Collusion, Efficiency, and Dominant Strategies

Alan Deckelbaum*
Math, MIT
deckel@alum.mit.edu

Silvio Micali
EECS, MIT
silvio@csail.mit.edu

March 23, 2014

Abstract

Green and Laffont [1] have proven that no collusion-resilient dominant-strategy mechanism guarantees efficiency in multi-unit auctions if a player is allowed to report only a single valuation. Chen and Micali [2] have bypassed this impossibility by slightly enlarging the strategy spaces, but via mechanisms that can impose impractically high fines. For unrestricted combinatorial auctions, the possibility of guaranteeing efficiency in collusion-resilient dominant strategies has remained open.

In this paper, we (a) generalize the notion of a collusion-resilient dominant-strategy mechanism by allowing for arbitrary strategy spaces, (b) construct such a mechanism for multi-unit auctions which is highly practical, and (c) prove that no such mechanism, practical or not, exists for unrestricted combinatorial auctions.

Our results hold in the very general collusion model of Green and Laffont, where the mechanism does not know who colludes with whom, a player may or may not have information about players with whom he does not collude, and players belonging to the same coalition can make side payments to each other and perfectly coordinate their strategies.

*Supported by Fannie and John Hertz Foundation Daniel Stroock Fellowship.

1 Introduction

Collusion is a major problem in mechanism design, and in auctions in particular. Dominant-strategy mechanisms assure that the players will choose the desired strategies, but not in the presence of collusion. When the players act independently, in unrestricted combinatorial auctions the VCG mechanism guarantees efficiency in dominant strategies [3, 4, 5]. Yet, Ausubel and Milgrom [6] show that just two collusive players can destroy the efficiency of the VCG mechanism. It is therefore legitimate to ask whether efficiency in unrestricted combinatorial auctions can be guaranteed in "collusion-resilient dominant strategies".

In a classical paper, Green and Laffont [1] put forward one such notion, *coalition incentive compatibility*. In their collusion model,

(a) Coalition members act so as to maximize the sum of their individual utilities;

(b) The members of each coalition can perfectly coordinate their actions; and

(c) Coalitions are secret and non-overlapping.

Note that the Green-Laffont collusion model is quite unrestricted. For example, the players of a coalition can make side-payments and enter binding contracts with each other. A mechanism need not know who colludes with whom, and a player need not know anything about the players outside his coalition. Non-overlapping coalitions is a restriction, but a natural one when the focus is on dominant-strategy implementation.[1]

In their collusion model, Green and Laffont prove that coalition-incentive-compatible mechanisms are unable to implement many social choice functions of interest. In particular, no efficient multi-unit auction mechanism satisfies their notion. However, while their collusion model is unrestricted, their notion of dominant-strategy collusion resilience is not, and their result applies only to mechanisms in which a player can only report a single valuation. In the Green-Laffont collusion model, Chen and Micali [2] propose a more general notion of a collusion-resilient dominant-strategy mechanism. They allow a player to report not only a single valuation, but also a set of other players (allegedly, the set of his colluders). They prove that one such mechanism guarantees efficiency in m-unit auctions, where m identical copies of the same good need to be allocated. Their mechanism never has negative revenue. However, two drawbacks limit its applicability. Namely, their mechanism is assumed to

(1) know an upper bound V to any value a player might have for a copy of the good, and

(2) be able to impose impractically large fines, even when all bids are low.

Specifically, even when all players bid low amounts for each copy of the good, their mechanism is assumed to be able to impose fines larger than $2mV$. First, this is problematic because no such upper bound may exist, let alone be known to the mechanism. Moreover, to ensure

[1]Indeed, a player belonging to two distinct coalitions may find it impossible to simultaneously maximize the utility of both of them. Should this case occur, his preferences would be unspecified.

that it exceeds any *possible* value a player may have for a copy of the good, V might be astronomically high, and $2mV$ even higher. Therefore, a player may not have sufficient funds to pay such a fine. If it is not credible that the envisaged fine is enforceable, then the whole dominant-strategy structure of their mechanism collapses.

The possibility of guaranteeing efficiency in unrestricted combinatorial auctions by means of a non-negative-revenue mechanism, whether practical or not, has remained totally open.

In this paper, we adopt the Green-Laffont collusion model, but propose a more general notion of collusion resilience in dominant strategies. As before, we demand both that (i) each independent player has a "best" strategy no matter what all other players do, and that (ii) each coalition C has a best strategy subprofile, no matter what the players outside C may do. What differentiates our approach from the prior ones is that we do not envisage any restrictions on the strategy spaces of the players. Envisaging an unrestricted strategy space is a simple extension, but has significant implications on whether or not efficiency is achievable in collusion-resilient dominant strategies.

On the positive side, we prove the usefulness of our general notion by constructing a *very practical* mechanism that guarantees efficiency in multi-unit auctions in the Green-Laffont collusion model. Our mechanism has non-negative revenue and avoids the two aforementioned drawbacks of [2]: it need not know any bound on the players' valuations and does not impose on any player fines larger than the values reported by that player. Note that our mechanism presupposes the 'decriminalization' of collusion. To be sure, when all players, collusive or not, use the best strategies available, the mechanism may deduce significant information about who colludes with whom. Criminalization of collusion, however, is not necessary. All we have to do is to design mechanisms that are resilient to collusion.

On the negative side, we prove that no collusion-resilient dominant-strategy mechanism can guarantee efficiency in unrestricted combinatorial auctions in the Green-Laffont collusion model. This impossibility is very strong, because it applies to mechanisms with arbitrary strategy spaces and to auctions with even two goods and three players.

Let us also mention that weaker notions of collusion resilience have been previously considered in the literature. As put forward by Schummer [7], *bribe proofness* requires that no two players can benefit by colluding together in a particular manner. The notion of *c-truthfulness*, due to Goldberg and Hartline [8], assumes that each coalition has at most c members. Other notions of collusion resiliency envisage that the players are incapable of making side payments; see [9, 10, 11, 12, 13, 14, 15, 16]. Laffont and Martimort [17] and Che and Kim [18] also study collusion resilience using equilibrium-based solution concepts. *Collusion leveraging*, as put forward by Chen, Micali, and Valiant [19], aims at leveraging the players' knowledge about the payoff types of their opponents.

2 Auctions

Auction Contexts In an auction, a player's *type* is also called a *valuation*.

In a multi-good auction, a valuation is a function mapping the empty subset to 0, and every other subset of the goods to a non-negative real number.

The set of players is $N = \{1, 2, \ldots, n\}$, the number of goods is m, the set of all possible valuations of a player i is Θ_i, and the true valuation of player i is θ_i^*. An outcome ω is a pair (A, P), where $A = (A_0, \ldots, A_n)$ is a vector of subsets of the goods and P is a profile of numbers. Vector A is referred to as the allocation of ω and must be such that $A_i \cap A_j = \emptyset$ whenever $i \neq j$. Component A_0 represents the set of unallocated goods, and, for $i > 0$, A_i represents the subset of the goods allocated to player i. Each P_i represents the price paid by player i. For each player i, i's utility function u_i maps a valuation θ_i and an outcome $\omega = (A, P)$ to $u_i(\theta_i, \omega) \triangleq \theta_i(A_i) - P_i$. For brevity, when the true valuation of a player i is clear, we may write $u_i(\omega)$ instead of $u_i(\theta_i^*, \omega)$. An allocation A is *efficient* if $\sum_i \theta_i^*(A_i) \geq \sum_i \theta_i^*(A_i')$ for all allocations A'.

In a m-unit auction, there are m identical copies of the same good for sale.

A valuation of a player i is a sequence of non-negative real numbers, $t_i = t_i^{(1)}, \ldots, t_i^{(m)}$, such that $t_i^{(1)} \geq \cdots \geq t_i^{(m)} \geq 0$. That is, $t_i^{(j)}$ represents i's marginal value for a j-th copy of the good, and valuations have non-increasing marginals. An allocation A is a sequence of $n+1$ non-negative values whose sum is m, $A = A_0, A_1, \ldots, A_n$, where A_0 is the number of unallocated copies and, for $i > 0$, A_i is the number of copies allocated to player i. An outcome ω is a pair (A, P), where A is an allocation and P a profile of prices. The utility of a player i with valuation t_i for an outcome $\omega = (A, P)$ is $\sum_{j=1}^{A_i} t_i^{(j)} - P_i$.[2] All other notions and notations for multi-good auctions (such as efficiency and $u_i(\omega)$) automatically extend to multi-unit ones.

Auction Mechanisms A mechanism M for an auction context specifies:

- For each player i, the set S_i of pure strategies available to i, and
- A function (traditionally also denoted by M) mapping each strategy profile in $S = S_1 \times \cdots \times S_n$ to an outcome (A, P).

For $s \in S$, we denote by $M(s)$ the outcome generated by M and by $u_i(M(s))$ the corresponding utility of player i. If M is probabilistic, $M(s)$ is a distribution over outcomes, and $u_i(M(s))$ is the corresponding expected utility of i. If the underling mechanism M is clear, we may write $u_i(s)$ instead of $u_i(M(s))$.

Additional Notation Let i be a player and A a subset of players. Then,

- For every profile x, x_A is the subprofile obtained by restricting x to A.
- S_A and Θ_A respectively are the Cartesian product $\prod_{i \in A} S_i$ and $\prod_{i \in A} \Theta_i$.
- $-i$ is the set $N \setminus \{i\}$, and $-A$ is the set $N \setminus A$.

[2] If $A_i = 0$, then, according to usual conventions, $\sum_{j=1}^{A_i} t_i^{(j)}$ is 0.

3 Collusive Rationality

Recall that a partition of a set T is is a collection of subsets of T, $T_1, \ldots, T_k$, such that $\cup_i T_i = T$ and $T_i \cap T_j = \emptyset$ whenever $i \neq j$.

The Green-Laffont model Before the mechanism is executed (but possibly after the mechanism has been announced), the players are free to form an arbitrary partition of N. The formed partition, the *collusive partition*, is denoted by $\mathbb{C}$. A set in $\mathbb{C}$ is called a *coalition*. A player i is *independent* if $\{i\} \in \mathbb{C}$.

If $C \in \mathbb{C}$, then C is common knowledge to its members, but a player in C need not have any additional information about $\mathbb{C}$. The mechanism has no information about $\mathbb{C}$.

Focussing on auctions, for each subset C of the players, the *collective utility function* of C, u_C, maps a valuation subprofile $\theta_C \in \Theta_C$ and an outcome ω to $\sum_{i \in C} u_i(\theta_i, \omega)$. When the valuation subprofile θ_C under consideration is clear, we may write $u_C(\omega)$ instead of $u_C(\theta_C, \omega)$. If C is a coalition, then its players can perfectly coordinate their actions,[3] and act so as to maximize u_C.

Collusive solution concepts A pure strategy subprofile $s_A \in S_A$ is *(weakly) dominant* for a subset of players A if $u_A(s_A, s_{-A}) = \max_{s'_A \in S_A} u_A(s'_A, s_{-A})$ for all $s_{-A} \in S_{-A}$.

A mechanism is *collusive dominant-strategy* if for every subset of players A and every true valuation subprofile $\theta^*_A \in \Theta_A$ there exists a pure strategy subprofile $d_A \in S_A$ which is dominant for A.

A collusive dominant strategy mechanism is (ex-post) *coalitionally rational* if for every subset of players A, every true-valuation subprofile $\theta^*_A \in \Theta_A$, and every strategy subprofile $s_{-A} \in S_{-A}$, the sum utility of A is non-negative under d_A. That is, $u_A(d_A, s_{-A}) \geq 0$.

Efficiency in collusive dominant-strategies A collusive dominant-strategy auction mechanism M is *efficient* if, for every collusive partition $\mathbb{C} = C_1 \cup C_2 \cup \cdots \cup C_k$ and every true valuation profile θ^*, the outcome $M(d_{C_1}, \ldots, d_{C_n})$ is an efficient allocation for θ^*.

4 A Practical Collusive Dominant-Strategy Mechanism for Multi-Unit Auctions

Our multi-unit auction mechanism modifies the one of Vickrey [3].

4.1 The Standard Vickrey Mechanism

In the mechanism of Vickrey [3] for m-unit auctions, a strategy consists of reporting a single valuation. Given a profile t of reported valuations, the mechanism constructs a sequence of $n \cdot m$ "value-owner" pairs $\{(t_i^{(k)}, i) : i \in N, k = 1, \ldots, m\}$, ordered in decreasing order with respect to the first ("value") component. We call the first m pairs in the sequence the "winning pairs" and all other ones the "losing pairs". For each player i, we let m_i be the

[3] E.g., the members of a coalition can enter binding agreements and make side payments to each other

number of winning pairs with owner i. The mechanism allocates m_i copies of the good to i, identifies the first m_i losing pairs whose owner is not i, and charges i the sum of the values of these pairs.

The Vickrey mechanism is efficient in dominant strategies in absence of collusion, but not otherwise [2].

4.2 The Intuition Behind Our Mechanism $\mathcal{M}$

To describe our mechanism $\mathcal{M}$, we find it useful to present first an 'auxiliary' mechanism $\mathcal{M}'$ guaranteeing efficiency under a solution concept that is very strong, but weaker than dominant strategies. We then show how to modify $\mathcal{M}'$ to obtain our collusive dominant-strategy, collusively rational, and efficient mechanism $\mathcal{M}$.

Mechanism $\mathcal{M}'$. In $\mathcal{M}'$, each player reports a sequence of m "value-beneficiary" *pairs*, $(v_1, b_1), \ldots, (v_m, b_m)$, where $(v_c, b_c) \in \mathbb{R} \times N$ signifies that the bidding player is willing to pay an amount v_c in order for player b_c to receive an additional copy of the good. For instance, if $n = 5$, $m = 3$ and player 2 reports $(\$10, 3), (\$8, 2), (\$4, 3)$, then player 2 declares that he is willing to pay to the mechanism \$10 for the first copy of player 3, \$8 for his own first copy, and \$4 for the second copy of player 3.

Given all these reports, $\mathcal{M}'$ transforms each "value-beneficiary" pair into a "value-beneficiary-owner" *triplet*. If, as in the example above, player 2 reports $(\$4, 3)$ as one of his pairs, then $\mathcal{M}'$ constructs $(\$4, 3, 2)$ as the corresponding triplet. After that, $\mathcal{M}'$ orders all $n \cdot m$ triplets in decreasing order by their value component, breaking ties by the beneficiary component (in any fixed order), and then by the owner component (in any fixed order). We call the first m triplets in the sequence the "winning triplets" and all others the "losing triplets". For each player i, let m_i be the number of winning triplets with beneficiary i, and let m'_i be the number of winning triplets with owner i. Then, $\mathcal{M}'$ allocates m_i copies of the good to i, identifies the first m'_i losing triplets whose owner is not i, and charges i the sum of the value components of these triplets.

An Intermediate Solution Concept. As $\mathcal{M}'$ is not our final mechanism, we do not define the solution concept under which $\mathcal{M}'$ can be proved to be efficient, nor do we discuss such a proof. We simply provide some intuition in a very informal manner.

Assume for a moment that each player only bids value-beneficiary pairs whose beneficiary belongs to his own coalition. Then, it can be proven that the following strategy subprofile is dominant for each coalition C: one member of C, i^*_C, bids the m pairs $(v_1, b_1), \ldots, (v_m, b_m)$, corresponding to the m highest marginal valuations of players in C, and all other members of C bid pairs with value 0. We refer to such a report as a "smart strategy subprofile". Clearly, when all coalitions choose smart strategy subprofiles, the allocation returned by $\mathcal{M}'$ is efficient.

However, if some player $j \notin C$ bids a pair (v, b) where $b \in C$, then the above strategy subprofile need not be dominant for C. For instance, suppose there are two copies of the same good, and two players, 1 and 2, both independent and both valuing \$10 a first copy and \$5 a second one. Assume that player 1 truthfully bids the sequence $(\$10, 1), (\$5, 1)$, but player 2, for whatever reason, bids $(\$8,1),(\$7,2)$. Under this bid profile, player 1 gets both goods (which of course is not an efficient allocation) and pays \$7, so that his utility is \$8. Player 1 would be better off, however, bidding $(\$5,1),(\$0,1)$, so as to get one copy for a price of zero and have a net utility of \$10. When player 1 is truthful, player 2 gets no goods, pays \$5, and his utility is $-\$5$. Notice that reporting $(\$8,1),(\$7,2)$ is, for player 2, dominated by reporting $(\$8,2),(\$7,2)$. More generally, in $\mathcal{M}'$, it is always best for a player to name only beneficiaries in his own coalition.

In sum, reporting a smart strategy subprofile is certainly "smart" but not quite dominant.[4]

From $\mathcal{M}'$ to $\mathcal{M}$. We modify $\mathcal{M}'$ by having each player i bid either (1) a "representative" player $j \in N$ or (2) a sequence of m "value-beneficiary" pairs $(v_1, b_1), \ldots, (v_m, b_m)$. In the first case, only the representative j has permission to declare a value-beneficiary pair with i as a beneficiary, while in the second case no other player has permission to declare a pair with i as a beneficiary. If any player declares a value-beneficiary pair with a disallowed beneficiary, that pair is discarded by the mechanism.

We must ensure, however, that i never has incentive to name a player outside of his collusive group as a representative. Namely: if any player k reports a pair with beneficiary i, and if k is not i's declared representative, then the mechanism not only discards the pair but also forces k to pay to i an amount of money equal to the value of the discarded pair. With this modification, i will never have incentive to name an outside player k as his representative: i would be better off discarding k's reported pair, receiving the money from k, and (if desired) reporting the pair himself.

4.3 Our Mechanism $\mathcal{M}$

Strategies. Our mechanism is of normal form. Every player i, simultaneously with his opponents, reports a strategy $s_i \in N \cup (\mathbb{R}_+ \times N)^m$. That is, for every strategy s_i of player i, either

- s_i is a player, $s_i \in N$, or
- s_i is a sequence of m pairs: $s_i = (v_1, b_1), \ldots, (v_m, b_m) \in (\mathbb{R}_+ \times N)^m$.

A player i reporting $s_i \in N$ is called *passive*. In this case, we refer to s_i as "i's representative," denoted by $Rep(i)$.

A player i reporting $s_i = (v_1, b_1), \ldots, (v_m, b_m) \in (N \times \mathbb{R}_+)^m$ is called *active*. In this case we refer to each (v_j, b_j) as a value-beneficiary pair, with value v_j and beneficiary b_j.

Choosing the Outcome. When the players report a strategy profile s, the mechanism chooses to allocate the copies of the goods and the price charged to each player by means of the following steps:

[4] Indeed, *if every coalition C' eliminates all strategy subprofiles that are dominated for C'*, then for every coalition C, a smart strategy subprofile is dominant for C *with respect to all surviving strategies.*

1. Initialize L to be an empty list of "value-beneficiary-owner" triplets.
2. For each active player i and value-beneficiary pair (v, b) reported by i:
   (a) If either `(i)` $b = i$ or `(ii)` b is passive and $Rep(b) = i$,
   then append the value-beneficiary-owner triplet (v, b, i) to L.
   (b) Otherwise, i makes a payment of v to b.
   (Call R_i the net payment of i after all executions of step 2b.)
3. Sort the value-beneficiary-owner triplets of L in decreasing order by their value component. (If needed, break ties first by beneficiary, and then by owner, in lexicographic order.)

   (Call the first m triplets "winning" and the remaining triplets "losing.")
4. Each player j receives m_j copies of the good, where m_j is the number of winning triplets with beneficiary j.[5]
5. For each active player i:
   (a) Let m'_i be the number of winning triplets with owner i.
   (b) Identify the first m'_i losing triplets in L whose owner is not i and (in addition to any payments from step 2b) have player i pay to the mechanism an amount Q_i equal to the sum of the value components of these triplets.[6]

Note that, in $\mathcal{M}$, the final price of a player i is $P_i = R_i + Q_i$. (As usual, if P_i is positive, then i disburses money, and otherwise he receives money.)

4.4 Analysis

Theorem 1. *Mechanism $\mathcal{M}$ is collusive dominant-strategy, coalitionally rational, and efficient.*

Proof. Consider an arbitrary collusive group A. We begin with three claims:

Claim 1: Arbitrarily fix a strategy subprofile s_A for the players in A and a strategy subprofile s_{-A} for the other players. Then there exists a strategy subprofile $\widetilde{s}_A \in S_A$ such that

(a) $u_A(\widetilde{s}_A, s_{-A}) \geq u_A(s_A, s_{-A})$, and

(b) for all $i \in A$,

if i is passive in $\widetilde{s}_i$, then $\widetilde{s}_i \in A$

(i.e., i chooses his representative in his own coalition)

if i is active in $\widetilde{s}_i$, then the beneficiaries of all pairs in $\widetilde{s}_i$ are in A.

[5] If L contains fewer than m winning triplets, the remaining goods are unallocated.

[6] If there are fewer than m'_i such triplets, instead charge i the sum of the values of all such triplets.

Proof of Claim 1: For all $i \in A$, we construct $\widetilde{s}_i$ by modifying the original strategy s_i according to the following (exhaustive) three cases.

1. Player i is passive in s_i and nominates a player $j \in A$ as his representative.

   Then $\widetilde{s}_i = s_i$.

2. Player i is active in s_i and $(v_1, j_1), \ldots (v_m, j_m)$ are his value-beneficiary pairs.

   Then player i is active in $\widetilde{s}_i$. He bids all pairs (v_l, j_l) from s_i for which $j_l \in A$. For the remaining pairs, if any, he bids $(0, i)$.

3. Player i is passive in s_i but nominates a player $j \notin A$ as his representative.

   Then player i is active in $\widetilde{s}_i$ and chooses his m value-beneficiary pairs as follows:

   - If j is passive in s_{-A}, then $\widetilde{s}_i$ consists of m pairs $(0, i)$.
   - If j is active in s_{-A}, letting $(v_1, i), \ldots, (v_k, i)$ denote all the value-beneficiary pairs of s_j with beneficiary i, set $\widetilde{s}_i$ to consist of the k pairs $(v_1, i), \ldots, (v_k, i)$ and the $m - k$ pairs $(0, i)$.

Let us first show that each player in A receives at least as many copies of the good under $(\widetilde{s}_A, s_{-A})$ as he does under (s_A, s_{-A}). Indeed, we shall argue that the triplets with beneficiary in A inserted into L under these two strategy profiles differ only in their owner components, with the exception that, under $(\widetilde{s}_A, s_{-A})$, additional triplets may be inserted into L with value 0 and beneficiary in A. This holds because of the following observations:

- For any $i \in A$ who is active in s_i, we have only removed his pairs (v, j) when $j \notin A$. Furthermore, for any pair (v, j) with $j \in A$, this pair "survives" into L under s only if either $j = i$ or j is passive with $Rep(j) = i$. In both of these cases, the pair will also survive into L under $(\widetilde{s}_A, s_{-A})$.

- For every $i \in A$ who is passive with $Rep(i) = j \notin A$ in s_i, there are up to m triplets $(v_1, i, j), \ldots, (v_k, i, j)$ which are inserted into L under (s_A, s_{-A}) but not under $(\widetilde{s}_A, s_{-A})$. The triplets $(v_1, i, i), \ldots, (v_k, i, i)$ are, however, appended to L under $(\widetilde{s}_A, s_{-A})$ but not under (s_A, s_{-A}).

Therefore, changing from (s_A, s_{-A}) to $(\widetilde{s}_A, s_{-A})$ affects only the owner component of triplets inserted into L with beneficiary in A, possibly appends additional triplets with value 0, and might remove triplets with beneficiary outside A. Since $\mathcal{M}$ breaks ties according to the beneficiary components of triplets before the owner components, every winning triplet under (s_A, s_{-A}) has a corresponding winning triplet under $(\widetilde{s}_A, s_{-A})$.

Finally, we claim that the net amount charged to A (i.e., $\sum_{i \in A} P_i$) under $(\widetilde{s}_A, s_{-A})$ is no more than the net amount charged to A under (s_A, s_{-A}). Indeed, if a player $i \in A$ is passive in s_i with $Rep(i) \in A$, then R_i is the same and $Q_i = 0$ in both scenarios. If i is passive in s_i with $Rep(i) \notin A$, then the amount he is newly charged in step 5b of $\mathcal{M}$ is no more than the additional amount he receives in step 2b. That is, in the original scenario Q_i was 0, and in the new scenario has become positive, yet the increase in Q_i is no more than the decrease in R_i. Thus, it is in i's interest to "discard" triplets with owner outside of A, receive payment equal to the discarded triplets' values, and declare the appropriate pairs himself.

Finally, if i is active in s_i, then R_i either stays the same or decreases in the second scenario, while Q_i may increase or decrease. However, denoting by A' the subset of the coalition A who are active in s_A, a simple case analysis shows that $\sum_{j \in A'} Q_j$ does not increase in $(\widetilde{s}_A, s_{-A})$. The only crucial case to consider is when, for some $i, j \in A'$, $\ell \in A$, and $k \notin A$, a triple (x, k, i) was winning in (s_A, s_{-A}), but after replacing (x, k, i) with $(0, i, i)$ a previously losing triple (x', ℓ, j) now wins in $(\widetilde{s}_A, s_{-A})$. In this case, Q_j increases by at most x', while Q_i decreases by at least x'.

Claim 2: Arbitrarily fix a strategy subprofile s_{-A} for $-A$. Let $\widetilde{s}_A$ be a strategy subprofile for A satisfying the second and third properties of Claim 1, namely:

- for $i \in A$, if i is passive in $\widetilde{s}_i$, then $\widetilde{s}_i \in A$.
- for $i \in A$, if i is active in $\widetilde{s}_i$, then the beneficiaries of all pairs in $\widetilde{s}_i$ are in A.

Then there exists a subprofile $s'_A \in S_A$ such that `(i)` a single player in A, i', is active, `(ii)` all other $j \in A$ are passive and $s'_j = Rep(j) = i'$, and `(iii)` $u_A(s'_A, s_{-A}) \geq u_A(\widetilde{s}_A, s_{-A})$.

Proof of Claim 2: Let $(v_1, b_1, q_1), \ldots, (v_k, b_k, q_k)$, where $k \leq m$, be the value-beneficiary-owner triplets which, under $(\widetilde{s}_A, s_{-A})$, are "winning" in L and have beneficiary $b_i \in A$.

Set s'_A to be the strategy subprofile where a single $i' \in A$ bids the pairs $(v_1, b_1), \ldots, (v_k, b_k)$ along with $m - k$ pairs $(0, i')$. Comparing the execution of step 2 of the mechanism under the strategy profiles (s'_A, s_{-A}) and (s_A, s_{-A}), we note the following differences: `(i)` the new triplets $(v_1, b_1, i'), \ldots, (v_k, b_k, i')$ and $m-k$ triplets $(0, i', i')$ are appended to L under (s'_A, s_{-A}) but not under (s_A, s_{-A}), `(ii)` the triplets $(v_1, b_1, q_1), \ldots, (v_k, b_k, q_k)$ are appended to L under (s_A, s_{-A}) but not under (s'_A, s_{-A}), and `(iii)` the losing triplets with beneficiary in A which are appended to L under (s_A, s_{-A}) are not appended under (s'_A, s_{-A}).

Since for every winning triplet (v_l, b_l, q_l) under (s_A, s_{-A}) with $b_l \in A$ there is a corresponding triplet (v_l, b_l, i') under (s'_A, s_{-A}) differing only in the owner component, and since the ordering on L breaks ties lexicographically by beneficiary before owner, it is clear that the allocation under (s'_A, s_{-A}) assigns every player in A at least as many copies of the good as under $(\widetilde{s}_A, s_{-A})$. (We note i' might receive additional copies under (s'_A, s_{-A}) if any of the $m - k$ new triplets $(0, i', i')$ are winning.) Furthermore, removing non-winning triplets from L does not affect the final allocation and does not increase the price Q_j charged to any player j in step 5b.

It is clear that the amount $Q_{i'}$ charged to i' in step 5b under (s'_A, s_{-A}) is no more than the net amount $\sum_{j \in A} Q_j$ charged to all of A in step 5b under $(\widetilde{s}_A, s_{-A})$. Indeed, by having a single player declare all of the pairs with beneficiary in A, more of the high-value triplets are skipped when computing the price in step 5b. In short, the subprofile s'_A avoids the scenario where a triplet owned by a player in A causes the price paid by a different player in A to increase. We note that any payments R_j received by players $j \in A$ from step 2b are the same under the two strategy profiles. $\square$

Claim 3: Let d_A be the following strategy subprofile for A:

- A single player $i^* \in A$ is active. He declares the m triplets $(v_1, j_i), \ldots, (v_m, j_m)$ corresponding to the m highest marginal valuations amongst players in A.[7]
- All other players $j \in A$ are passive and announce $d_j = Rep(j) = i^*$.

Then d_A is a best response of A against every $s_{-A} \in S_{-A}$.

Proof of Claim 3: By Claim 1 and Claim 2, it suffices to show that $u_A(d_A, s_{-A}) \geq u_A(s'_A, s_{-A})$, where s'_A is any strategy subprofile such that a single $i' \in A$ is active and all other $j \in A$ declare $Rep(j) = i'$.

We notice that under all such strategies s'_A, the player i' owns all the triplets in L which have beneficiary in A. The only remaining relevant features of the strategy subprofile are the pairs which i' bids. From the perspective of i', he is playing a standard Vickrey auction of a single good of limited supply, where his marginal valuations are the maximum of the marginal valuations in A. Thus, truthfully declaring the m highest marginal valuations is best for i'. A formal proof of this fact is nearly identical to the proof that the standard Vickrey auction is truthful. $\square$

Finally, notice that, for every profile of true valuations and every collusive partition $\mathbb{C} = A_1 \cup A_2 \cup \cdots \cup A_k$, when all coalitions A_i play the dominant subprofile d_A as in Claim 3, the resulting allocation is efficient. This holds since, under this strategy profile, one representative from each A_i bids truthfully on behalf of the m highest marginal valuations of A_i, no pairs are discarded in step 2b of $\mathcal{M}$, and the triplets with the m highest valuations are winners.

$\square$

5 A Simple Generalization of the Revelation Principle

In the next section we prove that no coalitionally rational and collusive dominant-strategy mechanism exists for unrestricted combinatorial auctions. Our proof applies to mechanisms with arbitrary strategy sets. (Already for our multi-unit auction mechanism $\mathcal{M}$ the pure strategies of a player i did not coincide with his set of possible valuations, Θ_i.) To simplify our analysis we argue, very much in the spirit of the revelation principle, that we need to consider only mechanisms with *very specific* strategy sets. Namely,

[7]For example, suppose $A = \{1, 2\}$, $m = 3$, player 1's value of obtaining a first, second, and third copy of the good are \$10, \$6 and \$2, respectively, and player 2's values of obtaining an additional copy of the good are \$9, \$7, and \$5. Then the active player would bid $(\$10, 1)$, $(\$9, 2)$, and $(\$7, 2)$. In the event of a tie among the marginal valuations, i^* can break the tie lexicographically according to the beneficiary.

Definition 1. *A collusive dominant-strategy mechanism M is* hyper-truthful *if:*

- *For each player i, $S_i = \{(A, \theta_A) : A \subseteq N, i \in A$ and $\theta_A \in \Theta_A\}$ and*
- *For every $A \subseteq N$, the strategy subprofile t_A where every $j \in A$ selects strategy $(A, \theta_A^*) \in S_j$ is dominant for A.*

In every collusive dominant-strategy mechanism, the optimal strategy subprofile for a set A of players depends only on the set A and the true type subprofile θ_A^* of the group, and not on the strategies or collusive structure of the outside players. Therefore, hyper-truthful mechanisms suffice in view of the following fact, whose proof is nearly identical to that of the revelation principle of Myerson [20].

Fact 1. *If there exists an efficient, coalitionally rational, collusive dominant-strategy mechanism M for an auction context $\mathcal{C}$, then there also exists an efficient, coalitionally rational, hyper-truthful mechanism M' for context $\mathcal{C}$.*

For example, let us show how our mechanism $\mathcal{M}$ can be transformed into a hyper-truthful mechanism $\mathcal{M}^*$. In an execution of $\mathcal{M}^*$, denote by (A_i, θ_{A_i}) the strategy selected by each player i. Then $\mathcal{M}^*$ works by simulating an execution of $\mathcal{M}$ as follows: If i is the lexicographically first player in A_i, then he is made active and made to bid the m highest marginal valuations in θ_{A_i}. Otherwise, i is made passive and made to declare his representative to be the lexicographically first player in A_i.

Remark Fact 1 is not, strictly speaking, a corollary of the revelation principle. This is so because a player's information about his own coalition is not an "original type." Indeed, the players partition themselves arbitrarily into collusive groups after a mechanism is announced. In principle, for one mechanism the players might want to collude in pairs, while for another mechanism they might want to collude in triples, and so forth. Nevertheless, the proof of the above fact follows from applying the revelation principle "conditionally on the collusive partition."

6 An Impossibility Result for Combinatorial Auctions

Theorem 2. *No collusive dominant-strategy mechanism can both be coalitionally rational and guarantee efficiency in combinatorial auctions with at least 3 players and 2 goods.*[8]

Proof. In light of Fact 1, it suffices to prove our thesis for mechanisms which guarantee efficiency in collusively dominant hyper-truthful strategies.

We proceed by contradiction. Assume the existence of an efficient, coalitionally rational, dominant-strategy hyper-truthful mechanism $\mathscr{M}$ for combinatorial auctions with 3 players and 2 goods. We derive a contradiction by showing that, when no upper-bound exists for a

[8]If $n > 3$ or $m > 2$, we can extend the proof below by making the extra players and goods "irrelevant".

player's value for a subset of the goods, then $\mathscr{M}$ must pay an infinite amount of money to the players when they report a special valuation profile. More precisely, for every $x > 0$, as long as every player may value $5x$ a subset of the goods, then $\mathscr{M}$ must return revenue $\leq -x$ whenever all players report that they belong to the same coalition and that they all value 0 every subset of the goods. (This would already be a contradiction for the case of bounded valuations, if the mechanism were required never to lose money.)

We obtain this contradiction via a sequence of 7 scenarios. In the first scenario, all players are independent and the true context and the bids are as follows:

Truth

Player	Coalition	Type		
		$\{a\}$	$\{b\}$	$\{a,b\}$
1	$\{1\}$	$5x$	0	$5x$
2	$\{2\}$	0	$5x$	$5x$
3	$\{3\}$	0	0	$2x$

Bids

Player	Declared Coalition	Type Subprofile
1	$\{1\}$	$(5x, 0, 5x)$
2	$\{2\}$	$(0, 5x, 5x)$
3	$\{3\}$	$(0, 0, 2x)$

Scenario 1

Faced with these bids, $\mathscr{M}$ must allocate good a to player 1, good b to player 2, and nothing to player 3, since it is efficient. Let us now argue that the prices charged in Scenario 1, $P_1^{(1)}$, $P_2^{(1)}$, and $P_3^{(1)}$, are ≤ 0. We note first that $P_3^{(1)} \leq 0$ follows immediately from individual rationality, since player 3 receives no goods. Suppose now that $P_1^{(1)} = \epsilon > 0$. Then player 1 would have incentive to change his bid to $\{1\}, (\epsilon/2, 0, \epsilon/2)$. In fact, by changing his strategy in this manner, while the bids of the other two players are unchanged, he would still be allocated the good (because the so modified bid profile could have been truthful and $\mathscr{M}$ is efficient), but would be charged at most $\epsilon/2$ (because the so modified bid profile could have been truthful and $\mathscr{M}$ is individual rational). This contradicts the fact that bidding truthfully is a dominant strategy for player 1 in Scenario 1. Therefore, $P_1^{(1)} \leq 0$. A similar argument shows that $P_2^{(1)} \leq 0$.

Consider now the following scenario, borrowed from Ausubel and Milgrom [6]:

Truth

Player	Coalition	Type		
		$\{a\}$	$\{b\}$	$\{a,b\}$
1	$\{1,2\}$	$x/2$	$x/2$	$x/2$
2	$\{1,2\}$	$x/2$	$x/2$	$x/2$
3	$\{3\}$	0	0	$2x$

Bids

Player	Declared Coalition	Type Subprofile
1	$\{1\}$	$(5x, 0, 5x)$
2	$\{2\}$	$(0, 5x, 5x)$
3	$\{3\}$	$(0, 0, 2x)$

Scenario 2

That is, players 1 and 2 are collusive but not truthful, so that the bids are identical to those of Scenario 1. Accordingly, $\mathscr{M}$ returns in Scenario 2 the same outcome as in Scenario 1. In particular, in Scenario 2 the collective utility of coalition $\{1, 2\}$ is at least x. Consider now the following scenario:

Truth

Player	Coalition	Type		
		$\{a\}$	$\{b\}$	$\{a,b\}$
1	$\{1,2\}$	$x/2$	$x/2$	$x/2$
2	$\{1,2\}$	$x/2$	$x/2$	$x/2$
3	$\{3\}$	0	0	$2x$

Bids

Player	Declared Coalition	Type Subprofile
1	$\{1,2\}$	$(\frac{x}{2},\frac{x}{2},\frac{x}{2})$ $(\frac{x}{2},\frac{x}{2},\frac{x}{2})$
2	$\{1,2\}$	$(\frac{x}{2},\frac{x}{2},\frac{x}{2})$ $(\frac{x}{2},\frac{x}{2},\frac{x}{2})$
3	$\{3\}$	$(0,0,2x)$

Scenario 3

The true context of Scenario 3 is identical to that of Scenario 2, but now players 1 and 2 are also truthful. Accordingly, since $\mathscr{M}$ is hyper-truthful, it must produce an outcome in which the collective utility of coalition $\{1,2\}$ is at least x. (Else, players 1 and 2 could increase their collective utility by bidding as in Scenario 2.) Furthermore, because $\mathscr{M}$ must return an efficient allocation when the bids are truthful, it must allocate both goods to player 3. Thus, to ensure that the collective utility of $\{1,2\}$ is at least x, $\mathscr{M}$ must pay to players 1 and 2 a total of at least x (i.e., $P_1^{(3)} + P_2^{(3)} \leq -x$). We now consider Scenario 4:

Truth

Player	Coalition	Type		
		$\{a\}$	$\{b\}$	$\{a,b\}$
1	$\{1,2\}$	0	0	0
2	$\{1,2\}$	0	0	0
3	$\{3\}$	0	0	$2x$

Bids

Player	Declared Coalition	Type Subprofile
1	$\{1,2\}$	$(\frac{x}{2},\frac{x}{2},\frac{x}{2})$ $(\frac{x}{2},\frac{x}{2},\frac{x}{2})$
2	$\{1,2\}$	$(\frac{x}{2},\frac{x}{2},\frac{x}{2})$ $(\frac{x}{2},\frac{x}{2},\frac{x}{2})$
3	$\{3\}$	$(0,0,2x)$

Scenario 4

Since the bids of Scenario 4 are identical to those of Scenario 3, $\mathscr{M}$ must return the same outcome. That is, the coalition $\{1,2\}$ obtains no items but receives a total payment of at least x. We now consider Scenario 5:

Truth

Player	Coalition	Type		
		$\{a\}$	$\{b\}$	$\{a,b\}$
1	$\{1,2\}$	0	0	0
2	$\{1,2\}$	0	0	0
3	$\{3\}$	0	0	$2x$

Bids

Player	Declared Coalition	Type Subprofile
1	$\{1,2\}$	$(0,0,0)$ $(0,0,0)$
2	$\{1,2\}$	$(0,0,0)$ $(0,0,0)$
3	$\{3\}$	$(0,0,2x)$

Scenario 5

Under the bids of Scenario 5, since $\mathscr{M}$ is efficient, both goods must be allocation to player 3. Furthermore, since $\mathscr{M}$ is hyper-truthful, the coalition $\{1,2\}$ must collectively receive at least as much utility as in Scenario 4, and thus must collectively receive a net payment of at least x. Finally, player 3 could only receive money from $\mathscr{M}$ (that is, $P_3^{(5)} \leq 0$), since otherwise (analogously to the argument above) he would have incentive to lower his bid and still receive the goods. We now consider Scenario 6:

Truth

Player	Coalition	Type		
		$\{a\}$	$\{b\}$	$\{a,b\}$
1	$\{1,2,3\}$	0	0	0
2	$\{1,2,3\}$	0	0	0
3	$\{1,2,3\}$	0	0	0

Bids

Player	Declared Coalition	Type Subprofile
1	$\{1,2\}$	$(0,0,0)$ $(0,0,0)$
2	$\{1,2\}$	$(0,0,0)$ $(0,0,0)$
3	$\{3\}$	$(0,0,2x)$

Scenario 6

The bids of scenario 6 are identical to those of scenario 5. Thus, the collective utility of coalition $\{1,2,3\}$ is $-P_1^{(5)} - P_2^{(5)} - P_3^{(5)}$, which is at least x. Finally, we consider Scenario 7:

Truth

Player	Coalition	Type		
		$\{a\}$	$\{b\}$	$\{a,b\}$
1	$\{1,2,3\}$	0	0	0
2	$\{1,2,3\}$	0	0	0
3	$\{1,2,3\}$	0	0	0

Bids

Player	Declared Coalition	Type Subprofile
1	$\{1,2,3\}$	$(0,0,0)^3$
2	$\{1,2,3\}$	$(0,0,0)^3$
3	$\{1,2,3\}$	$(0,0,0)^3$

Scenario 7

Since $\mathscr{M}$ is dominant-strategy hyper-truthful, the coalition $\{1,2,3\}$ must obtain at least as much net utility in Scenario 7 as they do in Scenario 6. Thus, the three players together must receive a net payment of at least x. ■

7 Final Remarks

Theorem 1 tells us that mechanism $\mathcal{M}$ guarantees perfect efficiency in an m-unit auction, in a tough collusion model, if all players are perfectly rational. Note that $\mathcal{M}$ does not require that the valuations of the members of a coalition C are common knowledge within C. It only requires that the m value-beneficiary pairs announced by the representative of a coalition C are indeed 'optimal' for C. Not even this value-beneficiary information, however, need be common knowledge within C.[9] Also note that $\mathcal{M}$ has some form of robustness. Informally speaking, when the players are not perfectly rational, but are rational enough to avoid fines and to have the representative of each coalition C report m value-beneficiary pairs that are sufficiently good for C, then $\mathcal{M}$ generates a sufficiently good social welfare.

Theorem 2 holds because the mechanism must work properly no matter how high a player's value for a subset of the goods may be. However, if (1) an upper bound V to every possible value that a player may have for a subset of the goods exists, and (2) negative revenue is not a problem, then there is an hyper-truthful, collusively rational, and collusive dominant-strategy mechanism that guarantees efficiency in unrestricted combinatorial auctions by "generating very negative revenue."[10]

[9] Possibly, the optimal m pairs could arise from some sort of iterative process within C.

[10] Very Informally, such a mechanism $\mathscr{M}'$ can be constructed by modifying the one of [2] as follows.

As collusion is a major problem in mechanism design, it is useful to develop strong notions of collusion resilience. It is also useful to construct practical mechanisms that, like $\mathcal{M}$ for multi-unit auctions, guarantee such resilience in dominant strategies for practical applications. The fact that no such mechanisms exist for unrestricted combinatorial auctions is a 'fact of life' (but such combinatorial auctions are also problematic in many other ways). This fact too, however, is useful to know.

In sum, it is important to understand which social choice functions are implementable in a collusion-resilient, and hopefully practical, way, and which are not so implementable, regardless of practicality.

References

[1] J. Green and J.-J. Laffont, "On coalition incentive compatibility," *Review of Economic Studies*, vol. 46, pp. 243–254, April 1979.

[2] J. Chen and S. Micali, "Collusive dominant-strategy truthfulness," *J. Economic Theory*, vol. 147, no. 3, pp. 1300–1312, 2012.

[3] W. Vickrey, "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, pp. 8–37, 1961.

[4] E. H. Clarke, "Multipart pricing of public goods," *Public Choice*, vol. 11, September 1971.

[5] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, pp. 617–631, July 1973.

[6] L. M. Ausubel and P. Milgrom, "The lovely but lonely vickrey auction," in *Combinatorial Auctions*, pp. 17–40, MIT Press, 2006.

[7] J. Schummer, "Manipulation through bribes," *Journal of Economic Theory*, vol. 91, no. 2, pp. 180–198, 2000.

[8] A. V. Goldberg and J. D. Hartline, "Collusion-resistant mechanisms for single-parameter agents," in *Symposium on Discrete Algorithms*, (Vancouver, British Columbia, Canada), pp. 620–629, SIAM, 2005.

If the reports of all players are consistent (i.e., all declared coalitions and valuations subprofiles match), then $\mathscr{M}'$ (1) gives a reward of kV to each player reporting to belong to a coalition of size k, and (2) "conceptually coalesces each coalition C to a single player, conceptually runs the VCG mechanism with these coalesced players so as to compute the subset of goods A_C and price P_C of each coalition C, charges P_C to —say— the lexicographically first player of C, and finally uses the declared valuation subprofile of C to allocate efficiently, and *ad personam*, the goods in A_C within C."

If, instead, a player i claims to be colluding with j but either (a) j's declared coalition differs from the coalition declared by i, or (b) i declares a different valuation function for j than j declares for himself, then $\mathscr{M}'$ imposes to each such player i a fine of $2n^2V$ payable to the mechanism and a huge fine payable to j, and no goods are allocated.

[9] E. Maskin, *Implementation and Strong Nash Equilibrium*, pp. 433–440. North Holland, 1979.

[10] H. Moulin and B. Peleg, "Cores of effectivity functions and implementation theory," *Journal of Mathematical Economics*, vol. 10, pp. 115–145, June 1982.

[11] S.-C. Suh, "Implementation with coalition formation: A complete characterization," *Journal of Mathematical Economics*, vol. 26, no. 4, pp. 409–428, 1996.

[12] H. Moulin, "Incremental cost sharing: Characterization by coalition strategy-proofness," *Social Choice and Welfare*, vol. 16, no. 2, pp. 279–320, 1999.

[13] S. Barbera and M. O. Jackson, "Strategy-proof exchange," *Econometrica*, vol. 63, pp. 51–87, January 1995.

[14] K. Jain and V. V. Vazirani, "Applications of approximation algorithms to cooperative games," in *Symposium on the Theory of Computing*, (Heraklion, Crete, Greece), pp. 364–372, ACM, 2001.

[15] H. Moulin and S. Shenker, "Strategyproof sharing of submodular costs: Budget balance versus efficiency," *Economic Theory*, vol. 18, pp. 511–533, Nov. 2001.

[16] M. Pal and E. Tardos, "Group strategyproof mechanisms via primal-dual algorithms," in *44th Annual Symposium on Foundations of Computer Science*, pp. 584–593, IEEE, 2003.

[17] J.-J. Laffont and D. Martimort, "Mechanism design with collusion and correlation," *Econometrica*, vol. 68, pp. 309–342, March 2000.

[18] Y.-K. Che and J. Kim, "Robustly collusion-proof implementation," *Econometrica*, vol. 74, pp. 1063–1107, July 2006.

[19] J. Chen, S. Micali, and P. Valiant, "Robustly leveraging collusion in combinatorial auctions," in *Innovations in Computer Science*, (Beijing, China), pp. 81–93, Tsinghua University Press, 2010.

[20] R. B. Myerson, "Incentive compatibility and the bargaining problem," *Econometrica*, vol. 47, pp. 61–73, January 1979.