Collusive Dominant-Strategy Truthfulness*

Jing Chen and Silvio Micali†
Computer Science and Artificial Intelligence Laboratory,
Massachusetts Institute of Technology,
Cambridge, MA 02139, USA

January 17, 2012

Abstract

We show that collusion and wrong beliefs may cause a dramatic efficiency loss in the Vickrey mechanism for auctioning a single good in limited supply. We thus put forward a new mechanism guaranteeing efficiency in a very adversarial collusion model, where the players can partition themselves into arbitrarily many coalitions, exchange money with each other, and perfectly coordinate their actions. Our mechanism bypasses classic impossibility results (such as those of Green and Laffont, and of Schummer) by providing the players with a richer set of strategies, making it dominant for every coalition C to instruct each of its members to report truthfully not only his own valuation, but also his belonging to C. Our mechanism is *coalitionally rational*, which implies being individually rational for independent players.

JEL classification: C70, D44, D70
Keywords: Vickrey auctions, Collusion, Efficiency

*We would like to thank Gabriel Carroll and Jason Hartline for pointing out relevant prior work. Thanks also to the editor and an anonymous referee for helping us clarify our paper. This work is supported in part by ONR Grant No. N00014-09-1-0597.
†Corresponding author. E-mail addresses: jingchen@csail.mit.edu (J. Chen), silvio@csail.mit.edu (S. Micali).

1 Introduction

The presence of collusion and wrong beliefs can clearly frustrate the aims of a mechanism designer. If all players act independently, beliefs do not come into play in a dominant-strategy mechanism, but collusion might continue to be a problem. Indeed, any equilibrium, including a dominant-strategy one, only guarantees that no individual player has any incentive to deviate from his envisaged strategy. However, two or more players may have plenty of incentive to *jointly* deviate from their equilibrium strategies.

Sometimes the ability of the players to collude is *constrained* by suitable assumptions. For instance, one may assume an upper bound to the number of possible colluders, the colluders' inability to keep secret their cooperation, their inability to make "side payments" to each other, and their mutual distrust. It has also been assumed that the "collusion structure", that is, the full description about who colludes with whom, may be known to the players. The collusion model envisaged in this paper is instead quite *unconstrained*.

Unconstrained Rational Collusion After a mechanism is announced, we assume that the players secretly partition themselves into arbitrarily many coalitions of arbitrary sizes. (An independent player forms a coalition of cardinality 1.) The members of the same coalition can make side payments to each other, and perfectly coordinate their actions —e.g., thanks to their ability to enter secret binding agreements. The members of each coalition C know each other perfectly well, but may have arbitrary beliefs (e.g., no belief at all, or false beliefs) about the preference and collusion structure of the players outside C. Our only constraint is that *all coalitions are rational.* That is, the members of the same coalition C act so as to maximize the sum of their individual utilities. Since they can separately compensate each other, this is indeed the rational thing for them to do.

Unconstrained Rational Collusion in Single-Good and Multi-Good Auctions The second-price mechanism, the Vickrey one [23], and the VCG [23, 6, 11] guarantee efficiency in dominant strategies in increasingly more general auctions: respectively, auctions of a single good, of a single good in multiple copies, and of multiple goods. Indeed, each mechanism can be considered a special case of the next. On one extreme, the second-price mechanism is so simple that it is "automatically" invulnerable to unconstrained rational collusion.[1] On the other extreme, the VCG, as shown by Ausubel and Milgrom [2], can be rendered totally inefficient by just two collusive players. So: what is the status of the Vickrey mechanism?

The Problem of Vickrey's Mechanism The Vickrey auction does not quite fall in either extreme. However, we prove by a simple example that its efficiency is highly vulnerable to a *combination* of unrestricted rational collusion and wrong beliefs. (Our example assumes some familiarity with the Vickrey mechanism, which is anyway recalled in Subsection 4.1.)

Consider a Vickrey auction (with ties broken at random) for a good available in two identical copies. As usual, we assume that a player's marginal value for a second copy is no greater than his value for a first copy: that is, all valuations are of the form (x, y), with $x \geq y$, where x represents a player's value for a first copy, and y his marginal value for a second one. There are 4 players: a and b, who form a coalition, and c and d, who are independent. Their respective valuations are (100,0), (100,0), (1,0), and (1,0). Accordingly, the maximum social welfare possible in this context is 200, and can be realized only by allocating one copy to a and the other to b.

The two collusive players know each other's valuations, and their beliefs are as follows: "c and d are independent and their respective valuations are (1000,0) and $(x, 0)$, where $x < 100$." Let us now argue that, with such beliefs, bidding truthfully is not a joint dominant strategy for a and b.

Because the Vickrey mechanism is dominant-strategy-truthful, the collusive players expect c and d to bid truthfully: that is, (1000,0) and $(x, 0)$. Therefore, a and b also expect that, if they too bid truthfully, then

[1] Informally, if the member with the highest valuation for the good in coalition C is player i, and if i's value for the good is v, then the "best collective strategy" for C is to have i bid v, and all other members bid 0. By so doing C may affect the auction's revenue, but not its efficiency.

at most one of them —say a— will win a copy of the good, in which case the bid of b will set a's price to 100. In sum, the colluders expect that, by bidding truthfully, their "collective utility" will be 0 in any case. This is not an attractive prospect for the two colluders. In accordance to their beliefs, a better strategy for a and b is for one of them to bid $(100, 0)$ and the other $(0, 0)$. If they bid so, however, and if c and d rationally —and thus truthfully— bid (1,0) and (1,0), then the Vickrey mechanism must allocate one copy to a collusive player and the other to an independent one, thus realizing a social welfare of 101 rather 200. In sum, the efficiency of Vickrey mechanism is vulnerable to unconstrained collusion with wrong beliefs.

Our Solution We modify the Vickrey auction so as to make it resilient to beliefs and collusion. Our mechanism actually *welcomes* colluders to the auction by providing them with special collusive bids, so that it becomes dominant for a coalition to report its presence and the true valuations of its members. Since independent players are just coalitions of size 1, truthful revelation is the best strategy for independent players and coalitions alike. In a sense, our *collusive dominant-strategy truthfulness* tries to harmonize the cooperative and the non-cooperative settings.

Our Conclusion So far, by forbidding and prosecuting collusion we have not eradicated it, and we have just pushed it underground where it continues to be quite disruptive. Perhaps it is time, at least in some applications, to try a new course: namely, bringing colluders into the open and incentivizing them to help us achieve our social goals.

Remarks

- *Novelty.* The basic idea of our solution consists of *enlarging the strategy space so as to include coalition information.* With this idea the mechanism becomes quite straightforward: It asks each bidder to report not only his own true valuation but also the coalition he belongs to. Truthfully reporting the coalition structure is made dominant by using proper cross checks and punishments. Once the truthful coalition structure has been elicited, the mechanism then runs a Vickrey auction among coalitions as if each coalition were a single bidder in the standard Vickrey auction.

  We note that all prior impossibility results about coalition incentive compatibility were deduced by applying the revelation principle to type spaces consisting of valuations only. Thus, although very simple, our basic idea might be novel and useful.

- *Practicality.* Our aims are primarily theoretical. Yet, our mechanism can be considered quite practical in the sense that it is conceptually simple; it is of normal form; it envisages very compact reports from the players (each report consisting of a traditional valuation plus a subset of the players); and it requires trivial computation from both the players and the mechanism. In addition, even our proof is a simple (and not too long) case analysis. The only potential concern is the big punishment that must be imposed to a player choosing to report untruthfully the coalition he belongs to.

- *Coalition Disjointness.* Coalition disjointness is not a totally arbitrary constraint. Since we adopt collusive dominant-strategy truthfulness as our solution concept and we assume that the members of each coalition try to maximize the sum of their individual utilities, in some cases a player belonging to multiple coalitions may find it impossible to act so as to maximize the collective utility of all of them.

- *The Hard Case.* Despite their ability of entering binding agreements, the players may fail to come to terms with one another, so that no coalition with more than one player might be formed. This of course would be an easy case, because faced with individually rational players the Vickrey mechanism is efficient. On the other extreme, if all players formed a single rational coalition, the Vickrey mechanism would continue to be efficient, because it makes possible for such players to bid so as to yield an efficient allocation with all prices equal to 0, which of course is the way for the "grand" coalition to maximize its utility. In our setting, the difficult case is when there are *multiple* coalitions and *nobody* has complete information about the collusion structure.

2 Related Work

Coalition Incentive Compatibility Our solution concept is closely related to coalition incentive compatibility, as put forward by Green and Laffont [10]. A coalition incentive-compatible auction mechanism *requires* that all strategies consist of valuations, and that every *agent*, whether an individual player or a coalition, has a dominant "course of action": an individual strategy for an independent player and a subprofile of individual strategies for a coalition. As it turns out, however, requiring strategies to coincide with valuations is a severe *restriction.* In fact, Green and Laffont prove the impossibility of maximizing social welfare via coalition incentive compatible mechanisms. More generally, the main results on coalition incentive compatibility consist of *impossibility results.* Our collusive dominant-strategy mechanism bypasses such impossibilities by endowing each individual strategy with an additional *coalitional component.*

Bribe-Proofness Schummer [21] proposes the notion of bribe-proofness. Different from coalition incentive compatibility, a social choice function is bribe-proof as long as no two players can jointly benefit by forming a coalition and deviating in a particular form. He shows that for a broad class of economies, a social choice function must be a constant function for being bribe-proof. Mizukami [16] further proves that this remains the case even under conditions weaker than that considered by Schummer. Once more, these results apply to mechanisms where the strategy space consists of valuations only, and thus do not apply to our mechanism.

Other Notions of Resiliency to Collusion Weaker notions of resiliency against rational collusion in settings of incomplete information have been investigated in the literature. In particular, Goldberg and Hartline [9] study *c-truthfulness*, where the cardinality of each coalition is upper-bounded by a value c. Laffont and Martimort [14] and Che and Kim [4] study collusion resiliency on a variety of solution concepts, ultimately all *equilibrium-based.* (The latter authors further allow the utility of a coalition to be the weighted sum of the individual utilities of its members.) Collusion resiliency when each coalition prefers an outcome ω to an outcome ω' if and only if each of its members prefers ω to ω', and when players cannot guarantee side-payments to one another, has been studied by [15, 18, 22] using *equilibrium*, and by [17, 3, 13, 19, 7, 20] using *group (or coalition) strategy-proofness. Collusion leveraging* has been investigated in [5], in order to leverage the players' knowledge about the payoff types of their opponents, via *rationally robust implementation*, a new solution concept not based on equilibria.

Consistency Checks As briefly mentioned, our mechanism elicits coalition information from the players, and relies on consistency checks —and punishments if needed— to ensure truthfulness. To be sure, the idea of consistency checks (but not for collusion structure) is quite classical in settings of complete information —see [12, 1, 8]. Consistency check for collusion structure was used in [5], under a weaker solution concept than ours, to generate revenue in combinatorial auctions.

3 Our Notions

Each game G can be decomposed into a context $\mathscr{C}$ and a mechanism M, $G = (\mathscr{C}, M)$, where $\mathscr{C}$ specifies the players, the outcomes, and the players' preferences over the outcomes,[2] while M specifies the available strategies and how they lead to outcomes. Our contexts and mechanisms are defined as follows.

[2]Since our solution concept is dominant strategies, we have no need to specify the players' beliefs.

3.1 Collusive Auction Contexts

A collusive auction context $\mathscr{C}$ has the following components

- The *players*, a finite set N identified with $\{1, \ldots, n\}$.
- The *number of copies* of the good for sale, a positive number m.
- The *outcomes*, the set of all pairs (A, P) where $A = (A_0, A_1, \ldots, A_n)$ is a *allocation*, a vector in $\{0, \ldots, m\}^{n+1}$ such that $\sum_i A_i = m$, and P a *price profile*, a profile of reals. For $i > 0$, A_i is the number of copies allocated to i, and A_0 is the number of copies left unallocated. Each P_i is i's price.
- The *valuation bound*, a number B upper bounding any player's marginal value for any copy.
- The *valuations*, the set of all functions $v : \{1, \ldots, m\} \to [0, B]$ such that $v(1) \geq \cdots \geq v(m)$.
- The *true valuations*, a valuation profile θ, where each $\theta_i(k)$ is i's marginal value for a kth copy.
- The *collusion structure*, a partition $\mathbb{C}$ of the players.

We refer to a player i as *independent* if $\{i\} \in \mathbb{C}$, and as *collusive* otherwise; and to a set $C \in \mathbb{C}$ as a *coalition*. Note that $\mathscr{C}$ is totally identified by just n, m, B, θ, and $\mathbb{C}$ alone: that is, $\mathscr{C} = (n, m, B, \theta, \mathbb{C})$.

The set of all (collusive auction) contexts with n players, m goods, and valuation bound B is $\mathscr{C}_{n,m,B}$.

Knowledge In a context $\mathscr{C} = (n, m, B, \theta, \mathbb{C})$, the values n, m, and B are common knowledge to all players. For each coalition $C \in \mathbb{C}$, the set C itself and the subprofile θ_C are common knowledge to the players in C. Nothing is assumed about any other knowledge (or belief) of the players.

A mechanism designer is assumed to know n, m, and B, but have no information about θ or $\mathbb{C}$.

Collective Utility, Collusive Rationality, and Social Welfare In a context $\mathscr{C} = (n, m, B, \theta, \mathbb{C})$, the *individual utility function* of a player i, u_i, and the *collective utility function* of a coalition $C \in \mathbb{C}$, u_C, are so defined: given an outcome (A, P),

$$u_i((A, P); \theta_i) \triangleq \sum_{k=1}^{A_i} \theta_i(k) - P_i \quad \text{and} \quad u_C((A, P); \theta_C) \triangleq \sum_{i \in C} u_i((A, P); \theta_i).$$

In any mechanism, each player i acts so as to maximize u_C if $i \in C \in \mathbb{C}$. For simplicity, we may omit the true valuation part and write the utilities as $u_i(A, P)$ and $u_C(A, P)$.

The *social welfare* of an allocation A (for $\mathscr{C}$), $SW(A)$, is $\sum_i \sum_{k=1}^{A_i} \theta_i(k)$. Allocation A is *efficient* if for any allocation A', $SW(A) \geq SW(A')$.

Remarks We find it convenient to make the collusive structure $\mathbb{C}$ part of the collusive context $\mathscr{C}$, but wish to emphasize that $\mathbb{C}$ may arise after a mechanism has been chosen. Indeed, in Theorem 1 "contexts are universally quantified after mechanisms."

We also find it simpler to assume that θ_C is common knowledge within a coalition C, but in our mechanism, for maximizing u_C it suffices that each player i in C knows his own valuation θ_i.

3.2 Collusive Auction Mechanisms

Definition A normal-form auction mechanism M for a context in $\mathscr{C}_{n,m,B}$ is *directly collusive* if the set of pure strategies for a player i, S_i, consists of the set of all pairs (v, C), where $v : \{1, \ldots, m\} \to [0, B]$ and $C \subseteq \{1, \ldots, n\}$.

If the mechanism M and the collusive context $\mathscr{C}$ under consideration are clear, and s and C respectively are a pure-strategy profile and a coalition, then $u_C(s)$ denotes the collective utility of C for the outcome produced by M under s.

Definition A directly collusive auction mechanism M for a collusive context $\mathscr{C} = (n, m, B, \theta, \mathbb{C})$ is *collusively dominant-strategy truthful* if, for all coalitions $C \in \mathbb{C}$ and all strategy subprofiles s_C and s_{-C},

$$u_C(t_C, s_{-C}) \geq u_C(s_C, s_{-C}),$$

where $t_i = (\theta_i, C)\ \forall i \in C$. We refer to the strategy profile t as *the truthful strategy profile* of auction $(\mathscr{C}, M)$.

We further say that M is *coalitionally rational* if, for all coalitions $C \in \mathbb{C}$ and all strategy subprofiles s_{-C},

$$u_C(t_C, s_{-C}) \geq 0.$$

Remark Because directly collusive mechanisms only specify strategies for *individual players,* the "collusive strategies" for coalition C are the subprofiles of strategies s_C, where $s_i \in S_i$ for each $i \in C$. Because an independent player i belongs only to the coalition $\{i\}$, coalitional rationality implies individual rationality for independent players. For coalitions of multiple players, given that their collective utility is non-negative when they bid truthfully, we leave it to the colluders to decide how to "split the proceeds."

4 Our Mechanism

We first recall Vickrey's mechanism, in a fashion convenient to pave the way for ours.

4.1 The Vickrey Mechanism (with Lexicographically Broken Ties)

The Vickrey mechanism for n players and m copies, $\textsc{Vick}_{n,m}$, is a direct mechanism not envisaging any collusion or any valuation upperbound B. Accordingly, $S_i = \{v : \{1, \ldots, m\} \to \mathbb{R}^+\}$, where $\mathbb{R}^+$ is the set of non-negative reals.

Given a profile of valuations V, $\textsc{Vick}_{n,m}$ orders the set of pairs $\{(i, V_i(k)) : i = 1, \ldots, n\ ;\ k = 1, \ldots, m\}$ according to the relation $>_V$ so defined: $(i, k) >_V (i', k')$ if and only if

$$V_i(k) > V_{i'}(k') \qquad \text{or} \qquad V_i(k) = V_{i'}(k') \text{ and } i < i' \qquad \text{or} \qquad V_i(k) = V_{i'}(k'),\ i = i', \text{ and } k < k'.$$

The first m pairs in the ordered sequences are called the "winning pairs", and the others the "losing pairs". Then, for each player i, $\textsc{Vick}_{n,m}$ allocates A_i copies of the good to i, where A_i is the number of winning pairs having i as their first components. Further, $\textsc{Vick}_{n,m}$ sets i's price, P_i, to 0 if $A_i = 0$, and to $\sum_{k=1}^{A_i} p_{i,k}$ otherwise, where $p_{i,k}$, *the price for i's kth copy of the good*, is the second component of the $(A_i - k + 1)$th losing pair whose first component is not i. (We assume that there are enough losing pairs, otherwise we could simply append sufficiently many $(0, 0)$ pairs to the end of the above ordered sequence.)

4.2 The Intuition Behind Our Mechanism

Our mechanism for n players and m copies with valuation upperbound B, $\mathcal{M}_{n,m,B}$ ($\mathcal{M}$ for short), allows each player to report not only his true valuation, but also the coalition he belongs to (if he is collusive). Assume for a moment that each player reports truthfully. Then, $\mathcal{M}$ considers each coalition C as a single *fictional* player, whose fictional valuation V_C is obtained by merging the valuations of C's members.[3] Having done this, $\mathcal{M}$ runs the Vickrey mechanism for m copies, with the fictional players and their fictional valuations, so as to compute an allocation of the copies of the good to the fictional players and the price paid by each fictional player. Having figured out this way that C should collectively receive —say— m_C copies and make a payment P_C, $\mathcal{M}$ must now decide how to distribute these copies and this payment to the members of C. Conceptually, $\mathcal{M}$ asks C how to do so. Practically, it uses the originally reported (non-fictional) valuations

[3] Notice that if there are m copies of the good for sale and C consists of c players, then C's fictional valuation will consist of a decreasing sequence of cm numbers. That is, the number of elements in V_C is larger than the number of copies at hand. We can however consider a valuation v as mapping $\{1, 2, \ldots, \}$, rather than $\{1, \ldots, m\}$, to $[0, B]$, where $v(k)$ is the marginal value of a kth copy of the good, *if available*. The Vickrey mechanism remains well-defined and dominant-strategy-truthful with such "infinite" valuations.

to compute the right answer. (It is quite clear that **if** the members of C report truthfully C, then it is in their interest to report truthfully also their individual valuations. Not doing so would be quite irrational of them, since they try to maximize the sum of their individual utilities.) As for prices, since again what matters is C's collective utility, $\mathcal{M}$ is at liberty of subdividing P_C in any way it wishes among the members of C. (For simplicity, we charge only the lexicographically first player in each coalition.)

Note that, if the coalitional components of the bids were ignored, and the Vickrey mechanism ran on just the reported valuations, then the produced allocation would be *the same* as that computed by $\mathcal{M}$ in the manner specified above, but the sum of the prices paid by the players in C might be *higher.* This is so because in $\mathcal{M}$ "the valuations of all the players in C are not used to set the price of a player i in C" while in the Vickrey case "only the valuation of i himself is not used to set i's price." In a sense, therefore, $\mathcal{M}$ guarantees efficiency while generating potentially less revenue: it gives a *discount* to coalitions of multiple players.

Discounts, of course, are attractive to every one. It is thus unclear whether it is dominant for —say— an independent player i to report truthfully $\{i\}$ as his coalitional component. Indeed, if a coalition C were somehow kind enough, in its report, to include him as a member (i.e., if C's members reported $C \cup \{i\}$ instead of C), then he might be better off reporting $C \cup \{i\}$ instead of $\{i\}$.

To guarantee collusive dominant-strategy truthfulness, $\mathcal{M}$ must therefore prevent such eventualities from being rationally entertained. Ideally, $\mathcal{M}$ should check the "consistency" of all reported coalitional components, and punish all "misreporters." The problem, however, is that this is not obviously implementable, and in a sense it is actually impossible.

For instance, if players a, b, c, and d form a coalition C, but a and b consistently report their coalition to be $\{a, b\}$, while c and d consistently report their coalition to be $\{c, d\}$, then there is no way for $\mathcal{M}$ to figure out that this is not true. The only thing that $\mathcal{M}$ can do, and actually does, is to ensure that it will not be in C's interest to "pretend to be two smaller coalitions", or a smaller coalition plus one or two independent players.

For some forms of inconsistent coalitional reporting, $\mathcal{M}$ identifies a guilty player and severely fines him. For others, it figures out a "correct coalitional report" and acts as if that were what the players actually reported.

For instance, if player a reports his coalition to be $\{a, b\}$; player b that his coalition is $\{a, b, c\}$; player c that her coalition is $\{b, c, d\}$; and player d that his coalition is $\{c, d\}$ (and if all other coalitional reporting seems to be in good order), then the mechanism proceeds to compute allocation, prices and fictional players as if a, b, c, and d all reported the coalition $\{a, b, c, d\}$, although such a coalition was not reported by anyone. Yet, the players should not worry: collusive dominant-strategy truthfulness will be guaranteed.

4.3 Formalization

Our mechanism is based on the following definition.

Definition Let $s = ((V_1, C_1), \ldots, (V_n, C_n))$ be a strategy profile in a directly collusive mechanism selling m copies of a single good to n players. A *disagreement* (in s) is an ordered pair of players (i, j) such that $C_i \ni j$ but $C_j \not\ni i$.

That is (i, j) is a disagreement if "i declares to collude with j, but j does not reciprocate."

In description below we denote by $x := y$ the operation that assigns value y to variable x.

Directly Collusive Mechanism $\mathcal{M}_{n,m,B}$

For all contexts in $\mathscr{C}_{n,m,B}$ and strategy profiles $s = ((V_1, C_1), \ldots, (V_n, C_n))$, compute an outcome (A, P) as follows.

(Punishing Procedure) *If there exists a disagreement, then (1) set $A_0 = m$ (i.e., allocate no copy), (2) initially set $P_x = 0$ for each player x, and then (3) make the following updates: for each disagreement (i, j),*

$$P_i := P_i + 2mB \quad \text{and} \quad P_j := P_j - mB \qquad \textit{(i.e., i pays a fine of mB to each of j and the seller).}$$

(Standard Procedure) *If there is no disagreement, then*

1. *Compute the player partition $\mathbb{P}$, consisting of the connected components of the directed graph having the players as nodes, and having a directed edge (i, j) whenever $j \in C_i$ —and thus $i \in C_j$.*
2. *Order the set $\{(i, V_i(k)) : i \in N, k = 1, \ldots, m\}$ according to $>_V$, denote by F_V^m the subsequence of the first m pairs, and then choose the allocation and price of each player i as follows:*
   - *If A_i pairs in F_V^m have i as their first component, then allocate A_i copies to i (as in Vickrey), and*
   - *If i is the lexicographically first player in a coalition $C \in \mathbb{P}$, set i's price to be*

   $$P_i = \begin{cases} 0 & \text{if } A_C = 0, \text{ where } A_C = \sum_{j \in C} A_j, \text{ and} \\ \sum_{k=1}^{A_C} p_{k,C} & \text{otherwise, where } p_{k,C} \text{ is the second component of the } (A_C - k + 1)\text{th pair} \\ & \text{not in } F_V^m \text{ with the first component not in } C. \end{cases}$$

   Otherwise set $P_i = 0$.

5 Analysis of Our Mechanism

Theorem 1. *For all n, m, and B, and for all collusive contexts $\mathscr{C}$ in $\mathscr{C}_{n,m,B}$,*

- *$\mathcal{M}_{n,m,B}$ is collusive dominant-strategy truthful,*
- *$\mathcal{M}_{n,m,B}$ is coalitionally rational, and*
- *$\mathcal{M}_{n,m,B}$ produces an efficient allocation under the truthful strategy profile.*

Proof. Let us more simply denote $\mathcal{M}_{n,m,B}$ by $\mathcal{M}$, and assume that $\mathcal{M}$ is played in a collusive context $(n, m, B, \theta, \mathbb{C})$. To prove that $\mathcal{M}$ is collusive dominant-strategy truthful, we first establish that, for the members of each coalition C, "bidding truthfully their coalitional component dominates bidding untruthfully", no matter what valuation component they might bid. Then, we shall establish that, "bidding truthfully both components dominates bidding truthfully only their coalitional component".

Claim 1. For all $C \in \mathbb{C}$, all subprofiles of valuations $(V_i)_{i \in C}$, all subprofiles of subsets of players $(\mathcal{C}_i)_{i \in C}$, and all strategy subprofiles s_{-C},

$$u_C(\mathbb{S}_C, s_{-C}; \theta_C) \geq u_C(\mathcal{S}_C, s_{-C}; \theta_C),$$

where $\mathbb{S}_C$ is the strategy subprofile $(V_i, C)_{i \in C}$ and $\mathcal{S}_C$ is the strategy subprofile $(V_i, \mathcal{C}_i)_{i \in C}$.

Proof of Claim 1. We consider the following four exhaustive cases.

Case 1: The Punishing Procedure is invoked for both $(\mathbb{S}_C, s_{-C})$ and $(\mathcal{S}_C, s_{-C})$

In this case, no copy is allocated, and thus C's utility coincides with $IN_C - OUT_C$, where IN_C is the sum of all fines paid to members of C by outside players, and OUT_C is the sum of all fines paid by members of C to outside players or the seller.

Recall that the punishing procedure is invoked only if there exists a disagreement (i, j) —that is, only if i declares j to collude with him, but j does not "reciprocate"— and that such a disagreement in particular results in i paying a fine of mB to each of j and the seller. Thus, under the strategy profile $(\mathbb{S}_C, s_{-C})$, $OUT_C = 0$; while under $(\mathcal{S}_C, s_{-C})$, $OUT_C \geq 0$. On the other hand, if such a disagreement (i, j) between an outsider i and a member j of C exists under $(\mathcal{S}_C, s_{-C})$, then it exists also under $(\mathbb{S}_C, s_{-C})$. Thus the total value of IN_C under $(\mathbb{S}_C, s_{-C})$ is greater than or equal to its value under $(\mathcal{S}_C, s_{-C})$.

Accordingly Claim 1 holds in Case 1.

Case 2: The Punishing Procedure is invoked for $(\mathcal{S}_C, s_{-C})$ but not for $(\mathbb{S}_C, s_{-C})$

In this case, under $(\mathcal{S}_C, s_{-C})$ there is at least one disagreement. Note that, for each such a disagreement (i, j) it cannot be $i \notin C$, otherwise (i, j) would be a disagreement under $(\mathbb{S}_C, s_{-C})$ too. Thus, i belongs to C and pays a fine of mB to each of j and the seller. Accordingly, the collective utility of C decreases by at least mB for each disagreement, implying that $u_C(\mathcal{S}_C, s_{-C}) \le -mB$, because in the case under consideration all utilities are determined by fines.

On the other side, under the standard procedure all utilities are due to the allocation of the copies and the prices paid for them. And whenever a member j of C receives his kth copy of the good, the collective utility of C increases by $\theta_j(k) \ge 0$ and decreases by $p_{j,k}$, where $p_{j,k} \le V_j(k) \le B$. Since there are m copies altogether, we have $u_C(\mathbb{S}_C, s_{-C}) \ge -mB$.

Accordingly Claim 1 holds also in Case 2.

Case 3: The Punishing Procedure is invoked for $(\mathbb{S}_C, s_{-C})$ but not for $(\mathcal{S}_C, s_{-C})$

In this case, under $(\mathbb{S}_C, s_{-C})$ there must exist a disagreement. Moreover, for each such disagreement (i, j), we must have (a) $i \notin C$, because the coalition component of $\mathbb{S}_x$ is C for each $x \in C$, and (b) $j \in C$, because otherwise (i, j) would be a disagreement also under $(\mathcal{S}_C, s_{-C})$, contrary to our hypothesis. Therefore, for each such disagreement, $i \notin C$ pays $j \in C$ a fine of mB, implying that $u_C(\mathbb{S}_C, s_{-C}) \ge mB$.

On the other side, because under the standard procedure all utilities are due to allocation and payments, and because whenever a player $i \in C$ gets his kth copy the value received is $\theta_i(k) \le B$ and the payment is $p_{i,k} \ge 0$, we have $u_C(\mathcal{S}_C, s_{-C}) \le mB$.

Accordingly, Claim 1 holds in Case 3 too.

Case 4: The Standard Procedure is invoked for both $(\mathbb{S}_C, s_{-C})$ and $(\mathcal{S}_C, s_{-C})$

In this case, all copies are allocated, both under $(\mathbb{S}_C, s_{-C})$ and $(\mathcal{S}_C, s_{-C})$. Moreover, since the underlying valuation profile V is the same for both strategy profiles, so is the allocation, A. However, the player partitions computed by $\mathcal{M}$, respectively denoted by $\mathbb{P}$ and $\mathcal{P}$, may be different: $\mathbb{P}$ and $\mathcal{P}$ must coincide outside C, but while $C \in \mathbb{P}$, C may not be a member of $\mathcal{P}$. However, if $C \notin \mathcal{P}$, then $C^1, \ldots, C^\ell \in \mathcal{P}$, where $\{C^1, \ldots, C^\ell\}$ is a partition of C. (This is so because otherwise a disagreement should have existed under either $(\mathbb{S}_C, s_{-C})$ or $(\mathcal{S}_C, s_{-C})$.)

Let $A_C = \sum_{i \in C} A_i$ and $A_{C^j} = \sum_{i \in C^j} A_i$ for each $j \le \ell$. If $A_C \ne 0$, then for the kth copy allocated to C, without loss of generality assume that it is allocated to player $i \in C^j$ and is the k'th copy allocated to C^j. We have that $A_C - k \ge A_{C^j} - k'$, because the set of copies allocated to C after this copy contains the set of copies allocated to C^j after this copy. Letting the prices charged for this copy respectively be $p_{k,C}$ and p_{k',C^j}, we have that

- $p_{k,C}$ is the second component of the $(A_C - k + 1)$th pair not in F_V^m with the first component not in C, and
- p_{k',C^j} is the second component of the $(A_{C^j} - k' + 1)$th pair not in F_V^m with the first component not in C^j.

Thus, $p_{k,C} \le p_{k',C^j}$, and the collusive price of C is no greater under $(\mathbb{S}_C, s_{-C})$ than under $(\mathcal{S}_C, s_{-C})$.

Accordingly, Claim 1 holds in this last case too, and is therefore proven. $\square$

Claim 2. For all $C \in \mathbb{C}$, all valuation subprofiles $(V_i)_{i \in C}$, and all strategy subprofiles s_{-C},

$$u_C(t_C, s_{-C}; \theta_C) \ge u_C(\mathbb{S}_C, s_{-C}; \theta_C),$$

where t_C is the truthful strategy subprofile $(\theta_i, C)_{i \in C}$, and $\mathbb{S}_C$ is the strategy subprofile $(V_i, C)_{i \in C}$.

Proof of Claim 2. This claim is indeed an immediate corollary of the dominant strategy truthfulness of the Vickrey mechanism, and of the fact that, when all players in C truthfully reveal C, our mechanism either (1) invokes the Punishing Procedure under both strategy profiles, so that IN_C is the same under both profiles and $OUT_C = 0$ always, or (2) invokes the Standard Procedure and essentially treats C as a single player in the Vickrey mechanism. $\square$

Claims 1 and 2 together imply that for all $C \in \mathbb{C}$, all strategy subprofiles s_C, and all strategy subprofiles s_{-C}, writing s_i as $(V_i, \mathcal{C}_i)$ for each $i \in C$, we have

$$u_C(t_C, s_{-C}; \theta_C) \geq u_C((V_i, C)_{i \in C}, s_{-C}; \theta_C) \geq u_C((V_i, \mathcal{C}_i)_{i \in C}, s_{-C}; \theta_C) = u_C(s_C, s_{-C}; \theta_C),$$

concluding the proof that $\mathcal{M}$ is collusive dominant-strategy truthful.

Next, let us point out that $\mathcal{M}$ is coalitionally rational. Indeed, for each coalition $C \in \mathbb{C}$, when its members bid truthfully, no matter what the other players do, either (1) the punishing procedure is invoked and $OUT_C = 0$, or (2) the standard procedure is invoked so as to produce an allocation A and the price charged to C (via its lexicographically first member) does not exceeds $\sum_{i \in C} \sum_{k=1}^{A_i} \theta_i(k)$. In either case therefore, the collective utility of C is non-negative.

Finally, let us point out that $\mathcal{M}$ returns an allocation maximizing social welfare at the truthful strategy profile. Indeed, when all players bid truthfully, (1) no disagreement exists, and thus $\mathcal{M}$ does not execute the punishing procedure, and (2) the standard procedure returns the same allocation A as the Vickrey mechanism does under the true-valuation profile, and thus A has maximum social welfare. $\blacksquare$

6 Applying Our Mechanism to Our Starting Example

Formal analysis notwithstanding, let us consider the auction context discussed in the example of Section 1 and see why our mechanism guarantees maximum social welfare even when the players have wrong beliefs. Recall that in our previous example, players a and b form a coalition, each of them has $(100, 0)$ as his true valuation, and they wrongly believe $(1000, 0)$ and $(x, 0)$ —with $x < 100$— to be the respective valuations of players c and d, where instead the true valuation of each of them is $(1, 0)$.

Given such beliefs, in our mechanism, having both a and b announcing $((100, 0), \{a, b\})$ turns out to be a best response for the coalition. Indeed, from the point of view of the coalition, this "joint" strategy ensures that player a gets a copy of the good and pays a price x: realizing that a and b declare themselves as a coalition, our mechanism will not set player a's price based on player b's announced valuation. On input the bids $((100, 0), \{a, b\})$, $((100, 0), \{a, b\})$, $((1, 0), \{c\})$, and $((1, 0), \{d\})$, our mechanism allocates one copy of the good to player a and the other to b, maximizing social welfare.

More generally, in our mechanism, differently from in Vickrey's mechanism, no player ever wants to underbid his own valuation, for all possible beliefs and all possible collusion structures.

7 Final Remarks

Collusion is only too human, and to develop a realistic theory of human interactions it will be crucial to consider mechanisms and solution concepts resilient to collusion. In this paper we have put forward such a solution concept, and showed its usefulness for auctions of a single good in limited supply. We do not know, however, how to export our results to other settings, such as combinatorial auctions and provision of a public good. Perhaps, one must be content with developing weaker notions of collusion resiliency for such settings. Yet, we should always guarantee stronger protection against collusion whenever we can.

References

[1] D. Abreu, H. Matsushima, Virtual implementation in iteratively undominated strategies: complete information, Econometrica 60(5) (1992), 993-1008.

[2] L.M. Ausubel, P. Milgrom, The lovely but lonely Vickrey auction, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), Combinatorial Auctions, MIT Press, 2006, pp. 17-40.

[3] S. Barbera, M.O. Jackson, Strategy-proof exchange, Econometrica 63(1) (1995), 51-87.

[4] Y. Che, J. Kim, Robustly collusion-proof implementation, Econometrica 74(4) (2006), 1063-1107.

[5] J. Chen, S. Micali, P. Valiant, Leveraging collusion in combinatorial auctions, Innovations in Computer Science (2010), 81-93.

[6] E.H. Clarke, Multipart pricing of public goods, Public Choice 11(1) (1971), 17-33.

[7] J. Feigenbaum, C.H. Papadimitriou, S. Shenker, Sharing the cost of multicast transmissions, Journal of Computer and System Sciences 63 (2001), 21-41.

[8] J. Glazer, M. Perry, Virtual implementation in backwards induction, Games Econ. Behav. 15 (1996), 27-32.

[9] A. Goldberg, J. Hartline, Collusion-resistant mechanisms for single-parameter agents, Symposium on Discrete Algorithms (2005), 620-629.

[10] J. Green, J. Laffont, On coalition incentive compatibility, Rev. Econ. Stud. 46(2) (1979), 243-254.

[11] T. Groves, Incentives in teams, Econometrica 41(4) (1973), 617-631.

[12] M.O. Jackson, T. Palfrey, S. Srivastava, Undominated Nash implementation in bounded mechanisms, Games Econ. Behav. 6 (1994), 474-501.

[13] K. Jain, V. Vazirani, Applications of approximation algorithms to cooperative games, Symposium on Theory of Computing (2001), 364-372.

[14] J. Laffont, D. Martimort, Mechanism design with collusion and correlation, Econometrica 68(2) (2000), 309-342.

[15] E. Maskin, Implementation and strong Nash-equilibrium, in: J.J. Laffont (Ed.), Aggregation and Revelation of Preferences, North-Holland, Amsterdam, 1979, pp. 433-440.

[16] H. Mizukami, On the constancy of bribe-proof solutions, Econ. Theory 22(1) (2003), 211-217.

[17] H. Moulin, Incremental cost sharing: characterization by coalition strategy-proofness, Soc. Choice Welfare 16(2) (1999), 279-320.

[18] H. Moulin, B. Peleg, Cores of effectivity functions and implementation theory, J. Math. Econ. 10 (1982), 115-145.

[19] H. Moulin, S. Shenker, Strategyproof sharing of submodular costs: budget balance versus efficiency, Econ. Theory 18(3) (2001), 511-533.

[20] M. Pal, É. Tardos, Group strategyproof mechanisms via primal-dual algorithms, Symposium on Foundations of Computer Science (2003), 584-593.

[21] J. Schummer, Manipulation through bribes, J. Econ. Theory 91(2) (2000), 180-198.

[22] S. Suh, Implementation with coalition formation: a complete characterization, J. Math. Econ. 26 (1996), 409-428.

[23] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, J. Finance 16(1) (1961), 8-37.