis of shill bidding in Lehmann, Lehmann, and Nisan (2006), which is the most closely related analysis in the literature.

Definition 2 A valuation v is **submodular** if for all $A \subseteq B \subseteq N$ and $x \in N \setminus B$:

$$v(B \cup x) - v(B) \leq v(A \cup x) - v(A)$$

Submodularity means that there is a decreasing marginal utility of additional goods as the set of goods already acquired increases. Therefore submodularity states that the different goods in N are substitutes for one another. Lehmann, Lehmann, and Nisan (2006) establish the following theorem:

Theorem 4 (Lehmann, Lehmann, and Nisan 2006) *If v_b is submodular, then, regardless of her valuation, there is no profitable shill bid for Ann against v_b in the VCG auction for N .*

What is the relationship between submodularity and SubTop? It is an immediate consequence of Theorems 3 and 4 that SubTop is weaker than submodularity. However, what is the precise relation between Theorems 3 and 4? The following well-known theorem—which can be found in Fujishige (2005) and Lehmann, Lehmann, and Nisan (2006), who also refer to Topkis (1998)—answers this question:

Theorem 5 *The following conditions are equivalent:*

1. v_b is submodular.
2. For all $A, B \subseteq N$ with $A \cap B = \emptyset$ and all $\mathcal{P} \in \Pi(B)$:

$$v_b(A \cup B) - v_b(A) \leq \sum_{P \in \mathcal{P}} [v_b(A \cup P) - v_b(A)]$$

In expressing the result in terms of partitions, I have phrased the result differently than in Fujishige (2005) and Lehmann, Lehmann, and Nisan (2006); the way the theorem is more commonly expressed is to say that v_b is submodular if and only if for all $A \subseteq N$, the function $v_b(\cdot|A)$ is subadditive, where $v_b(\cdot|A) : 2^{N \setminus A} \rightarrow \mathbb{R}$ is defined by $v_b(B|A) := v_b(A \cup B) - v_b(A)$. In other words, an alternative definition of submodularity of a valuation is that the marginal value of additional goods is subadditive conditional on any package already acquired. We can now directly compare SubTop and submodularity:

- **Submodularity** For all $A, B \subseteq N$ with $A \cap B = \emptyset$ and all $\mathcal{P} \in \Pi(B)$:

$$v_b(A \cup B) - v_b(A) \leq \sum_{P \in \mathcal{P}} [v_b(A \cup P) - v_b(A)] \quad (8)$$

- **SubTop** For all $B \subseteq N$ and $\mathcal{P} \in \Pi(B)$:

$$v_b(N) - v_b(N \setminus B) \leq \sum_{P \in \mathcal{P}} [v_b((N \setminus B) \cup P) - v_b(N \setminus B)]$$

Both conditions impose a set of inequalities of the form (8): the sum of marginal values of the cells of a partition of B conditional on A is (weakly) greater than the marginal value of B conditional on A . However, submodularity imposes these inequalities for all A and B such that $A \cap B = \emptyset$, whereas SubTop only imposes the subset of these inequalities corresponding to situations where $A = N \setminus B$. In other words, SupTop imposes the condition that the marginal value of adding the cells of a partition of B one at a time is greater than the marginal value of B only when Bob already has *all goods outside of B* , whereas submodularity imposes this condition regardless of which goods outside of B Bob already has. This shows the precise way in which submodularity is a stronger condition than SubTop. For any $D \subseteq N$, define the **VCG auction for D** as the application of the VCG mechanism for allocating all goods in D . In this case we may assume either that bidders submit valuations v_b for all packages $B \subseteq N$, but that the marginal value for goods outside of D are ignored, or that bidders submit the restriction of v_b to packages contained within D . It is now an immediate consequence of Theorems 3 and 5 that:

Theorem 6 (*The Incentive to Shill in SubAuctions*) *The following conditions are equivalent:*

1. v_b is submodular.
2. For all $D \subseteq N$, and for all valuations for Ann, there is no profitable shill bid for Ann against v_b in the VCG auction for D .

This follows because when we eliminate all goods outside of D submodularity implies that v_b satisfies SubTop with respect to the set of goods D . This theorem provides a characterization of submodularity in terms of the incentive to shill in subauctions, and therefore provides a sort of converse to the result of Lehmann, Lehmann, and Nisan (2006). However, notice that the hypothesis of this converse is that Ann does not have a profitable shill bid in the VCG auction for D , for all $D \subseteq N$, not just in the VCG auction for N . If one wants instead the exact equivalent of the statement that Ann does not have a profitable shill bid in the VCG auction for N , then one needs to look to Theorem 3. An alternative analysis of the incentive to shill was provided by Yokoo, Sakurai, and Matsubara (2004). Lehmann,

Lehmann, and Nisan (2006) discuss the relation of their result to that of Yokoo, Sakurai, and Matsubara (2004).

I conclude this section by presenting an example of a valuation that satisfies SubTop but not submodularity.

Example 2 *Suppose that $N = \{1, 2, 3\}$, and:*

$$v_b(A) := \begin{cases} 3, & \text{if } |A| \geq 2; \\ 1, & \text{if } |A| = 1; \\ 0, & \text{if } A = \emptyset. \end{cases}$$

First, we would like to verify SubTop, namely that the expression on the left hand side (LHS) is less than the expression on the right hand side (RHS) in inequality (7). If $B = N$, and $\mathcal{P} = \{\{1\}, \{2\}, \{3\}\}$, then $LHS = 3$ and $RHS = 3$. If $B = N$ and $\mathcal{P} = \{\{1, 2\}, \{3\}\}$, then $LHS = 3$ and $RHS = 4$. If $B = \{1, 2\}$ and $\mathcal{P} = \{\{1\}, \{2\}\}$, then $LHS = 2$ and $RHS = 4$. These three cases are representative of all nontrivial cases, and so it follows that v_b satisfies SubTop. So by Theorem 3, there does not exist a valuation for Ann such that she has an incentive to shill in the VCG auction for N .

In contrast, notice that $v_b(\{1\}) - v_b(\emptyset) = 1 < v_b(\{1, 2\}) - v_b(\{2\}) = 2$. So v_b is not submodular. On the other hand Theorem 6 implies that there is some subset D of N and some valuation for Ann such that Ann has a profitable shill bid in the VCG auction for D . In particular, let $D = \{1, 2\}$. Suppose that Ann has a single minded valuation for $\{1, 2\}$ that assigns a very high value to this bundle. Then Ann will have an incentive to enter the auction for D using two shills one of which bids single mindedly for the first item with a very high bid, and the other of which does the same for the second item. In this case, Ann wins both items and submits a payment of 2 rather than a payment of 3, which is what would happen under truthful bidding.

4.2 Comparison to a Result by Ausubel and Milgrom (2002)

Theorem 3 is also related to a result due to Ausubel and Milgrom (2002). For any valuation v , define the demand correspondence induced by v by:

$$D(p; v) := \operatorname{argmax}\{v(A) - \sum_{x \in A} p_x : A \subseteq N\}$$

v is a **gross substitutes** valuation if for all $p, p' \in \mathbb{R}_+^N$ with $p \leq p'$ and $A \in D(p)$, there exists a $B \in D(p')$ such that $\{x \in A : p_x = p'_x\} \subseteq B$. Gul and Stacchetti (1999) showed that the every gross substitutes valuation is submodular. We can think of the set of all valuations $V = \{v \in \mathbb{R}_+^{2^N} : v(\emptyset) = 0, \forall A, B \subseteq N, A \subseteq B \Rightarrow v(A) \leq v(B)\}$ as a subset of $\mathbb{R}_+^{2^N}$. Lehmann, Lehmann, and Nisan (2006) showed that—considered in this way—while

the set of gross substitutes valuations has zero Lebesgue measure, the set of submodular valuations has positive Lebesgue measure, so in this sense the latter is a much larger set. Let V_{GS} be the set of gross substitutes valuations.:

Theorem 7 (Ausubel and Milgrom 2002) *Suppose that $(v_1, \dots, v_n) \in V_{GS}^n$. Then assuming that all bidders other than 1 bid truthfully, then bidder 1 has no incentive to use skills. On the other hand, for any V_* with $V \supseteq V_* \supsetneq V_{GS}$, there exists a profile $(v_1, v_2, \dots, v_n) \in V_*^n$ such that bidder 1 has an incentive to use skills.*¹¹

In contrast to the characterization of the incentive to shill in terms of the aggregate valuation of Ann's opponent, the result of Ausubel and Milgrom (2002) provides a maximal domain such that if valuations are drawn from this domain, there is no incentive to shill. Ausubel and Milgrom (2006) argue that this sort of a maximal domain characterization is superior to a characterization in terms of the aggregate valuation of Ann's opponents because the a priori knowledge that one is likely to have about bidders is likely to concern the domain from which valuations are drawn rather than the aggregate valuation of the opponent.

Say that a set $U \subseteq V$ of valuations is a **maximal domain** for a set $W \subseteq V$ of valuations if (i) for all $n \in \mathbb{N}$ and all $(u_1, \dots, u_n) \in U^n$,

$$w(A) := \max\left\{\sum_{i=1}^n u_i(X_i) : \forall i, j = 1, \dots, n, i \neq j \Rightarrow X_i \cap X_j \neq \emptyset, \bigcup_{k=1}^n X_k = A\right\} \quad (9)$$

is such that $w \in W$, and (ii) for all $U_* \subseteq V$, with $U_* \supsetneq U$, there exists $n \in \mathbb{N}$ and $(u_1, \dots, u_n) \in U_*^n$, such that w defined by (9) does not belong to W . The following result explains the relationship between the characterization of the incentive to shill found in this paper and the characterization of Ausubel and Milgrom (2002).

Theorem 8 1. *The gross substitutes valuations are a maximal domain for the submodular valuations.*

2. *The gross substitutes valuations are a maximal domain for the SubTop valuations.*

Proof. In Appendix. \square

This result may be of independent interest.

5 Pure Complements

The results of the previous section show roughly that when goods are substitutes, there is no incentive to use skills, but when goods are not substitutes, there may be an incentive

¹¹Actually, Ausubel and Milgrom (2002) show something stronger, namely that every set V_* of valuations that contains either the additive valuations or the unit demand valuations and at least one valuation which violated the gross substitutes property is such that one can use valuations from V_* to construct a profile of valuations such that some bidder has an incentive to use skills. It is possible to strengthen the statement of Theorem 8 in a similar way. The proof in the appendix already accommodates such a strengthening.

to use skills. This conclusion is in line with previous research on the topic. When goods fail to be substitutes for Bob, then they may be pure complements, or else a mixture of substitutes and complements. While the literature has distinguished the case of substitutes from its negation, skill bidding in the case of pure complements has not been explored. The assumption of pure complements is at the opposite extreme from the assumption pure substitutes, and it is the case in which the incentive to skill bid takes its purest form. In this section, I will examine this case.

Definition 3 A *singleton valuation* is a valuation for a single item $x \in N$. In other words, a singleton valuation is a valuation of the form $v^{\{x\},r}$ for some $x \in N$ and $r \in \mathbb{R}_+$.

For notational simplicity, I will write $v^{x,r}$ instead of $v^{\{x\},r}$. The following example shows that even when Ann has a profitable skill bidding strategy, Ann may not have a profitable skill bid strategy in which all of her skills submit singleton bids. In fact, even when skill bidding is worthwhile, the best skill bid strategy in which uses only singleton bids may be worse than truthful bidding.

Example 3 Let $N = \{1, 2, 3\}$. Suppose that Bob's valuation is given by:

	1	2	3	12	13	23	123
v_b	3	3	1	4	4	4	6

The top row gives the package. For example, 12 represents the package containing items 1 and 2. The bottom row represents the value of each package. For example, the value of package 12 is 3. Notice that Bob's valuation involves a mix of substitutes and complements. In particular the marginal value of item 2 is lower if Bob already has item 1 than if Bob has nothing, but the marginal value of item 3 is higher if Bob already has item 1 than if Bob has nothing. Assume that Ann's true valuation is single minded for the package 123, and we assume that Ann's value for 123 is sufficiently large that she would win 123 if she reported truthfully, and hence in any skill bid which outperforms truthful bidding, Ann wins 123. Then Ann's utility if she reports truthfully is:

$$v_a(123) - 6.$$

Now suppose that Ann sponsors two skills, one of which submits bid v^{12,z^*} and the other of which bids v^{3,z^*} , where z^* is a very large number. Then again Ann will win the package 123 and her utility will be:

$$v_a(123) - 4 - 1 = v_a(123) - 5,$$

where 4 is the opportunity cost of bundle 12 and 1 is the opportunity cost of item 3. So in this case, skill bidding is better than truth-telling for Ann. Next suppose that Ann submits

three skill bids, v^{1,z^*} , v^{2,z^*} , and v^{3,z^*} . Then here utility will be:

$$v_a(123) - 3 - 3 - 1 = v_a(123) - 7.$$

This is worse than truth-telling. If it were possible for Ann to have a profitable skill bid in which all of her skills bid singleton valuations, then three distinct skill bidders for Ann must win three distinct items, in which case—if they place very high bids on each of these items—Ann’s payment will be 7. Given that Ann would still win the items, Theorem 1 shows that Ann would not do better by having her skills submit lower bids. So this example shows that it is possible for skill bidding to be profitable even when there is no profitable skill bidding strategy in which all skills submit singleton valuations.¹²

Recall from above that for any $A \subseteq B$, we define:

$$v_b(A|N \setminus B) = v_b((N \setminus B) \cup A) - v_b(N \setminus B)$$

$v_b(A|N \setminus B)$ is the marginal value of A given that Bob already has everything outside of B . $v_b(\cdot|N \setminus B)$ is **superadditive** if, for all nonempty $C, D \subseteq B$:

$$C \cap D = \emptyset \Rightarrow v_b(C \cup D|N \setminus B) \geq v_b(C|N \setminus B) + v_b(D|N \setminus B)$$

Say that v_b is **strictly superadditive** if the above inequality is strict for all nonempty pairwise disjoint C and D contained in N . Superadditivity represents a notion of complements. The whole package $C \cup D$ is worth more than the sum of its parts C and D . The following theorem characterizes the solution to the CMP when $v_b(\cdot|N \setminus B)$ is superadditive.

¹²Consider the following condition:

$$\forall B \subseteq N \quad v_b(N) - v_b(N \setminus B) \leq \sum_{x \in B} [v_b((N \setminus B) \cup x) - v_b(N \setminus B)] \quad (10)$$

Clearly (10) is implied by SubTop. An example similar to the one just presented shows that (10) does not imply SubTop. In particular, suppose that v_b is given by:

	1	2	3	12	13	23	123
v_b	3	3	1	4	6	6	6

If $B = N$, then the left hand side (LHS) of (10) is equal to 6 and the right hand side (RHS) is equal to 7. If $B = \{1, 2\}$, the LHS equals 5 and the RHS equals 10. If $B = \{1, 3\}$ or $B = \{2, 3\}$, the LHS equals 3 and the RHS equals 4. For all other sets $B \subseteq N$, (10) is satisfied trivially. On the other hand:

$$v_b(N) - v_b(\emptyset) = 6 > 5 = [v_b(\{1, 2\}) - v_b(\emptyset)] + [v_b(\{3\}) - v_b(\emptyset)]$$

So v_b does not satisfy SubTop.

Theorem 9 Assume $v_b(\cdot|N \setminus B)$ superadditive. If $(r_x : x \in B) \in \mathbb{R}_+^B$ is sufficiently large:

1. Then $(v^{x,r_x} : x \in B)$ solves the CMP for B .
2. The value of the CMP for B is given by:

$$p_B^{\text{Skill}} = \sum_{x \in B} [v_b((N \setminus B) \cup x) - v_b(N \setminus B)]$$

If $v_b(\cdot|N \setminus B)$ is strictly superadditive then all solutions to the CMP are of the form $(v^{x,r_x} : x \in B)$ —given that redundant skill bidders that do not win any items are eliminated.

Proof. Superadditivity of $v_b(\cdot|N \setminus B)$ implies that for all $\mathcal{P} \in \Pi(B)$:

$$\sum_{x \in B} v_b((N \setminus B) \cup x) - v_b(N \setminus B) \leq \sum_{P \in \mathcal{P}} v_b((N \setminus B) \cup P) - v_b(N \setminus B)$$

This, together with Theorem 1 implies the result. \square

This theorem says that in the case of pure complements—represented as superadditivity of $v_b(\cdot|N \setminus B)$ —the cheapest way for Ann to win a given package B is to sponsor one skill per item x in B , and have that skill place a high bid for the single item x . In other words, Ann has an incentive to totally disintegrate into one identity per item. The case of pure complements demonstrates the incentive to skill in its purest form. In contrast to the general case, we are able to display not only a program for the optimization problem, but the solution to the optimization problem as well.

Definition 4 A valuation v is **supermodular** if for all $A \subseteq B \subseteq N$ and $x \notin B$:

$$v(B \cup x) - v(B) \geq v(A \cup x) - v(A)$$

v is **strictly supermodular** if all of the above inequalities are strict.

Supermodularity is the dual of submodularity which was defined above. A valuation is supermodular exactly if it exhibits increasing marginal utility of additional goods as the set of goods already acquired increases. Thus, supermodularity represents a notion of complements.

Given that we have solved the CMP (Theorem 9), it is straightforward to characterize the optimal skill bidding program in the case of complements as well:

Theorem 10 Assume that v_b is supermodular. Then there exists $B^* \subseteq N$ such that for sufficiently large profile $(r_x : x \in B)$: $(v^{x,r_x} : x \in B)$ is an optimal skill-bidding strategy. If $v_b(\cdot|N \setminus B^*)$ is strictly superadditive, then skill bidding is strictly better than truthful bidding.

Proof. This theorem is an immediate consequence of Theorem 9 and the fact that superadditivity of $v_b(\cdot|N \setminus B)$ for all $B \subseteq N$ is equivalent to the supermodularity of v_b . This fact is simply the dual of Theorem 5. \square

Observe that *strict* supermodularity of v_b is equivalent to strict superadditivity of $v_b(\cdot|B \setminus N)$ for all $B \subseteq N$, and therefore—by Theorem 10—strict supermodularity is a sufficient condition for skill bidding to strictly outperform truthful bidding.

6 Mixture of Substitutes and Complements: Computational Complexity

In the previous section, we presented a solution to the CMP when goods are pure complements. We saw in Section 4 that in the case of pure substitutes, it is optimal not to use skills at all. What can be said about the intermediate case when there is a mix of substitutes and complements? This section discusses evidence that the general problem of optimal skill bidding is computationally intractable.¹³ In the course of doing this, I will provide an interesting economic interpretation of the problem of optimal skill bidding. In particular, I will show that the CMP is equivalent to the *winner determination problem*, which is the problem of finding an efficient allocation in a combinatorial auction.

6.1 The Winner Determination Problem

Consider a profile $(v^{S_i, r_i} : i \in I)$ of single-minded valuations. As above v^{S_i, r_i} is the single minded valuation for package S_i at value r_i . We refer to v^{S_i, r_i} as *the valuation for bidder i* . Now consider the problem of finding an efficient allocation given valuation profile $(v^{S_i, r_i} : i \in I)$. This problem is known as the *winner determination problem for single-minded bidders* (WDSMB). For simplicity, let us assume that $i \neq j \Rightarrow S_i \neq S_j$, since if there are multiple bids for the same package, we can easily eliminate all but the highest bid. We can formulate WDSMB as an integer program:

$$\max \sum_{i \in I} r_i x(S_i) \tag{11}$$

$$\forall j \in N \quad \sum_{S_i: i \in I, S_i \ni j} x(S_i) \leq 1 \tag{12}$$

$$\forall i \in I \quad x(S_i) \in \{0, 1\} \tag{13}$$

$x(S_i) = 1$ means that package S_i is allocated to bidder i , and $x(S_i) = 0$ means that S_i is not allocated to bidder i . The constraints of the form (12) say that each item j can be

¹³That, in general, the problem of optimal skill bidding is NP-hard follows from the work of Sanghvi and Parkes (2004). Conitzer and Sandholm (2006) is also closely related. Section 6.3 discusses this work.

allocated to at most one bidder. Lehmann, O’Callaghan, and Shoham (2002) has shown that WDSMB is NP-hard.¹⁴

6.2 Reduction from The Winner Determination Problem

We now examine the CMP. The general strategy of this section will be to embed WDSMB into a special case of the CMP. This will show on the one hand, that solving the CMP implicitly involves solving an efficient allocation problem, and on the other hand, that the CMP is computationally hard.

We begin by restricting attention to the special case of the CMP in which $B = N$. Then since $N \setminus B = \emptyset$ and $v_b(\emptyset) = 0$, the CMP becomes:

$$\min\left\{\sum_{P \in \mathcal{P}} v_b(P) : \mathcal{P} \in \Pi(N)\right\}$$

In other words, we would like to find a partition of N which minimizes the sum of values of the partition cells is minimized. This may be represented as an integer program:¹⁵

$$\min \sum_{S \subseteq N} v_b(S)x(S) \tag{14}$$

$$\forall j \in N \quad \sum_{S \ni j} x(S) = 1 \tag{15}$$

$$\forall S \subseteq N \quad x(S) \in \{0, 1\} \tag{16}$$

Notice that this program has exponentially many variables (that is, exponential in $|N|$). Let us consider a subproblem which allows us to express the problem with a smaller number of variables. First restrict attention to valuations v_b such that:

$$\forall S \subseteq N, |S| - 1 \leq v_b(S) \leq |S| \tag{17}$$

Now consider the valuation \hat{v} such that for all $S \subseteq N$, $\hat{v}(S) = |S|$ for all $S \subseteq N$. \hat{v} is the additive valuation such that $\hat{v}(x) = 1$ for all $x \in N$. Of course \hat{v} satisfies (17). Now we will express each valuation v_b in terms of its deviation from \hat{v} . That is to say, suppose that for any valuation v_b satisfying (17), we define the function $\hat{v}_b : 2^X \rightarrow [0, 1]$ by the equation

$$v_b(S) = \hat{v}(S) - \hat{v}_b(S)$$

Since \hat{v}_b and v_b determine one another, we can take \hat{v}_b as our input instead of v_b . Notice

¹⁴See also Blumrosen and Nisan (2007).

¹⁵Recall that v_b represents the aggregate bid of Ann’s opponent. However for the purposes of analyzing the complexity of the CMP, it may be more conceptually straightforward to think in terms of the special case in which Ann has just one opponent with valuation v_b . Of course, if this special case is complex, so is the problem in general.

that since $\widehat{v}_b(S) \in [0, 1]$ for all S , it follows that $v_b = \widehat{v} - \widehat{v}_b$ is monotone and any valuation satisfying (17) can be generated in this way. One way of specifying \widehat{v}_b is by a profile $(S_i, r_i)_{i \in I}$ where for all $i \in I, S_i \subseteq N, r_i \in [0, 1]$, and for all $i, j \in I, i \neq j \Rightarrow S_i \neq S_j$. Then suppose that $(S_i, r_i)_{i \in I}$ corresponds to the function \widehat{v}_b defined by $\widehat{v}_b(S_i) = r_i$ for all $i \in I$ and $\widehat{v}_b(S) = 0$ otherwise. Then for $v_b = \widehat{v} - \widehat{v}_b$, and $(x(S) : S \subseteq N)$ satisfying (15):

$$\sum_{S \subseteq N} v_b(S)x(S) = \sum_{S \subseteq N} (\widehat{v}(S) - \widehat{v}_b(S))x(S) = \sum_{S \subseteq N} |S|x(S) - \sum_{i=1}^n r_i x(S_i) = |N| - \sum_{i=1}^n r_i x(S_i)$$

Given this observation we may rewrite (14)-(16) on input $(S_i, r_i)_{i \in I}$ as (11)-(13) where we may change the equality in (15) to an inequality in (12) because if $T := \bigcup\{S_i : i \in I, x(S_i) = 1\} \neq N$, then either (i) $N \setminus T = S_j$ for some $j \in I$ or (ii) $N \setminus T \neq S_j$ for all $j \in I$. In case (i), it follows that we are not at an optimal solution under the relaxed constraint (12) (assuming $r_i > 0$ for all $i \in I$, which is without loss). In case (ii), we can set the variable $x(N \setminus T)$ which does not occur in the objective (11) equal to 1.

So we have just re-interpreted the CMP as the WDSMB. More precisely, when we assume (17) and represent valuations by their deviation from the additive valuation \widehat{v} , the CMP (for $B = N$) is equivalent to the WDSMB. As mentioned above, WDSMB is NP-hard. It follows from this simple translation given here that any algorithm that solves CMP solves WDSMB, and so:

Theorem 11 *The CMP is NP-hard.*

Here it is understood that we restrict attention to the CMP for all goods N , that we assume (17), and we represent valuations in terms of their deviation from the additive valuation $\widehat{v}(S) = |S|$. Of course, lifting these restrictions would not undo the complexity result. It is interesting that there is such a close relationship between the problem of skill bidding and the problem of efficient allocation.

The interpretation of this result presents some interesting issues. On the one hand, we have just presented evidence that it is difficult for Ann to find the optimal skill bidding strategy. This may be thought of as evidence that the problem of skill bidding may not be as severe as one might otherwise think, because Ann may fail to find a profitable skill bid even when one exists. On the other hand, it is also difficult to find an efficient allocation in a combinatorial auction, and therefore to implement the outcome of the VCG mechanism. So this is an additional problem for the VCG mechanism. Is the VCG mechanism better or worse off for being both hard to manipulate and hard to implement?¹⁶

¹⁶For a related discussion, see Sanghvi and Parkes (2004).

6.3 Relation to a Results by Sanghvi and Parkes (2004) and Conitzer and Sandholm (2006)

Sanghvi and Parkes (2004) studied the following decision problem, which they called the **false-name manipulation problem**: given a profile of bids for all agents other than some bidder i (who plays the role of Ann in this paper), does i have a strategy using shills which gives i a utility of ϵ more than truthful bidding? Sanghvi and Parkes (2004) showed that the false-name manipulation problem is NP-hard, by reduction from EXACT-COVER-BY-THREE-SETS, a known NP-hard problem. It is an immediate consequence of this result that the problem of optimal shill bidding is also NP-hard, as Sanghvi and Parkes (2004) point out. The contribution of Section 6 relative to Sanghvi and Parkes (2004) is the elucidation of the relationship between the problem of optimal shill bidding and the problem of efficient allocation within a combinatorial auction. This relationship is economically interesting. The new complexity proof presented in this paper by reduction to the winner determination problem for single minded-bidders established a simple translation between the problems of shill bidding and efficient allocation, which shows how the winner determination problem can be embedded in the optimal shill bidding problem. A further contribution with regard to the understanding of the complexity of optimal shill bidding comes from the context set by Section 5, which showed that with pure complements, the problem of optimal shill bidding has a simple solution. Of course, it is well known that there is also a simple solution with substitutes: simply *tell the truth*. Thus the analysis of this paper shows that it is the *mixture* of substitutes and complements which makes the problem complex.

Conitzer and Sandholm (2006) discuss the complexity of collusive strategies in the VCG mechanism. They focus on the case in which all bidders are single-minded, and they ask the question of when a cartel can win all goods and make a payment of zero. They show that answering the question of whether a collusive strategy achieving this exists is NP-complete. Comparing their result with the discussion in this section, the most important difference is as follows. Conitzer and Sandholm (2006) assume that there is a fixed cartel with a fixed number of bidders, and what makes their question hard is the question of whether there are enough members in the cartel to win all items for free. If one always has access to an unlimited number of shills, as in this paper, then this question becomes computationally trivial: it is possible to win all items for no payment against a collection of single-minded bidders if and only if there are no bids on individual items (Conitzer 2008). More generally, allowing for valuations which are not single-minded, it is possible to win all items for free if and only if the aggregate bid of the opponent assigns a value of zero to each package containing only a single item. To summarize, the analysis of Conitzer and Sandholm (2006) shows that the optimal shill bidding strategy is NP-complete if there is a fixed bound on the number of shills one may use, but it has no consequences for the complexity of finding

the optimal shill bidding strategy without such a bound. This paper—as well as the analysis of Sanghvi and Parkes (2004)—did not assume such a bound.

6.4 Interim Summary: Substitutes, Complements, and a Mixture of the Two

The following table summarizes our knowledge of shill bidding in the VCG mechanism under different assumptions on the aggregate bid:

Substitutes	Truthful Bidding	1 Identity
Complements	Incentive to Disintegrate	1 Identity per Item Won
Mix of Complements And Substitutes	Partial Incentive to Disintegrate	1 Identity per Package

In the case of substitutes, it is optimal not to shill bid at all, but rather to bid truthfully using only one identity. In the case of complements, it is optimal for Ann to totally disintegrate into one bidder per item. When there is a mix of complements and substitutes, Ann has a partial incentive to disintegrate, but typically this will fall short of one bidder per item; rather Ann will split up the bundle for which she bids into multiple packages and split up into one bidder per package. While the previous literature has separated the case of pure substitutes from its negation, this is the first paper to study the case of pure complements. This case is interesting because it is the case in which the incentive to use shills is at its purest.

7 Shill Bidding and Dominance

In this section, I explore the issue of dominant strategies in the presence of shill bidding. In order to discuss this issue, it is necessary to have some means of dealing with situations in which there are multiple efficient allocations. It is well known that if Ann bids truthfully, then her utility does not depend on which efficient allocation the VCG mechanism selects. However, if Ann does not bid truthfully, or uses shills, then when there are multiple efficient allocations, Ann’s utility may be different depending on which tie-breaking rule the VCG mechanism uses. One way to proceed would be to fix a given tie-breaking rule and analyze dominance given this rule. However, this is somewhat artificial and complicates the analysis. Therefore, I will proceed in another way. Let $\bar{v} := (v_j; j \in J)$ be a profile of shill bids for Ann, let v_b be a bid for Bob, and let v_a be Ann’s true valuation. Let $\bar{U}_a(\bar{v}, v_b; v_a)$ be the highest utility that Ann could receive given any efficient allocation—based, of course, on reported valuations—when she uses shill bid profile \bar{v} , Bob’s bid is v_b and Ann’s true valuation is v_a . Likewise, let $\underline{U}_a(\bar{v}, v_b; v_a)$ be the lowest utility that Ann could receive given any

efficient allocation in these circumstances. Whenever there is a unique efficient allocation given the reports, $\bar{U}_a(\bar{v}, v_b; v_a) = \underline{U}_a(\bar{v}, v_b; v_a)$. Also, as explained above $\bar{U}_a(v_a, v_b; v_a) = \underline{U}_a(v_a, v_b; v_a)$.

Definition 5 1. *Shill bid profile* $\bar{v} = (v_j : j \in J)$ **dominates** $\bar{v}' = (v'_j : j \in J')$ **for** v_a if for all v_b , $\bar{U}_a(\bar{v}, v_b; v_a) \geq \bar{U}_a(\bar{v}', v_b; v_a)$ and $\underline{U}_a(\bar{v}, v_b; v_a) \geq \underline{U}_a(\bar{v}', v_b; v_a)$, and (ii) there exists at least one v_b such that either $\underline{U}_a(\bar{v}, v_b; v_a) > \underline{U}_a(\bar{v}', v_b; v_a)$ or $\bar{U}_a(\bar{v}, v_b; v_a) > \bar{U}_a(\bar{v}', v_b; v_a)$.

2. *Sill bid profile* $\bar{v} = (v_j : j \in J)$ is **dominant for** v_a if for all v_b , $\bar{U}_a(\bar{v}, v_b; v_a) \geq \bar{U}_a(\bar{v}', v_b; v_a)$ and $\underline{U}_a(\bar{v}, v_b; v_a) \geq \underline{U}_a(\bar{v}', v_b; v_a)$.

Notice that unlike the standard definition, the definition of a dominant strategy allows for the possibility that there is another strategy that always performs equally well. The reason for this is that one can always add shills that value all packages at 0 without affecting one's payoff. Therefore a notion of dominance which implies that for a strategy to be dominant, it must be always at least as good and sometimes better than any other strategy is trivially too demanding in this setting.

For any index set I , let $\mathbb{R}_{++}^I := \{r \in \mathbb{R}^I : \forall i \in I, r_i > 0\}$. By a **non-shill bidding strategy**, I mean a strategy under which Ann submits a bid under a single identity.

Theorem 12 *Consider* $B \subseteq N$ *with* $v_a(B) > 0$. *For any* $\mathcal{P} \in \Pi(B)$ *with* $|\mathcal{P}| \geq 2$, *and* $(z_P)_{P \in \mathcal{P}} \in \mathbb{R}_{++}^{\mathcal{P}}$, *there does not exist any non-shill bidding strategy that dominates* $\bar{v} := (v^{P, z_P})_{P \in \mathcal{P}}$ *for* v_a . *In particular, truthful bidding does not dominate* \bar{v} *for* v_a .

Proof. Since truthful bidding dominates any non-shill bidding strategy, it is sufficient to show that truthful bidding does not dominate \bar{v} . In particular, suppose that Bob is comprised of two bidders, Carol and Dan. Carol values each item x in $N \setminus B$ at a value greater than $v_a(N)$, and she assigns a value of zero to all items in N . Carol's valuation is additive.¹⁷ Dan's valuation is $v^{B, r}$, where $0 < r < \min\{\sum_{P \in \mathcal{P}} r_P, v_a(B)\}$. If Ann bids truthfully, she will win package B and pay r .¹⁸ In contrast, if Ann submits shill profile $(v^{P, r_P} : P \in \mathcal{P})$, she will win B and make a zero payment. \square

Without shills, it is well known that truthful bidding is a dominant strategy. Theorem 12 presents a large class of strategies which are sometimes better than truthful bidding in the VCG mechanism.

Theorem 13 *There does not exist a shill bidding strategy that dominates truthful bidding.*

Proof. In Appendix.

¹⁷In other words, Carol values any package at the sum of the values of the items in the package.

¹⁸More precisely, Ann will win a package C such that $v_a(C) = v_b(B)$ and pay r . It is possible that $B \neq C$ if there exists $x \in B$ such that $v_a(B \setminus x) = v_a(B)$.

Corollary 1 *When shill bidding is possible, Ann does not have a dominant strategy.*

Proof. During the course of the proof of Theorem 12, we construct situations in which shill bidding strictly outperforms truthful bidding. Since truthful bidding dominates any other strategy not involving shills, it follows that if there is a dominant strategy it must involve the use of shills. However, Theorem 13 shows that for any strategy using shills is either sometimes worse than truthful bidding, or always leads to the same utility as truthful bidding; in neither case can such a strategy be dominant. \square

8 The Riskiness of Shill Bidding

In this section, I explore the riskiness of shill bidding. Notice that the canonical optimal shill bidding strategies found in Theorems 1 and 2 may often involve bids which are above Ann’s valuation for certain packages. Such strategies may be risky in the sense that if Ann is wrong about Bob’s bid, then she may end up paying more for a package than it is worth to her. In a setting where Ann is uncertain about Bob’s bid, she may not want to use such a strategy. Given any shill bid profile $\bar{v} = (v_j : j \in J)$, $\min_{\hat{v}_b} \underline{U}_a(\bar{v}, \hat{v}_b; v_a)$ is the lowest utility that Ann can receive if she submits shill bid profile \bar{v} when her true valuation is v_a . Observe that $\min_{\hat{v}_b} \underline{U}_a(v_a, \hat{v}_b; v_a) = 0$. In other words, if Ann bids truthfully, the lowest utility that Ann can get for any bid that Bob might submit is zero. Because it is always possible that Ann will be outbid on all items for which she bids, it must be the case that $\min_{\hat{v}_b} \underline{U}_a(\bar{v}, \hat{v}_b; v_a) \leq 0$ for all \bar{v} and v_a . One may regard a shill bidding strategy \bar{v} as risky for Ann when her true value is v_a if $\min_{\hat{v}_b} \underline{U}_a(\bar{v}, \hat{v}_b; v_a) < 0$. In other words, in the worst case, the strategy delivers a strictly lower utility to Ann than does truthful bidding.

In general, there are two motives for shill bidding:

1. Sponsoring shills may reduce the prices that a bidder would have paid on items that they would have won without shills.
2. Sponsoring shills may allow a bidder to win items that the bidder would not otherwise have won at acceptable prices.

Fix a bid v_b for Bob, and suppose that Ann chooses a shill bidding strategy that is a best reply to v_b . It is easy to construct examples in which an optimal shill bidding strategy \bar{v} against v_b does not alter the allocation and but merely reduces prices achieving 1 in a way which is not “risky” in the sense described above, i.e., $\min_{\hat{v}_b} \underline{U}_a(\bar{v}, \hat{v}_b; v_a) = 0$. What is more surprising is that it is possible to construct situations in which there is an optimal shill bidding strategy \bar{v} that achieves 2 as well, attaining for Ann a package which is more valuable than the one she would have attained under truthful bidding but which is not risky, i.e., $\underline{U}_a(\bar{v}, \hat{v}_b; v_a) = 0$. This is shown by the following example.

Example 4 $N = \{1, 2, 3\}$

$$v_a(A) = \begin{cases} 5, & \text{if } \{1, 2\} \subseteq A; \\ 4, & \text{if } 3 \in A, \{1, 2\} \not\subseteq A; \\ 2, & \text{if } 3 \notin A, |A| = 1. \\ 0, & \text{otherwise.} \end{cases}$$

$$v_b(A) = \begin{cases} 5, & \text{if } A \subseteq \{1, 2\}; \\ 3, & \text{if } 3 \in A, \{1, 2\} \not\subseteq A; \\ 0, & \text{otherwise;} \end{cases}$$

Then the efficient allocation gives package $\{1, 2\}$ to Bob and item 3 to Ann for a social utility of 9. Ann pays a VCG price of 0 for a net utility of 4. Bob's VCG payment is 1 and Bob receives a net utility of 4. So the total utility to the bidders is 8 and the payment to the seller is 1. Suppose that instead Ann entered the auction under two skills, and bid for items 1 and 2 with single-minded bids of 2 for each. Then Ann would win $\{1, 2\}$ and Bob would win item 3 for a social utility of 8, which is lower. Ann's payment is then 0, and her utility is therefore 5. (Of course, this must be Ann's optimal skill bidding strategy). Bob's payment is 0 and his net utility is 3. The seller receives nothing. Notice that in this situation, using skill bidders upsets the efficient allocation and even delivers a package to Ann which she values more than the package that she would receive under truthful bidding, Ann's worst case utility is 0, so from a minimax perspective, she is no worse off than she would be under truthful bidding.

It would be interesting to study the optimal skill bidding strategy subject to the constraint that Ann does not risk overpaying for any package, or in other words, subject to the constraint that $\min_{\hat{v}_b} \underline{U}_a(\bar{v}, \hat{v}_b; v_a) = 0$. As we have just seen, sometimes this constraint would not be binding.

9 Collusion

This section uses the results about skill bidding derived above to draw some conclusions about collusion. Consider some coalition J with valuations $(v_j : j \in J)$. Let v_J be the aggregate valuation for J . Let us assume away any internal problems of enforcing the collusive arrangements, so that J can efficiently collude. In the VCG mechanism, J bids against Bob, who—as above—is assumed to be J 's aggregate opponent. Throughout this section, I fix Bob's aggregate valuation v_b , and assume, for expositional simplicity, that Bob bids truthfully—or more precisely, that the bidders comprising Bob bid truthfully. I also assume, for expositional simplicity, that there is a unique efficient allocation. Let B_j^* be the package assigned to bidder j at the efficient allocation, and let $B^* = \bigcup \{B_j^* : j \in J\}$

be the set of items assigned to members of J collectively at the efficient allocation.

In what follows it will be useful to define the dual of SubTop:

- **Supermodularity at the Top (SupTop):** For all $B \subseteq N$ and $\mathcal{P} \in \Pi(B)$:

$$v_b(N) - v_b(N \setminus B) \geq \sum_{P \in \mathcal{P}} [v_b((N \setminus B) \cup P) - v_b(N \setminus B)] \quad (18)$$

Just as SubTop is weaker than submodularity, SupTop is weaker than supermodularity. Whereas SubTop is a substitutes property, SupTop is a complements property.

9.1 VCG prices vs. Merged Prices

For any member $j \in J$ and any $B \subseteq N$, the VCG-price for B , assuming that all others bid truthfully is:

$$p_B^{VCG,j} = v_{(J \setminus j) \cup b}(N) - v_{(J \setminus j) \cup b}(N \setminus B),$$

where $v_{(J \setminus j) \cup b}$ is the aggregate valuation for the collection consisting of Bob and all members of J other than j . Next define:

$$p_{B^*}^{VCG,J} := \sum_{j \in J} p_{B_j^*}^{VCG,j}$$

So $p_{B^*}^{VCG,J}$ is the total payment made by members of J if they all bid truthfully. Notice here that if $B_j^* = \emptyset$ then $p_{B_j^*}^{VCG,j} = 0$. Since truthful bidding is a dominant strategy, $p_{B^*}^{VCG,J}$ represents the total payment made by J if players behave noncooperatively. This is so given any assumptions about player's knowledge of others' bids.

Next, for any package $B \subseteq N$, define:

$$p_B^{\text{Merged}} := v_b(N) - v_b(N \setminus B)$$

p_B^{Merged} is the price that the entire coalition J would have to pay for B , if J merged and presented a single sufficiently high bid for B . p_B^{Merged} is also the price that the coalition J would pay for any package B that it wins if J submitted a only a single bid equal to the aggregate valuation v_J of J .

9.2 Collusion in The Second Price Auction For a Single Item

Even in the absence of any issues involving shill bidding, it is well known that the VCG mechanism is prone to collusion (Graham and Marshall 1987). For example consider an auction for a single item. Then the VCG mechanism is the second price auction. With only one item up for auction all conditions such as gross substitutes, submodularity, and SubTop which are sufficient for the absence of shill bidding are trivially satisfied, and therefore there

is no incentive to use skills. However there is still an incentive to collude. This is expressed by the fact that in a second price auction for a single item x , if it is efficient to assign x to some member of J , then:

$$p_x^{\text{Merged}} \leq p_x^{VCG,J}, \quad (19)$$

and, in fact, the inequality is sometimes strict. Assuming that all bidders outside of the coalition J submit truthful bids and that the highest value for x resides within J , $p_x^{VCG,J}$ is the second highest value for x among all bidders. This is because $p_x^{VCG,J}$ represents the price that J would have to pay for x if its members did not collude. On the other hand, p_x^{Merged} represents the highest value outside of J . When both the highest and second highest value reside within J , then the highest value outside of J is less than the second highest value among all bidders and therefore $p_x^{\text{Merged}} < p_x^{VCG,J}$. In this case there is a strict incentive to collude. In fact the following is true:

- (*) Abstracting away from internal problems of enforcing collusive behavior, in the second price auction (the VCG mechanism for auctioning a single item), it is a dominant strategy for the coalition to submit a single bid equal to its aggregate value for the item (which is equal to the valuation of the member j of J who values it most).

9.3 Competition and Integration Effects

In a second-price auction, the benefit to merging comes from suppression of competition. In the combinatorial version of the second price auction, it is always beneficial to suppress competition, but merging does more than just suppress competition. In the simple second price auction, merging is only beneficial when some member of the coalition would win the item on auction, and the effect of merging is simply to eliminate losing bids within the coalition, and therefore potentially lower the price of the winner. In a combinatorial auction, we may think of an analogous strategy. Suppose that the coalition knows the aggregate opposing bid v_b . Then all losing bidders may simply withhold their bids. Moreover, all winning bidders may withhold their bids on all packages that they do not win. However, even after this is done, it may be possible for the coalition to further reduce the prices for the packages that it wins. All winning bidders should raise their bids on the packages they win until they arrive at a state at which they would still win the the same package in all marginal economies excluding some other winning member of the coalition. This may not initially be so because of complementarities in Bob's bid v_b ; once Bob receives a new package in the marginal economy excluding some winning coalition member j , this may raise his marginal value for k 's package, and therefore Bob may win k 's package as well. If all winning bidders j in J bid $r > v_b(N)$ on the packages B_j^* that they win at the efficient allocation, this would be sufficient to win these packages at all the relevant marginal economies as well. Of course a profile of bids below $v_b(N)$ may also be sufficient. In the scenario that we have

been discussing the coalition enters bid profile $(v^{B_j^*, r} : j \in J, B_j^* \neq \emptyset)$ instead of $(v_j : j \in J)$, and then makes a total payment of:

$$p_{B^*}^{\text{Suppressed}} = \sum_{j \in J} [v_b((N \setminus B^*) \cup B_j^*) - v_b(N \setminus B^*)]$$

Theorem 14 $p_{B^*}^{\text{Suppressed}}$ is the smallest payment possible for the coalition J when facing aggregate bid v_b subject to the constraint that for all $j \in J$, bidder j wins B_j^* .

The proof of this theorem is the same as part of the proof of Theorem 1 and therefore is omitted. The superscript ‘‘Suppressed’’ refers to suppressed competition. $p_{B^*}^{\text{Suppressed}}$ represents the smallest payment the coalition can make without upsetting the efficient allocation, and without having to reallocate items among themselves after the auction. In other words, each bidder keeps whatever items he wins during the auction. So $p_{B^*}^{\text{Suppressed}}$ represents the payment of the coalition if the members of the coalition agreed to suppress competition.¹⁹

We can then decompose the effect of merging on the price that the coalition has to pay into two terms:

$$p_{B^*}^{\text{Merged}} - p_{B^*}^{\text{VCG}, J} = \underbrace{(p_{B^*}^{\text{Merged}} - p_{B^*}^{\text{Suppressed}})}_{\text{Integration Effect}} + \underbrace{(p_{B^*}^{\text{Suppressed}} - p_{B^*}^{\text{VCG}, J})}_{\text{Competition Effect}}$$

The Competition Effect represents the reduction in payment due to suppressing competition. The Integration Effect represents the change in payment due to merging once competition has already been suppressed.

Theorem 15 *The Competition Effect is always nonpositive. The Integration Effect can be positive or negative. If goods are substitutes (i.e., v_b satisfies SubTop), then the Integration Effect is nonpositive. If goods are complements (i.e., v_b satisfies SupTop) then the Integration Effect is nonnegative.*

Proof. That the Competition Effect is always nonpositive is a restatement of Theorem 14. The question of whether the Integration Effect is negative or positive is equivalent to the resolution of the following inequality:

$$v_b(N) - v_b(N \setminus B^*) \stackrel{\leq}{\geq} \sum_{j \in J} [v_b((N \setminus B^*) \cup B_j^*) - v_b(N \setminus B^*)]$$

SubTop implies that $\stackrel{\leq}{\geq}$ becomes \leq , and SupTop implies that $\stackrel{\leq}{\geq}$ becomes \geq . \square

¹⁹Note that suppression of competition among a coalition of bidders in this sense does not mean that removal of any member of the coalition would have no effect on the prices paid by others. It means only that prices for the coalition members cannot be further reduced without altering the allocation selected by the VCG mechanism.

9.4 Merging is the Inverse of Shilling

Theorem 15 shows that suppression of competition always reduces the coalition’s payment, but once competition has been suppressed, whether the members of the coalition would like to merge or not depends on whether goods are complements or substitutes for Bob. Theorem 15—as well as Theorem 16 below—formalizes an observation made by Milgrom (2004) that when goods are substitutes, the VCG mechanism often creates an incentive for mergers. The proof of Theorem 15 depends on the following simple intuition: once competition has been suppressed shill bidding is simply the inverse of collusion. If goods are substitutes for Bob, this dissuades a single bidder from sponsoring shills, and for the same reason, this encourages a coalition to merge. On the other hand if goods are complements for Bob, then this encourages a single bidder to sponsor shills—and in the case of pure complements, as was shown by Theorem 9, this encourages a single bidder to sponsor one shill per item he wins. Looked at from the standpoint of a coalition, the incentive to use shills translates into the disincentive to merge.²⁰ The following observation captures the notion that shill bidding and merging are inverses of one another.

Observation 1 *The following are equal:*

1. *Coalition J ’s utility to submitting profile $(\hat{v}_j : j \in J)$ against v_b .*
2. *Ann’s utility to submitting shill bid profile $(\hat{v}_j : j \in J)$ against v_b , when Ann’s true valuation is equal to v_J .*

In particular, J attains the same utility to submitting a single bid v_J as Ann would attain from truthful bidding if her valuation were v_J . On the other hand, by playing straightforwardly, and submitting its true profile of valuations $(v_j : j \in J)$, coalition J attains the same utility as Ann would attain through the shill bid profile $(v_j : j \in J)$.

Theorem 15 and Observation 1 have a variety of consequences. One such consequence is the generalization of (19) and (*) from the VCG mechanism for single item auctions to the VCG mechanism for combinatorial auctions, as well as an analysis of the limits of this generalization:

²⁰The relationship between the incentive to split into multiple identities and the incentive for multiple identities to merge has been studied in other contexts; for instance, Moulin (2008) studies split-proofness and merge-proofness in scheduling problems.

Theorem 16 1. If v_b satisfies *SubTop*, then for the package B^* which is assigned to J in the efficient allocation, $p_{B^*}^{\text{Merged}} \leq p_{B^*}^{\text{VCG},J}$.

2. Assume members of J know a priori that v_b satisfies *SubTop*—which would be true, for example, if members of J know that all bidders that comprise Bob submit gross substitutes valuations. Then it is a dominant strategy for coalition J to submit a single bid equal to the aggregate valuation v_J of J .

3. Suppose that v_b fails to satisfy *SubTop*. Then there exists a coalition J with some profile of valuations $(v_j : j \in J)$ such that for a package B^* which would be allocated to J at an efficient allocation, $p_{B^*}^{\text{VCG},J} < p_{B^*}^{\text{Merged}}$.

4. Assume that v_b fails to satisfy *SubTop*. Then there exists a coalition J with some profile of valuations $(v_j : j \in J)$ such that it is suboptimal for the coalition to submit a single bid equal to the aggregate valuation v_J of J .

Proof. In Appendix.

The difference between the situations in which merging is attractive and the situations in which it is unattractive is accounted for by the integration effect. When the Integration Effect is sufficiently large and of opposite sign from the Competition Effect, then merging will be worse than truthful bidding. Note again that the conditions which make collusion attractive ($p_{B^*}^{\text{Merged}} < p_{B^*}^{\text{VCG},J}$) are the same as the conditions which make the shill bid profile $(v_j : j \in J)$ unattractive to Ann. Thus *SubTop*—which prevents the VCG mechanism from being prone to shills, makes the VCG mechanism prone to certain forms of collusion. The next theorem makes this more precise. Say that a coalition of bidders J with valuations $(v_j : j \in J)$ is **minimally competitive when facing** v_b if $p_{B^*}^{\text{VCG},J} = p_{B^*}^{\text{Suppressed}}$, where B^* is the package that would be efficiently allocated to J when J faces Bob with valuation v_b . Say that v_b is **additive** if for all $B \subseteq N$, $v(B) = \sum_{x \in B} v(\{x\})$. If v_b is additive, then Bob views all goods as being independent. Say that a coalition J has an **incentive to merge against** v_b , if the coalition J could do better by submitting a single bid on behalf of all of its members than by having all of its members bid truthfully in the VCG mechanism. It is important to note that it is possible that a coalition J may not have an incentive to merge against Bob, while some sub-coalition K of J may have an incentive to merge against the aggregate of Bob's bid and the truthful bids of the players in $K \setminus J$. The following theorem considers only merger of the entire coalition J , and not of its subcoalitions as just described.

Theorem 17 *The following are equivalent:*

1. v_b is additive.
2. There exists neither (a) a bidder who has an incentive to shill against v_b , nor (b) a minimally competitive coalition that has an incentive to merge against v_b .

In order to prove this theorem, it is necessary to prove the following lemma:

Lemma 1 *The following are equivalent:*

1. v_b is additive.
2. v_b satisfies both SubTop and SupTop.

It is well known that among valuations v_b such that $v_b(\emptyset) = 0$, the additive valuations are the intersection of the submodular and supermodular valuations. Since SupTop is strictly weaker than supermodularity and SubTop is strictly weaker than submodularity, it is conceivable that the intersection of the SubTop and SupTop valuations is larger than the set of additive valuations. However, Lemma 1 shows that indeed the intersection of SubTop and SupTop valuations is exactly the set of additive valuations. The proofs of Lemma 1 and Theorem 17 are in the appendix.

Theorem 17 shows that unless v_b is additive, the VCG mechanism is prone either to shill bidding or to merging by a minimally competitive coalition of bidders. Notice that the set of additive valuations is a maximal domain for itself, so that it is also true that when all bidder valuations are additive, then the VCG mechanism is prone to neither shill bidding nor merger by a minimally competitive coalition, but if it is possible that some bidder fails to satisfy this property, then the VCG mechanism will be prone to one of these manipulations. It is also worth noting that when v_b violates both SubTop and SupTop, then the VCG mechanism may be simultaneously prone to both manipulations, depending on who Bob is facing. Unlike shill bidding, while merging reduces the seller’s revenue, it is not harmful in terms of efficiency. This issue is discussed further in Section 9.6. A key intuition which is articulated by Theorem 17 is that while complements create an incentive for disintegration, substitutes create an incentive for integration.

9.5 The Optimal Collusive Strategy

Ignoring the problem of providing incentives within the coalition and making the extreme assumption that the coalition knows the aggregate bid it is facing, the problem of finding the optimal collusive strategy is essentially the same as that of finding the optimal shill bidding strategy. For any package $B \subseteq N$, define:

$$p_B^{\text{Shill},|J|} := \min_{P \in \mathcal{P}} \left\{ \sum_{P \in \mathcal{P}} [v_b((N \setminus B) \cup P) - v_b(N \setminus B)] : \mathcal{P} \in \Pi(B), |\mathcal{P}| \leq |J| \right\}$$

$p_B^{\text{Shill},|J|}$ is the price that the coalition J would have to pay if it submitted the collection of bids which would minimize their payment for B . I assume that—unlike in the case of shill bidding—the coalition cannot manufacture identities, and therefore can submit at most $|J|$ bids. The superscript “Shill” on the price represents the fact that this price is found

by solving an optimization problem which is similar to the skill bidding CMP. The only difference is that there is now an additional constraint $|\mathcal{P}| \leq |J|$, representing the fact that the coalition cannot manufacture identities, and therefore must select a partition with no more cells than there are members of J . The validity of the formula for $p_B^{\text{Skill},|J|}$ follows from Theorem 1 and its proof. The coalition's optimal strategy can be found by solving the optimization problem determining $p_B^{\text{Skill},|J|}$, and then solving:

$$\max\{v_j(B) - p_B^{\text{Skill},|J|} : B \subseteq N\} \quad (20)$$

Observe that the constraint $|\mathcal{P}| \leq |J|$ is trivially satisfied and hence never binding when $|J| \geq |N|$, or in other words, the coalition contains at least as many members as there are goods on auction; in this case, the program for optimal collusion (20) is identical to the program for optimal skill bidding (5).

9.6 Efficiency

Theorem 18 *Fix a bid v_b for Bob. Then the following are equivalent:*

1. v_b satisfies *SubTup*.
2. Every coalition J has an optimal strategy against v_b which leads to an efficient allocation.
3. For all collections of bidders I with aggregate valuation v_b (more precisely, $v_b(B) = \max\{\sum_{i \in I} v_i(X_i) : \forall i, j \in I, i \neq j \Rightarrow X_i \cap X_j = \emptyset, \bigcup_{i \in I} X_i = B\}$), every coalition J has an optimal strategy against v_b which does not lower the utility of any member of I (assuming that bidders in I bid truthfully) in comparison to the situation in which members of J were to bid truthfully.

Proof. In Appendix. \square

In the statement of the theorem, it is possible to replace 2 by:

- Every optimal strategy of every coalition against v_b leads to an efficient allocation.

This would strengthen the theorem when 2 is viewed as a necessary condition and weaken it when 2 is viewed as a sufficient condition.

Theorem 18 is strongly related to a result of Ausubel and Milgrom (2002), which shows that for any collection of bidders with gross substitutes valuations (*) there does not exist a profitable joint deviation of losing bidders, but for any larger domain one can find a counterexample to (*). Of course, a profitable joint deviation by losing bidders will always create inefficiency. The relation of Theorem 18 to this result of Ausubel and Milgrom (2002) is partly explained by Theorem 8.

One consequence of Theorem 18 is that while violations of SupTop may lead to incentives for a coalition—even a minimally competitive coalition—to merge, this leads neither to inefficiency nor lowers the utility of bidders outside of the coalition. Of course, such behavior may be detrimental to the seller. In contrast, a failure of SubTop may lead to collusive behavior, which—like shill bidding—leads to inefficiency. Such collusive behavior will not in general require the coalition to manufacture identities, but will cause the coalition to distort its bidding behavior so as to win items at lower prices, and in so doing, to win items that the coalition would not have won otherwise.

10 Conclusion

This paper has studied the problem of optimal shill bidding in the VCG mechanism. The problem was decomposed into two parts: a cost minimization problem (Theorem 1) and an optimal shill bidding problem (Theorem 2). I identified a necessary and sufficient condition on aggregate bids such that there is an incentive for some outside bidder to use shills, namely Submodularity at the Top. In line with the previous literature, it was shown that when goods are substitutes, there is no incentive to shill. However, a new finding was also presented, namely, that when goods are pure complements, then the optimal shill bidding strategy involves total disintegration: there is an incentive to sponsor one shill per item that one expects to win. The incentive to disintegrate is attenuated when there is a mixture of substitutes and complements; disintegration may only be partial with several shill bidders bidding on several packages, but not necessarily on individual items. With a mixture of substitutes and complements, the problem becomes quite complex, and I provided a new proof by reduction from the winner determination problem with single-minded bidders that optimal shill bidding is NP-hard.²¹ This last result is interesting economically, as it establishes a strong relationship between optimal shill bidding and the problem of finding an efficient allocation in a combinatorial auction. I also showed that when shill bidding is possible, there is no longer a dominant strategy in the VCG mechanism; in particular, no strategy employing shills is dominant. Finally, the analysis of this paper had consequences for other forms of collusion.

Several interesting questions remain. In particular, the optimization problem that I study assumes that the bidder considering the use of shills knows the aggregate bid of the opponent. It would be interesting to study the optimization problem which results when the bidder instead has a probability distribution over aggregate bids. This would certainly change the character of the solution. The analysis of this paper has been of a worst case variety. It was shown that under extreme assumptions about a bidder’s knowledge, there is an incentive to shill. It remains to be seen whether this incentive would translate into

²¹Section 6.3 discusses an earlier result implying the NP-hardness of optimal shill bidding problem by Sanghvi and Parkes (2004).

skill bidding in equilibrium. Of course any equilibrium analysis involves some specification of a probability distribution over valuations of bidders. One relatively trivial observation is that there is a degenerate distribution in which some specific valuation profile is given probability one, and therefore it follows from the analysis of this paper that skill bidding can occur in equilibrium for some probability distribution over valuations.²² Putting this degenerate case aside, there is the issue of the riskiness of skill bidding. The canonical skill bidding strategies presented in Theorems 1 and 2 may often involve bids for certain packages above one's valuation for those packages. While such strategies are not dominated, they are generally risky in the sense that in using them, a bidder exposes himself to the risk that he will have to pay more for a package than it is worth to him. This risk might deter skill bidding in equilibrium. However, in Section 8, it was shown that there are circumstances in which a bidder may benefit from skill bidding even to the extent that he may win items which it is inefficient for him to win but in which he does not expose himself to the risk of overbidding. This issue along with the general analysis of equilibria in such complex auctions deserves further exploration.

11 Appendix

Proof of Theorem 8

Theorem 6 of Gul and Stacchetti (1999) shows that the aggregate valuation w formed by a collection of gross substitutes valuations is submodular. Since SubTop is weaker than submodularity, w also satisfies SubTop. Since SubTop is weaker than submodularity in order to establish both parts 1 and 2 of Theorem 8, it is sufficient to show that for any strict superset of the gross substitutes valuations, one can construct an aggregate valuation which violates SubTop. This is achieved by the following two lemmas.

Lemma 2 *Suppose that v satisfies SubTop. Then for all $A \subseteq N$, if $A, N \in D(p; v)$, then for all B with $A \subseteq B \subseteq N$, $B \in D(p; v)$.*

Proof. That $A, N \in D(p; v)$ implies that:

$$v(N) - v(A) = \sum_{x \in N \setminus A} p_x \tag{21}$$

²²For example, suppose that there are two bidders, Ann and Bob. Suppose that Bob has a strictly supermodular valuation. Then, Ann will have an incentive to use skills against Bob (if Bob bids truthfully), and indeed one skill per item that she will ultimately win, but if she does this, then the aggregate valuation of Ann's skills will be additive, and so Bob will not have an incentive to use skills, and so, it will be optimal for Bob to bid truthfully.

SubTop implies that for all P, Q such that $\{P, Q\} \in \Pi(N \setminus A)$,

$$v(N) - v(A) \leq [v(A \cup P) - v(A)] + [v(A \cup Q) - v(A)] \quad (22)$$

In particular, we may choose P so that $A \cup P = B$. (21) and (22) imply that:

$$\sum_{x \in N \setminus A} p_x \leq [v(A \cup P) - v(A)] + [v(A \cup Q) - v(A)], \quad (23)$$

Moreover, $A \in D(p; v)$ implies:

$$\sum_{x \in P} p_x \geq v(A \cup P) - v(A) \quad \text{and} \quad \sum_{x \in Q} p_x \geq v(A \cup Q) - v(A) \quad (24)$$

(23) and (24) together imply:

$$\sum_{x \in P} p_x = v(A \cup P) - v(A) \quad \text{and} \quad \sum_{x \in Q} p_x = v(A \cup Q) - v(A)$$

which in turn implies that $B = A \cup P \in D(p; v)$. \square

Lemma 3 *Suppose that v is not a gross substitutes valuation. Then there exists a gross substitutes valuation u such that:*

$$w(B) := \max\{v(B \setminus A) + u(A) : A \subseteq B\} \quad (25)$$

violates SubTop.

Proof. Suppose that v is not a gross substitutes valuation. It follows from Theorem 1 of Gul and Stacchetti (1999) that there exists $p \in \mathbb{R}_+^N$, $A, B \in D(p; v)$, and $X \subseteq A \setminus B$ such that:

(*) for all $Y \subseteq B \setminus A$, $(A \setminus X) \cup Y \notin D(p; v)$.

Define valuation u by:

$$u(C) := \sum_{x \in C \cap [N \setminus A]} p_x$$

u is an additive valuation and hence is a gross substitutes valuation. Define w via (25). Then, for all $C \subseteq N$:

$$\begin{aligned} w(C) - \sum_{x \in C} p_x &= \max\{v(E) + \sum_{x \in C \setminus E} p_x : A \cap C \subseteq E \subseteq C\} - \sum_{x \in C} p_x \\ &= \max\{v(E) - \sum_{x \in E} p_x : A \cap C \subseteq E \subseteq C\} \end{aligned} \quad (26)$$

It follows that:

$$\begin{aligned}
\max\{w(C) - \sum_{x \in C} p_x : C \subseteq N\} &= \max\{\max\{v(E) - \sum_{x \in E} p_x : A \cap C \subseteq E \subseteq C\} : C \subseteq N\} \\
&= \max\{v(E) - \sum_{x \in E} p_x : E \subseteq N\}, \tag{27}
\end{aligned}$$

where the least inequality follows from the fact that for all $E \subseteq N$, $A \cap C \subseteq E \subseteq C$ when $C = E$. Next observe that for any $C \in D(p; v)$:

$$\begin{aligned}
w(C) - \sum_{x \in C} p_x &\geq v(C) - \sum_{x \in C} p_x \\
&= \max\{v(E) - \sum_{x \in E} p_x : E \subseteq N\} \\
&= \max\{w(E) - \sum_{x \in E} p_x : E \subseteq N\}
\end{aligned}$$

So $D(p; v) \subseteq D(p; w)$. In particular $A \in D(p; w)$, which implies that:

$$w(N) - \sum_{x \in N} p_x = \max\{v(E) - \sum_{x \in E} p_x : A \subseteq E \subseteq N\} = v(A) - \sum_{x \in A} p_x = w(A) - \sum_{x \in A} p_x$$

It follows that $N \in D(p; w)$. Notice also that since $B \in D(p; v)$, $B \in D(p; w)$. Now assume for contradiction that w satisfies SubTop. Then by Lemma 2, and the fact that $B \subseteq (A \cup B) \setminus X \subseteq N$, N also satisfies SubTop, $(A \cup B) \setminus X \in D(p; w)$. But:

$$\begin{aligned}
w((A \cup B) \setminus X) - \sum_{x \in (A \cup B) \setminus X} p_x &= \max\{v(E) - \sum_{x \in E} p_x : A \setminus X \subseteq E \subseteq (A \cup B) \setminus X\} \\
&< \max\{v(E) - \sum_{x \in E} p_x : E \subseteq N\} \tag{28} \\
&= \max\{w(E) - \sum_{x \in E} p_x : E \subseteq N\},
\end{aligned}$$

where the first equality follows from (26), the inequality from (*), and the last equality from (27). (28) contradicts $(A \cup B) \setminus X \in D(p; w)$. It follows that w violates SubTop. \square

Proof of Theorem 13

Fix v_a , and assume for contradiction that there exists some shill bidding strategy $\bar{v} = (v_j : j \in J)$ that dominates truthful bidding. Let v_J be the aggregate valuation induced by \bar{v} . First suppose that there exists some $B \subseteq N$ and some $x \in B$ such that $v_J(B) - v_J(B \setminus x) < v_a(B) - v_a(B \setminus x)$. Suppose that Bob is composed of two bidders, Carol and Dan. Carol's valuation is $v^{N \setminus B, r^*}$ where $r^* > \max\{v_J(N), v_a(N)\}$. Dan's valuation is

$v^{x,r}$, where $v_J(B) - v_J(B \setminus x) < r < v_a(B) - v_a(B \setminus x)$. Then if Ann bids truthfully, her utility will be $v_a(B) - r > v_a(B \setminus x)$, whereas $v_a(B \setminus x)$ will be an upper bound on Ann's utility if she reports \bar{v} , a contradiction. Next suppose that for all $C \subseteq N$ and all $y \in C$, $v_J(C) - v_J(C \setminus y) \geq v_a(C) - v_a(C \setminus y)$, and for some $B \subseteq N$ and $x \in B$, $v_J(B) - v_J(B \setminus x) > v_a(B) - v_a(B \setminus x)$. Then suppose that Carol is as above, and that Dan's utility is $v^{x,r}$ for $v_J(B) - v_J(B \setminus x) > r > v_a(B) - v_a(B \setminus x)$. Then Bob's utility from truthful bidding is exactly $v_a(B \setminus x)$ and Bob's utility from submitting \bar{v} is exactly $v_a(B) - r < v_a(B \setminus x)$,²³ a contradiction. The only case that remains is when $v_J(B) - v_J(B \setminus x) = v_a(B) - v_a(B \setminus x)$ for all $B \subseteq N$ and all $x \in B$. Since $v_J(\emptyset) = v_a(\emptyset) = 0$, it follows that $v_J = v_a$. This implies that for any aggregate bid v_b , the set of packages that Ann could win in some efficient allocation given that she bids truthfully is the same as the set of packages that Ann could win in some efficient allocation given that she submits shill bid profile \bar{v} . The fact that \bar{v} dominates truthful bidding then implies that there must exist aggregate bid v_b for Bob and resulting efficient allocations such that Ann wins the same package B under both \bar{v} and v_a and Ann pays less under \bar{v} than under v_a . This implies that at least two of Ann's shill bidders win items in B , and therefore that B contains at least two elements. We may assume that for all $x \in B$, $v_J(B) - v_J(B \setminus x) = v_a(B) - v_a(B \setminus x) > 0$.²⁴ Suppose that for each shill bidder $j \in J$, B_j is the package that shill bidder j wins, so that $\bigcup_{j \in J} B_j = B$. Now suppose that Bob is composed of two bidders, Carol and Dan. Carol's valuation is as above. Dan's valuation v_d is such that $0 < v_d(B) < \min\{v_a(B) - v_a(B \setminus x) : x \in B\}$ and that $\sum_{j \in J} v_d(B_j) > v_d(B)$. Notice that when Ann bids against Dan and Carol with valuations as just described, and submits shill bid profile \bar{v} , it must still be efficient to allocate B_j to shill bidder j for all $j \in J$. Ann's payment will then be at least $\sum_{j \in J} v_d(B_j)$. On the other hand, if Ann bids truthfully she will win B and her payment will be exactly $v_d(B)$, which is lower, and so she will be better off, a contradiction. \square

Proof of Theorem 16

During the course of the proof of Theorem 16, it will be useful to think of Ann's problem of selecting an optimal shill bidding strategy side-by-side with coalition J 's problem of finding the optimal collusive strategy. In the proof that follows p_B^{VCG} is the VCG price that Ann would have to pay in the VCG mechanism when bidding against Bob, and is given by (2). From a mathematical point of view $p_B^{VCG} = p_B^{\text{Merged}}$, although these two terms have different interpretations. p_B^{VCG} is to be contrasted with $p_B^{VCG,j}$ for $j \in J$, where the latter is the VCG price that j has to pay when bidding against both Bob *and* the other members of J .

1: This an immediate consequence of Theorem 15.

²³Here we have used the fact that for all $C \subseteq N$ and all $y \in C$, $v_J(C) - v_J(C \setminus y) \geq v_a(C) - v_a(C \setminus y)$

²⁴This follows from the fact that if $v_a(B) - v_a(B \setminus x) = 0$, then it is also efficient for Ann to win $B \setminus x$, and x must have zero marginal value to all bidders at the efficient allocation, so whether Ann receives x or not does not affect her payment (both when she bids truthfully and when she bids \bar{v}).

2: Theorem 3 implies that under SubTop, Ann never has an incentive to use skills. It follows that if there existed a skill bidding strategy which was sometimes better than truthful bidding for Ann, there would also be a strategy not using skills that was sometimes better than truthful bidding. This would contradict the fact that without skills, truthful bidding is a dominant strategy. The result now follows from Observation 1.

3: If SubTop fails, then there exists $B \subseteq N$ such that $p_B^{\text{Skill}} < p_B^{\text{VCG}}$. Let \mathcal{P}^* solve the CMP for B . Then $|\mathcal{P}^*| = 2$. Then set $J := \mathcal{P}^*$ and for all $j = P \in \mathcal{P}^*$ set $v_j := v^{P, r_P}$ for sufficiently large r_P (in the sense that conditions 1 and 2 of Theorem 1 are satisfied). When coalition J has valuations $(v_j : j \in J)$, it is efficient to allocate exactly package B to J . Note finally that $p_B^{\text{VCG}} = p_B^{\text{Merged}}$, and by construction, $p_B^{\text{Skill}} = p_B^{\text{VCG}, J}$. In other words, we have just constructed the coalition J so that the valuations of the different members of J just happen to coincide with the valuations of the skill bidders that solve the CMP for B . Part 4 is also proven by using this construction. \square

Proof of Lemma 1

Throughout the course of this proof, I will relax the assumptions that the valuation v_b is monotone and that $v_b(B) \geq 0$ for all $B \subseteq N$. If the equivalence holds without these assumptions, then of course it holds with these assumptions as well. Note that I will maintain the assumption that $v_b(\emptyset) = 0$.

It is well known that a valuation v_b (with $v_b(\emptyset) = 0$) is additive if and only if it is both submodular and supermodular. Since SubTop is weaker than submodularity and SupTop is weaker than supermodularity, it follows that any additive valuation satisfies both SubTop and SupTop.

If v_b satisfies both SubTop and SupTop, then for all $B \subseteq N$ with $1 \leq |B| \leq |N| - 2$:

$$v_b(N) - v_b(B) = \sum_{x \in N \setminus B} [v_b(B \cup x) - v_b(B)] \quad (29)$$

(29) is equivalent to:

$$v_b(B) = \frac{[\sum_{x \in N \setminus B} v_b(B \cup x)] - v_b(N)}{|N \setminus B| - 1} \quad (30)$$

Applying (30) recursively, it follows that given $(v_b(B) : B \subseteq N, |B| = |N| - 1)$, it is possible to derive $(v_b(B) : B \subseteq N, 1 \leq |B| \leq |N| - 2)$. Moreover, if v_b satisfies SubTop and SupTop, it must also satisfy (29) for $B = \emptyset$, and since $v_b(\emptyset) = 0$, this means that:

$$v_b(N) = \sum_{x \in N} v_b(x),$$

so this means that if we know $(v_b(B) : B \subseteq N, |B| = |N| - 1)$, we can derive $v_b(N)$ as well.

Next observe that for any additive valuation v_b and for any $x \in N$ and $A \subseteq N$ with

$|A| = |N| - 1$:

$$v_b(x) = \frac{[\sum_{B \ni x: |B|=|N|-1} v_b(B)] - (|N| - 2)v_b(N \setminus x)}{|N| - 1} \quad (31)$$

On the other hand, if given $(v_b(B) : B \subseteq N, |B| = |N| - 1)$, we define $v_b(x)$ by (31), then a simple calculation shows that for all $B \subseteq N$ with $|B| = |N| - 1$, $\sum_{x \in B} v_b(x) = v_b(B)$. It follows that for every profile $(v_b(B) : B \subseteq N, |B| = |N| - 1)$, there exists exactly one additive valuation w such that for all $B \subseteq N$ with $|B| = |N| - 1$, $w(B) = v_b(B)$.

Because (i) any profile $(v_b(B) : B \subseteq N, |B| = |N| - 1)$ uniquely determines an additive valuation and also uniquely determines a valuation satisfying both SubTop and SupTop, and (ii) any additive valuation satisfies both SubTop and SupTop, it follows that the set of additive valuations equals the set of valuations satisfying both SubTop and SupTop. \square

Proof of Theorem 17

If v_b is additive, then it satisfies SubTop, so from Theorem 3, there is no incentive to shill. Next consider the incentive to merge. Since it is a dominant strategy for a single agent to bid truthfully in the VCG mechanism if shills are unavailable, it follows that conditional on submitting only one bid for an entire coalition J , it is optimal for the coalition to submit a bid equal to the aggregate valuation v_J of J . However, if the coalition is minimally competitive against v_b , so that the Competition Effect is null. Moreover, any additive valuation satisfies SubTop. It now follows from Theorem 15 that the integration effect is nonnegative—meaning that the coalition cannot reduce its payment by merging—and hence submitting the single bid v_J is no better than having each member of J submit a truthful bid.

Next suppose that v_b is not additive. Then Lemma 1 implies that v_b violates either SubTop or SupTop. If v_b violates SubTop, then by Theorem 3, there is an incentive to shill. If v_b violates SupTop, then there exists B and partition \mathcal{P} such that (18) is violated. Then a noncompetitive coalition $J = \mathcal{P}$ such that it would be efficient for $j = P \in J$ to win package P would have an incentive to merge. \square

Proof of Theorem 18

First suppose that v_b satisfies SubTop. Theorem 16 implies that it is optimal for J to submit v_J , which leads to an efficient allocation, not only in the original economy, but also in the marginal economies excluding any member i of Bob; moreover, social welfare calculated based on submitted bids in these economies is equal to true social welfare. Since i 's utility depends only on social welfare in the original economy and the economy excluding i , it now follows that 1 implies both 2 and 3.

If $|J|$ is large enough, $p_B^{\text{Shill}, |J|} = p_B^{\text{Shill}}$. If v_b violates SubTop, this implies that there exists $B \subseteq N$ such that $p_B^{\text{Shill}} < p_B^{\text{Merged}}$. Suppose that $v_J = v^{B,r}$ for r such that $p_B^{\text{Shill}} <$

$r < p_B^{\text{Merged}}$. Then at the unique efficient allocation, Bob receives N . However, given the optimal shill bidding strategy, J would receive B , implying inefficiency, so 1 implies 3. Next suppose that I contains only a single member—who we may also call Bob—with valuation v_b . Notice that in this example, Bob’s utility is reduced from $v_b(N) - r$ to $v_b(N \setminus B)$ which is smaller by the definition of r . This shows that 1 implies 3. \square

References

- AUSUBEL, L., AND P. MILGROM (2002): “Ascending Auctions with Package Bidding,” *Frontiers of Theoretical Economics*, 1(1), Article 1.
- (2006): “The Lovely but Lonely Vickrey Auction,” in *Combinatorial Auctions*, ed. by P. Crampton, Y. Shoham, and R. Steinberg, pp. 17–40. MIT Press.
- BICHELER, M., A. DAVENPORT, G. HOHNER, AND J. KALAGNANUM (2006): “Industrial Procurement Auctions,” in *Combinatorial Auctions*, ed. by P. Crampton, Y. Shoham, and R. Steinberg, pp. 593–612. MIT Press.
- BLUMROSEN, L., AND N. NISAN (2007): “Combinatorial Auctions,” in *Algorithmic Game Theory*, ed. by N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, pp. 267–299. Cambridge University Press.
- CANTILLON, E., AND M. PESENDORFER (2006): “Auctioning Bus Routes: The London Experience,” in *Combinatorial Auctions*, ed. by P. Crampton, Y. Shoham, and R. Steinberg, pp. 573–591. MIT Press.
- CAPLICE, E., AND Y. SHEFFI (2006): “Combinatorial Auctions for Truckload Transportation,” in *Combinatorial Auctions*, ed. by P. Crampton, Y. Shoham, and R. Steinberg, pp. 539–572. MIT Press.
- CONITZER, V. (2007): “Limited Verification of Identities to Induce False-Name-Proofness,” in *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK’07)*, pp. 102–111.
- (2008): Personal Communication.
- CONITZER, V., AND T. SANDHOLM (2006): “Failures of the VCG Mechanism in Combinatorial Auctions and Exchanges,” in *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multi Agent Systems*, pp. 521–528.
- DAY, R., AND P. MILGROM (2008): “Core-Selecting Package Auctions,” *International Journal of Game Theory*, pp. 393–407.

- FUJISHIGE, S. (2005): *Submodular Functions and Optimization*, vol. 58 of *Annals of Discrete Mathematics*. Elsevier, second edn.
- GRAHAM, D., AND R. MARSHALL (1987): “Collusive Behavior at Single-Object Second-Price Auctions,” *Journal of Political Economy*, 95, 1217–1239.
- GREEN, J., AND J.-J. LAFFONT (1977): “Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods,” *Econometrica*, 45, 427–438.
- GUL, F., AND E. STACCHETTI (1999): “Walrasian Equilibrium with Gross Substitutes,” *Journal of Economic Theory*, 87, 95–124.
- HOLMSTROM, B. (1979): “Groves Schemes on Restricted Domains,” *Econometrica*, 55, 303–328.
- LEHMANN, B., D. LEHMANN, AND N. NISAN (2006): “Combinatorial Auctions with Decreasing Marginal Utilities,” *Games and Economic Behavior*, 55, 270–296.
- LEHMANN, D., L. O’CALLAGHAN, AND Y. SHOHAM (2002): “Truth Revelation in Approximately Efficient Combinatorial Auctions,” *Journal of the ACM*, 49, 577–601.
- MATSUO, T., I. TAKAYUKI, R. DAY, AND T. SHINTANI (2006): “A Robust Combinatorial Auction Mechanism against Shill Bidders,” in *Automated Agents and Multi-Agent Systems (AAMAS-06)*.
- MILGROM, P. (2000): “Putting Auction Theory to Work: The Simultaneous Ascending Auction,” *Journal of Political Economy*, 108, 254–272.
- (2004): *Putting Auction Theory to Work*. Cambridge University Press.
- MOULIN, H. (2008): “Proportional Scheduling, Split-Proofness, and Merge-Proofness,” *Games and Economic Behavior*, 63, 567–587.
- RASSENTI, S., V. SMITH, AND R. BULFIN (1982): “A Combinatorial Auction Mechanism for Airport Time Slot Allocation,” *The Bell Journal of Economics*, 13, 402–417.
- RASTEGARI, B., A. CONDON, AND K. LEYTON-BROWN (2007): “Revenue Monotonicity in Combinatorial Auctions,” in *Association for the Advancement of Artificial Intelligence (AAAI-07)*, pp. 122–127.
- ROTHKOPF, M. (2007): “Thirteen Reasons why the Vickrey-Clarke-Groves Process is Not Practical,” *Operations Research*, 55, 191–197.
- SAKURAI, Y., M. YOKOO, AND S. MATSUBARA (1999): “A Limitation of the Generalized Vickrey Auction in Electronic Commerce: Robustness Against False-Name Bids,” *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, pp. 86–92.

- SANGHVI, S., AND D. PARKES (2004): “Hard-to-Manipulate VCG-Based Auctions,” Technical Report, Harvard University.
- TOPKIS, D. (1998): *Supermodularity and Complementarity*. Princeton University Press.
- YOKOO, M. (2006): “Pseudonymous Bidding in Combinatorial Auctions,” in *Combinatorial Auctions*, ed. by P. Crampton, Y. Shoham, and R. Steinberg, pp. 265–294. MIT Press.
- YOKOO, M., AND A. IWASAKI (2007): “Making VCG More Robust in Combinatorial Auctions via Submodular Approximation,” in *Twenty-Second Conference on Artificial Intelligence*, pp. 763–769.
- YOKOO, M., Y. SAKURAI, AND S. MATSUBARA (2000): “The Effect of False-Name Declarations in Mechanism Design: Towards Collective Decision Making on the Internet,” in *Proceedings of the Twentieth International Conference on Distributed Computing Systems, ICDCS-2000*.
- (2004): “The Effect of False-Name Bids in Combinatorial Auctions: A New Fraud in Internet Auctions,” *Games and Economic Behavior*, 46, 174–188.