

Robustly Leveraging Collusion in Combinatorial Auctions*

Jing Chen¹ Silvio Micali¹ Paul Valiant²

¹CSAIL, MIT, Cambridge, MA 02139, USA

²Computer Science Division, UC Berkeley, Berkeley, CA 94720, USA

jingchen@csail.mit.edu silvio@csail.mit.edu pvaliant@mit.edu

Abstract: Because of its devastating effects in auctions and other mechanisms, collusion is prohibited and legally prosecuted. Yet, colluders have always existed, and may continue to exist. We thus raise the following question for mechanism design:

What desiderata are achievable, and by what type of mechanisms, when any set of players who wish to collude are free to do so without any restrictions on the way in which they cooperate and coordinate their actions?

In response to this question we put forward and exemplify the notion of a *collusion-leveraging mechanism*. In essence, this is a mechanism aligning its desiderata with the incentives of all its players, including colluders, to a significant and mutually beneficial extent. Of course such mechanisms may exist only for suitable desiderata.

In unrestricted combinatorial auctions, where classical mechanisms essentially guarantee 0 social welfare and 0 revenue in the presence of just two colluders, we prove that it is possible for collusion-leveraging mechanisms to guarantee that the sum of social welfare and revenue is always high, even when all players are collusive.

To guarantee better performance, collusion-leveraging mechanisms in essence “welcome” collusive players, rather than pretending they do not exist, raising a host of *new* questions at the intersection of cooperative and non-cooperative game theory.

Keywords: mechanism design; collusion leveraging; player-knowledge benchmark; and combinatorial auction

1 Introduction

Collusion is a major problem for traditional mechanisms for a very simple and fundamental reason. Traditional mechanism design guarantees a desired property \mathbb{P} at equilibrium. But, by definition, an equilibrium only guarantees that no *individual* player has incentive to deviate from his envisaged strategy, while two or more players may have plenty of incentive to coordinate a *joint* deviation. And when they do so in the course of a mechanism, the desired property \mathbb{P} typically no longer holds. The problem of collusion is both particularly acute and well documented in auctions. Both physical and legal protection against it are routinely employed: auction rooms are often monitored by a variety of surveillance equipment, and collusion is outlawed and criminally punished. But with limited results.

In this paper we thus put forward a new and purely mathematical approach to collusion in combinatorial auctions.

1.1 Prior Work

Restricted Collusion and Restricted Auctions

Some protection against collusion can be obtained by starting with the assumption of some *restriction* on the coordination ability of colluders. For instance, *group strategy-proof* (or equivalently, *coalition strategy-proof*) mechanisms [2, 14, 19–21], work under the assumption that colluders are incapable of making side-payments to each other. Alternatively, some collusion protection can be obtained for restricted auctions: in particular, single-parameter auctions [10]. (Some collusion protection is also available for other restricted games, such as with two players of two possible types, or Bayesian games, where additional information about the players is available to the mechanism designer. See [5, 6, 15, 16].) But all such protection vanishes when the colluders’ coordination is *unrestricted*, the auction is *combinatorial*, and the mechanism designer knows *nothing* about the players.

Combinatorial Auctions and the Ausubel-Milgrom Example

In auctions of multiple goods, each player i has a true valuation TV_i for the goods for sale: a function specifying i ’s true value $TV_i(S)$ for each possible sub-

*Work done when all three authors were at the Computer Science and Artificial Intelligence Laboratory at MIT

set S of the goods. Such an auction is called *combinatorial* when the players' valuations are arbitrary and unrelated functions. Combinatorial auctions are therefore the most general form of auctions, but also the most difficult one when it comes to collusion. In fact, their rich structure can be easily exploited by collusive players. Notably, Ausubel and Milgrom [1] have shown that just two (sufficiently informed) collusive players may drive to 0 the *social welfare* (as well as the *revenue!*) of the famous VCG mechanism. This is so despite the fact that the VCG is dominant-strategy truthful, in essence the best form of equilibrium, and that at equilibrium it maximizes social welfare.

Implementation in Undominated Strategies and Rationally Robust Implementation

The classical notion of *implementation in undominated strategies* [13], and its *feasible* version [3], although not applied to unrestricted combinatorial auctions, are ancestors of *rationally robust implementation*, a notion put forward by [7, 8], and adopted in this paper as our solution concept. Rationally robust implementation is recalled in Section 3, but its zest is first best conveyed by lying as follows: a mechanism provides a rationally robust implementation of a given property \mathbb{P} if it guarantees \mathbb{P} not at an equilibrium, but at any profile of strategies surviving iterated elimination of strictly dominated strategies.

Robust leveraging of external independent knowledge

Traditional mechanisms leverage only the *internal* knowledge of the players. In an auction, this would be the knowledge that each player i has of his own true valuation TV_i . However, very little revenue can be guaranteed by traditional mechanisms in combinatorial auctions, with or without collusion [17]. Any hope to guarantee more revenue (without assuming that the seller/designer has some convenient knowledge about the players, such as some suitable Bayesian information) rests on a mechanism's ability to leverage also the players' *external* knowledge. In an auction, this is essentially the knowledge that each player i has about the others' valuations. Quite realistically, in this paper we work with the original, imperfect external knowledge model of [7]: *guaranteed* (or *lower-bounded*) external knowledge. In essence,

Each i knows a lower-bound, $V_{j,S}^i$, for each $TV_j(S)$.

Notice that such external knowledge is *not* an assumption, since at worst $V_{j,S}^i$ could be 0. The mechanism of [7] leverages this external knowledge in a combinatorial auction in a very robust way. Namely, no matter how many collusive players there may be, no matter how many secret coalitions they may be partitioned

in, and no matter how the members of each coalition may coordinate their actions, the revenue of their rationally robust implementation is always greater than or equal to

$$1/2 \text{ of MIEW} = \max_{i \in I} \text{MIEW}_i, \text{ where } \text{MIEW}_i \text{ is the maximum external welfare known to an independent } i.^1$$

Of course, the more precise the external knowledge of the players, the better the performance one could guarantee. As shown in another paper [4], to appear in ICS 2010, when the players' external knowledge is perfect, one can guarantee perfect revenue too, even in a dreadfully collusive setting.

Note that the external welfare known to a player i can be interpreted as the best way known to i to sell the goods to the other players. Since the seller and/or the mechanism designer is assumed to know nothing, and is thus less informed than any of the players, being able to sell the goods roughly as well as some of the players could —let alone the “best-informed” independent player!— is a non-trivial guarantee.

1.2 Our Work

As we have seen, *coordinated* collusive players constitute a major obstacle to mechanism design in general and to combinatorial auctions in particular. As we have seen too, all work so far has focussed on *preventing* collusion from damaging an auction, either by trying to (1) “force” collusive players to behave independently [2, 10, 14–16, 19–21], or (2) “neutralize” collusive players from the auction [7, 17]. In this paper we put forward a more ambitious question:

Is it possible for a mechanism to leverage collusion?

We believe this question to be central to mechanism design. If we really want to leverage the players' knowledge, then we should be able to treat colluders as a potential reservoir of knowledge to be harvested. To explain both what our question means and what we can prove about it, we need to informally clarify a few things: our collusion model, our solution concept, the property we strive to achieve, the knowledge we try to leverage, our benchmark, our notion of collusion leveraging, and then the extent to which we can provably leverage collusion.

Rational, unrestricted, dynamic, unpunishable, and secret collusion

Mechanism design relies on the players' rationality, and for it to leverage also the knowledge of collusive

¹That is, I stands for the set of independent (i.e., not collusive) players, and MIEW_i is the maximum of $\sum_{j \neq i} V_{j,A_j}^i$ taken over all partitions A of the goods among the players, where A_j denotes the set of the goods that A assigns to player j .

players, coalitions of players must be rational too. In this paper we assume that a rational coalition is a subset of the players coordinating their actions so as to maximize the sum of the (individual) utilities of its members. Perhaps other models of rational coalitions can be analyzed in the future. But if we want to understand collusion leveraging, we have to start somewhere. And ours is not a random place to start, for two main reasons.

- (a) *“Maximizing the money coming in” is the best way for collusive players to enrich themselves.* This is important because players collude in order to further improve their individual utilities. Of course, different members of a coalition may have different bargaining powers, and any collusive gain might ultimately be split in different proportions. But limiting the amount of money coming in never is the rational thing to do for a coalition!
- (b) *It is the traditional model.* Indeed, most of the papers that need to specify a “joint utility function” for a coalition (e.g., [10]) adopt the same model.

In all other respects, our collusion model is totally unrestricted. In particular,

- *No player is afraid to collude.* (Even if collusion is severely punished, we model the players as believing with probability 1 that they will never be caught.)
- *The composition, size, and total number of coalitions is totally unrestricted.* (All players belonging to the same coalition is not ruled out. Coalitions of size 1 correspond to players who have chosen to remain independent.)
- *Members of the same coalition can cooperate in any way they want.* (In particular, they could make side-payments to one another, or enter contracts with each other that are perfectly binding —possibly with respect to quite different “enforcement systems.”)
- *Coalitions may be secret.* The members of a coalition are perfectly capable of keeping its existence secret, if this is to their advantage.
- *Coalitions can form dynamically.* Of course, coalitions may pre-exist the choice of a mechanism. (E.g., husband and wife, or brother and sister, may have decided to collude in any case.) But we want to protect even against a more dangerous case. Namely, we let the players choose, if they so wish, to form coalitions by means of the following 3-stage process: (1) All players are initially independent; (2) A specific mechanism

is announced, and then (3) The players partition themselves into coalitions in any way they want.

Note: In this paper, we do not specify the process of coalition formation. Indeed, it is a strength of our mechanism \mathcal{M} that it works *no matter how coalitions are formed*. But it is important to point out that our \mathcal{M} can handle dynamic coalitions. In fact, it should be appreciated that any mechanism leveraging dynamic coalitions also leverages “static” ones, while the viceversa needs not to hold.

Our model thus has two noteworthy consequences.

1. *Whether the players possess “the means to collude” is not an issue.* We view this as no big loss. Realistically, with the advent of modern communication networks, an auctioneer’s ability to credibly deny his players all means of colluding is vanishing fast anyway.
2. *The “Law” is no longer a credible ally.* Whether we like to admit it or not, our legal system has been aiding mechanism design in several ways. In particular, it has boosted the meaningfulness of equilibria: the law takes care of multi-player deviations, leaving only single-player deviations to be dealt with by mathematical analysis. But as mechanisms start being played over the Internet, legal help is vanishing too. If a combinatorial auction is conducted over the Internet, who has proper jurisdiction? Even if players were required to make high “safety deposits” in *our own* country (so as to vouch for our ability to punish them and to enforce the final outcome), and even if we clarified which countries have jurisdiction over which players, collusion should continue to worry us. Countries tolerating mass murderers may not care about energetically prosecuting colluders. Accordingly, we are seeking to address collusion by mechanisms relying only on *Mathematics*, rather than, explicitly or implicitly, a “combination” of *Mathematics* and *police/jail/torture/execution/et cetera*.

In sum, our chosen approach is of a safe, Machiavelian realism: namely, *any set of players who wishes to collude, does*. If a coalition does not come into being it is only because its potential members found more profitable alternatives, or because they could not bargain successfully and failed to reach agreement on how to split their potential gains.

Rational Play

As already said, we adopt rationally robust implementation as our solution concept. A bit more precisely, this notion of implementation guarantees a

property \mathbb{P} by identifying (1) a mechanism M and (2) a corresponding *refined* subset of strategies S_x for each *agent* (player or coalition) x such that

- If everybody is rational, each x will never want to choose a strategy outside S_x ; and
- No matter what strategy in S_x each x actually chooses, \mathbb{P} is guaranteed to hold.

Even this sketchy summary makes it clear that rationally robust implementation does not rely on equilibria. (Indeed, unless each subset of refined strategies has cardinality 1, an arbitrary profile of refined strategies may not be an equilibrium.) More generally, rationally robust implementation does not rely on the players' *beliefs* on how the mechanism will be played. Indeed, it is *robust*.

Total Performance

Traditional auctions are designed to maximize either *social welfare* or *revenue* (i.e., either the sum of the players' true values for the goods allocated to them, or the sum of the prices paid by the players). Our goal is to maximize *total performance*, that is, the sum of the two. There are compelling reasons for choosing this goal.

1. *It is an achievable goal.* As self-serving as this may sound at a superficial level, we note that, in the presence of collusive players capable of cooperating without restriction, it is a *non-trivial goal*.²
2. *It is a natural goal.* If we were guaranteed, by some means, that there will be no collusion in our auction, we would be only too happy to run the VCG mechanism and generate perfect social welfare. But as already mentioned, it was insightfully shown by Ausubel and Milgrom [1] that in the VCG mechanism two collusive players who do not value the goods very much can bid very high and get all goods while paying nothing, thereby destroying both social welfare and revenue. In light of their example, sacrificing some potential social welfare and converting it to revenue is a quite natural antidote to collusion. Indeed, we *do not prevent* collusive or independent players who value the goods very low from bidding very high and getting all the goods, but we *do guarantee* that by so doing they will pay through their noses.

²In particular, the total performance of a mechanism M designed to guarantee as much revenue as possible in the presence of collusive players may be quite poor. This is so because M may only yield modest revenue while sacrificing social welfare a lot, so that the total performance of M may be just twice a modest revenue.

3. *It is a desirable goal.* A traditional motivation behind the maximization of social welfare is that of a benevolent government, solely interested in the happiness of its citizens, rather than in revenue. To be sure, the VCG mechanism perfectly achieves this classical goal by imposing prices to the players. But such prices are almost an "afterthought," or a "necessary evil": they are just a means to maximize social welfare. But what is wrong with revenue? A benevolent government transforms it into roads, hospitals and other infrastructure from which everyone benefits. Taking this point of view, maximizing the sum of revenue and social welfare in the presence of collusion is a more meaningful goal for a benevolent government.

Independent of the above reasons, revenue alone cannot be meaningfully pursued in our setting. When all players are allowed to collude without restriction and the seller is not assumed to have any suitable (e.g., Bayesian) knowledge about them, no meaningful revenue-only benchmark can be guaranteed. For example, if all players in an auction collude together, then it is reasonably clear that no constant fraction of their value for the items may be extracted as revenue, since the mechanism essentially has to accept any price the coalition names for themselves.

Knowledge Model

In our combinatorial auctions we adopt the knowledge model of [7]. Again, this means that each player i not only knows his own true valuation, but, without any loss of generality, also a (possibly trivial) lower-bound on the other players' valuations. That is, the *guaranteed knowledge* of each i consists of a valuation profile K^i such that (1) $K^i_i = TV_i$ and (2) for all other players j and all subsets of the goods S , $0 \leq K^i_j(S) \leq TV_j(S)$.³

The following "union" operation on such guaranteed knowledge is crucial for us.

Definition 1. *If K is a guaranteed-knowledge profile and C a subset of the players, then K^C denotes the valuation profile such that, for any player i and any subset S of the goods, $K^C_i(S) = \max_{j \in C} K^j_i(S)$.*

In essence, K^C is "the most accurate guaranteed knowledge that the players in C could compute after

³Again following [7], we stress that K^i is not the only thing i knows. And indeed a player is free to use any additional knowledge when playing our mechanisms. However, if a mechanism is capable of leveraging the players' guaranteed knowledge, it is able to do so no matter what additional knowledge the player may have.

truthfully sharing their individual guaranteed knowledge." Notice that K_i^C coincides with TV_i for any member i of C .

Knowledge-Monotone Benchmarks

A guaranteed-knowledge benchmark is a function B mapping any possible guaranteed knowledge profile to a non-negative real number. For the sake of meaningfulness, we focus solely on *knowledge-monotone* benchmarks: that is, we demand that "the better the knowledge of the players, the better the mechanism's performance." A bit more formally, we impose a partial order on guaranteed knowledge as follows.

Definition 2. For any guaranteed knowledge K and \hat{K} we say that $K \geq \hat{K}$ if $K_j^i(S) \geq \hat{K}_j^i(S)$ for all players i and j and any subset S of the goods. We say that a guaranteed-knowledge benchmark B is *knowledge-monotone* if $B(K) \geq B(\hat{K})$ whenever $K \geq \hat{K}$.

Our Benchmark

Recall that the maximum social welfare of a valuation profile V , $MSW(V)$, is the maximum of $\sum_j V_j(A_j)$, taken over all partitions A of the goods among the players —where A_j denotes the set of the goods that A assigns to player j . Let us now define the characteristic benchmark of this paper.

Definition 3. (Maximum Known Welfare)

Letting $\text{MKW}_i(K) = MSW(K^i)$ for each guaranteed-knowledge profile K and player i , we define the maximum known welfare benchmark, MKW , as follows:

$$\text{MKW}(K) = \max_i \text{MKW}_i(K).$$

More generally, for any subset S of the players, we define $\text{MKW}_S = \max_{i \in S} \text{MKW}_i$.

Note that MKW indeed is a knowledge-monotone benchmark. Note too that each MKW_i consists of the maximum social welfare when the players' true valuations are precisely as in K^i , and thus consists of the maximum social welfare i knows he can guarantee if he were in charge of assigning the goods. Accordingly, MKW consists of "the maximum social welfare that the best-informed player knows how to guarantee." Since, following the purest form of mechanism design, we assume that all knowledge lies with the players (and none with the designer), MKW is a non-trivial benchmark, and achieving it within a constant factor (as we do) is significant. Of course we could conceive and construct auction instances whose maximum known welfare is rather low. But this is missing the point. When all knowledge lies with the players,

Enabling an ignorant seller to assign the goods roughly as well as the best informed player is an attractive goal.

(To be sure, player-knowledge benchmarks have generated some common confusion. At least some of it has been clarified in Section 5.1 of [8].)

From MIEW to MKW in a Collusion-Resilient Way

Note that MKW is a benchmark more demanding than MIEW . Indeed, MIEW is only defined over the external knowledge of independent players. By contrast, MKW allows any player i to assign goods to any player, including himself. Thus, MKW captures the *total* (i.e., both internal and external) relevant knowledge of *all players* (whether independent or collusive).

However, if we are satisfied to leverage just the knowledge of the independent players, then the two benchmarks can be related in various ways. In particular, the following holds. For any c between 0 and 1, one can easily transform a collusion-resilient mechanism M guaranteeing revenue $\geq c\text{MIEW}$ into a collusion-resilient mechanism M' guaranteeing (1) a total performance $\geq \frac{c}{c+1}\text{MKW}_I$, where I is the set of independent players, and (2) revenue greater than or equal to a fraction $\frac{c}{c+1}$ of the total knowledge of the "second-best-informed independent player." Essentially, M' runs M with probability $\frac{1}{c+1}$ and a "second-price" auction A' with complementary probability. In such an A' , each player bids a value together with a subset of the goods. The winner is the player with the highest value. He pays the second highest value, and gets the subset specified in his bid. All other goods remain unallocated, and all other players pay nothing. In particular, we can transform the mechanism in [7] to a new one guaranteeing a total performance $\geq \frac{\text{MKW}_I}{3}$.

Indeed, *collusion resiliency* is quite different from and quite simpler than *collusion leveraging*. The main point of this paper is not to leverage the total knowledge of just the independent players, but that of all players. Let us thus see what this should mean.

Collusion Leveraging

A basic way to express that a mechanism M achieves a fraction c of MKW is provided by the following property:

In every rational play with guaranteed knowledge K , the total performance of M is $\geq c \cdot \text{MKW}(K)$.

Let us now put forward a more demanding way to express that M achieves a fraction c of MKW (or any other knowledge-monotone benchmark) in a dynamic collusion model. Recall that such a model envisages a multi-stage process: in the initial stage, each player is

independent and has his own knowledge; in the second stage, M is announced; in the third stage the players partition themselves into collusive sets as they see fit; and finally M is played. Recall too that members of the same collusive set cooperate so as to maximize the sum of their individual utilities. To this end, they may need to share some of their knowledge. Accordingly, the guaranteed knowledge profile in the first and third stages may be quite different.

Definition 4. We say that M is a collusion-leveraging mechanism with total performance $c\text{MKW}$ if

$$\begin{aligned} & \text{in any rational play with initial guaranteed} \\ & \text{knowledge } K, M \text{ 's total performance is} \\ & \geq c \cdot \text{MKW}(\overline{K}) \end{aligned}$$

where \overline{K} is the (fictitious) knowledge profile such that, for every coalition C in the final stage, $\overline{K}^i = K^C \forall i \in C$, and $\overline{K}^i = K^i$ for all independent players i .

(It is actually possible to formally strengthen—and achieve in our case— collusion leveraging by mandating another property: *collusion rewarding*. Informally, a mechanism should make it preferable—subject only to the ability to agree on how to split the proceeds—for any subset of players to collude.)

REMARKS

- By knowledge monotonicity, $\text{MKW}(\overline{K}) \geq \text{MKW}(K)$. That is, our benchmark can only go up when players collude.
- Collusion leveraging does *not* demand that members of the same coalition share their knowledge. Rather, it states that such members *ultimately* behave *as if* they shared their knowledge. In particular, the members of a coalition might choose their best joint strategies via a secure multi-party computation, in which each one of them uses his own true knowledge as his own private input. This way, they are able to *de facto* choose their best strategies while preserving the privacy of their individual knowledge to the maximum possible extent, and thus without “sharing all of it” in any reasonable sense of the term.
- Collusion leveraging is a goal beyond those considered in the past. In our terminology, the traditional effort was directed either at preventing collusion (i.e., to achieve a desired benchmark evaluated at K , rather than at \overline{K}) or at neutralizing collusion (i.e., to achieve a desired benchmark evaluated at the subprofile of K cor-

responding to the independent players in the final stage, rather than at the full \overline{K}).

- Because the benchmark of a collusion-leveraging mechanism increases with collusion, such a mechanism might as well explicitly envisage the presence of collusive players. Indeed, our mechanism \mathcal{M} of Section 4 goes as far as making special “collusive strategies” available to coalitions of players. In some sense, therefore,

Our mechanism \mathcal{M} is at the intersection of cooperative and non-cooperative game theory.

- By providing strategies for collusive players, our \mathcal{M} de facto assumes that collusion is *legalized*. Despite going against a long tradition, this choice is quite logical in our adversarial collusive setting. Indeed, if coalitions can form whenever the players want,

Insisting on mechanisms envisaging only independent players is counter-productive.

Such insistence only ties the hands of the mechanism designer, and thus ultimately hurts performance!

Our Main Result

The main result of our paper is the following.

Informal Thm 1: *There exists a collusion-leveraging combinatorial-auction mechanism \mathcal{M} whose total performance is $\frac{\text{MKW}}{6}$.*⁴

Open Questions

Our paper raises a totally new class of questions. In particular,

- While we consider benchmarks of strictly increasing meaningfulness, the fraction of them we are able to achieve strictly decreases: namely,
 - 1/2 of the maximum external welfare known to any independent player (i.e., [7]);
 - 1/3 of the maximum total welfare known to any independent player (as discussed above); and
 - 1/6 of the maximum total welfare known to any player (i.e., Theorem 1).
 Does this “anti-correlation” arise intrinsically, or is it due to our currently poor tools in a new environment?
- Eric Maskin (private communication) has asked whether it would be possible to handle functions of social welfare and revenue more sophisticated

⁴By slightly changing our mechanism and complicating its analysis, we can improve the total performance of our mechanism to $\lfloor \frac{\text{MKW}}{1+\sqrt{2}} \rfloor / 2$.

than total performance. For instance, can we engineer mechanisms so as to “initially” privilege revenue, and then (i.e., after “enough” revenue has been generated) social welfare? Good question, but we are not ready for it yet!

- Could we guarantee better performance if “better knowledge” (e.g., a mixture of guaranteed and Bayesian knowledge) were available? Here, in line with mechanism design in its purest form, we mean that more accurate knowledge is available to the players, not to the seller/designer!
- Are there tight impossibility results for collusion leveraging? What are the right structural results for rationally robust implementation? Can we leverage collusion to a larger extent by better understanding “collusion formation?”

In sum, there is a lot more to understand and much more work to look forward to!

1.3 Our Related Forthcoming Work

A related work, and in fact one predating and inspiring this paper, is an unpublished manuscript of [18]. Their paper too aims at leveraging the knowledge of collusive players in a combinatorial auction, but in a quite different model. On one hand, they assume that at most one coalition of players exists, that the knowledge that the coalition has about the valuations of the independent players is within an approximation factor k , and that the mechanism designer is aware of the value of k . (Note that in many auctions it is reasonable to assume that competitors can estimate within a factor of 2 each other’s valuations for the goods.) On the other hand, they can leverage such knowledge by relying on a simpler solution concept: namely, their mechanism is *dominant-strategy truthful* for independent players and works in “undominated strategies” for the members of the coalition.

2 Preliminaries

Combinatorial Auctions.

In a combinatorial auction with n players and multiple goods for sale,

- The *true valuation of a player i* consists of a function TV_i mapping every subset S of the goods to a non-negative integer, where $TV_i(S)$ represents the true value that i has for S .
- An *allocation A* consists of a partition of the goods, $A = A_0, A_1, \dots, A_n$, where A_0 represents the unallocated goods, and A_i (for $i > 0$) the subset of goods allocated to player i .

- An *outcome Ω* consists of an allocation A and a *price profile P* , a vector of integers indexed by the players. If positive, P_i represents the amount paid by i , else $-P_i$ represents the amount received by i .

We say that a combinatorial auction is *unrestricted* to stress that the function TV is not assumed to be of any special form: each value $TV_i(S)$ is independent of $TV_j(S')$ for any $(i, S) \neq (j, S')$.

As is standard, goods are non-transferable and a player i ’s individual utility depends solely on how much he pays and on which goods he receives: in an outcome (A, P) , it consists of $TV_i(A_i) - P_i$.

Extensive-Form Public-Action Auction Mechanisms.

We focus solely on auction mechanisms of extensive form. Thus our mechanisms must specify the *decision nodes* (of a game tree), the player(s) *acting* at each node, the set of *actions* available to each acting player at each node, and the auction outcome (i.e., the allocation A and the price profile P) associated to each *terminal node* —leaf of the game tree. Our mechanisms may actually specify multiple players to act simultaneously at some decision nodes. Our mechanisms also are of *public action*: that is, each action becomes common knowledge as soon as it is played.⁵

A player i ’s *strategy* specifies i ’s action at each decision node in which i acts. A *play* of a mechanism M consists of a profile of strategies. If σ is such a play, then

- $H(\sigma)$ denotes the *history* of the play, that is the sequence of decision nodes together with the terminal node of the game tree reached when executing M with each player i choosing his actions according to σ_i .
- $M(\sigma)$ denotes the auction outcome (A, P) associated to $H(\sigma)$.

If M is probabilistic, then both $H(\sigma)$ and $M(\sigma)$ are distributions, respectively over histories and auction outcomes.

For each player i , a mechanism M must provide a particular *opt-out* strategy OUT_i , and must satisfy the following

opt-out condition: for each player i and each strategy subprofile σ_{-i} for players other than i , $u_i(M(\text{OUT}_i \sqcup \sigma_{-1})) = 0$ (with probability 1 if M is probabilistic).

⁵We refrain from using the more standard term “perfect-information” to avoid confusion. Our setting is in fact of “incomplete information.” That is, a player’s true valuation is not exactly known to his opponents. And mechanisms of “perfect information and incomplete information” would be too much...

Generalized Contexts and Auctions.

A traditional context for a combinatorial auction can be fully specified by the true-valuation profile TV alone. Indeed, the outcome set, and the players' utility functions are uniquely determined once TV is specified. Following [7] ([8] for a better treatment), we enrich such a traditional context by including the knowledge that each player has about the valuations of the other players, as well as the collusion structure.

The external knowledge is formalized without any recourse to any Bayesian information. It can be "zero", but when this is not the case, a mechanism should try to leverage it to its designer's advantage.

The collusion structure too can be "empty" in the sense that all players can be independent.

Definition 5. A generalized auction context consists of three components:

1. The true-valuation profile TV .
2. The collusion structure (\mathbb{C}, I) , where \mathbb{C} is a partition of the players, and I the set of all players i such that $\{i\} \in \mathbb{C}$.

We refer to a player in I as independent, to a player not in I as collusive, to any $C \in \mathbb{C}$ of cardinality > 1 as a collusive set. We use the term agent to denote either an independent player or a collusive set. Since each player i , collusive or not, belongs to a single set in \mathbb{C} , for uniformity of presentation we may denote by \mathbb{C}_i the set to which i belongs.

If \mathcal{A} is an agent, then the internal knowledge of \mathcal{A} is $TV_{\mathcal{A}}$, and the utility of \mathcal{A} in an outcome $\Omega = (A, P)$, $u_{\mathcal{A}}(\Omega)$, is $\sum_{i \in \mathcal{A}} TV_i(A_i) - P_i$.

3. The external-knowledge vector EK : for each agent $\mathcal{A} \in \mathbb{C}$, $EK_{\mathcal{A}}$ is a set of valuation subprofiles, for the players outside \mathcal{A} , such that $TV_{-\mathcal{A}} \in EK_{\mathcal{A}}$.

If \mathcal{C} is a generalized auction context whose components have not been explicitly specified, then by default we assume that $\mathcal{C} = (TV^{\mathcal{C}}, (\mathbb{C}^{\mathcal{C}}, I^{\mathcal{C}}), EK^{\mathcal{C}})$. We say that (\mathcal{C}, M) is a generalized auction if M is an auction mechanism, and \mathcal{C} a generalized auction context.

Let us now define the relevant knowledge of an agent. Essentially this is the outcome with maximum welfare known to its members.

Definition 6. (Relevant Knowledge) Given a generalized context \mathcal{C} and an agent \mathcal{A} , we define $RK_{\mathcal{A}}^{\mathcal{C}}$, the relevant knowledge of \mathcal{A} , to be the outcome with maximum revenue among all outcomes (A, P) such that, for all player j

1. If $j \in \mathcal{A}$, then $P_j = TV_j^{\mathcal{C}}(A_j)$.
2. If $j \notin \mathcal{A}$, then $V_j(A_j) \geq P_j$ for all $V \in EK_{\mathcal{A}}^{\mathcal{C}}$.

The maximum known welfare of \mathcal{A} , $\text{MKW}_{\mathcal{A}}$, is the revenue of $RK_{\mathcal{A}}^{\mathcal{C}}$. The maximum known welfare of \mathcal{C} , $\text{MKW}^{\mathcal{C}}$, is $\max_{\mathcal{A} \in \mathbb{C}^{\mathcal{C}}} \text{MKW}_{\mathcal{A}}$.

Remarks

- A collusion structure specifies separately the set I for convenience and clarity only.
- Recall that in a collusion-leveraging mechanism the members of the same collusive set will de facto behave *as if* they are sharing their knowledge. Accordingly, in a collusion-leveraging mechanism, saying that \mathcal{A} knows x means that all players $i \in \mathcal{A}$ know x .
- $RK_{\mathcal{A}}^{\mathcal{C}}$ is the knowledge of agent \mathcal{A} that our mechanism is capable of using to the designer's advantage, while $TV_{\mathcal{A}}^{\mathcal{C}}$ and $EK_{\mathcal{A}}^{\mathcal{C}}$ are the knowledge that \mathcal{A} uses to choose rationally the actions of its members.
- Since we envisage a dynamic collusion formation, the generalized context \mathcal{C} is that arising at the end of our third stage, where coalitions are already formed.
- When the generalized auction context \mathcal{C} under consideration is clear, we may "not use it as a superscript." For instance, we may simply write MKW instead of $\text{MKW}^{\mathcal{C}}$.

3 Distinguishable Domination and Rationally Robust Implementation

We adopt the same solution concept and implementation notion of [8] (see their paper for motivations and basic properties of the notions). Their notations and definitions are reported below essentially verbatim, except for some slight adjustments to our setting.

Through out this paper, whenever we say that S is a vector of strategy (sub)sets in a generalized auction (\mathcal{C}, M) , we mean that each $S_{\mathcal{A}}$ is a (sub)set of agent \mathcal{A} 's strategies. For such an S , we define the Cartesian closure of S as $\overline{S} = \prod_{\mathcal{A} \in \mathbb{C}^{\mathcal{C}}} S_{\mathcal{A}}$, and we define $\overline{S}_{-\mathcal{A}} = \prod_{C \in \mathbb{C}^{\mathcal{C}}, C \neq \mathcal{A}} S_C$.

Definition 7. (Distinguishable Strategies.) In a generalized auction $G = (\mathcal{C}, M)$, let S be a vector of deterministic-strategy subsets, and let $\sigma_{\mathcal{A}}$ and $\sigma'_{\mathcal{A}}$ be two different strategies for some agent \mathcal{A} . Then we say that $\sigma_{\mathcal{A}}$ and $\sigma'_{\mathcal{A}}$ are distinguishable over S if $\exists \tau_{-\mathcal{A}} \in \overline{S}_{-\mathcal{A}}$ such that

$$H(\sigma_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}) \neq H(\sigma'_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}).^6$$

⁶If $H(\sigma_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}})$ and $H(\sigma'_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}})$ are distributions over the histories of G , then the inequality means that the two distributions are different.

If this is the case, we say that $\tau_{-\mathcal{A}}$ distinguishes $\sigma_{\mathcal{A}}$ and $\sigma'_{\mathcal{A}}$ over S ; else, that $\sigma_{\mathcal{A}}$ and $\sigma'_{\mathcal{A}}$ are equivalent over S .

Definition 8. (Distinguishably Dominated Strategies.) Let $G = (\mathcal{C}, M)$ be a generalized auction, \mathcal{A} an agent, $\sigma_{\mathcal{A}}$ and $\sigma'_{\mathcal{A}}$ two strategies of \mathcal{A} , and S a vector of deterministic strategy subsets. We say that $\sigma_{\mathcal{A}}$ is distinguishably dominated (by $\sigma'_{\mathcal{A}}$) over S —equivalently that $\sigma'_{\mathcal{A}}$ distinguishably dominates $\sigma_{\mathcal{A}}$ over S — if

1. $\sigma_{\mathcal{A}}$ and $\sigma'_{\mathcal{A}}$ are distinguishable over S ; and
2. $\mathbb{E}[u_{\mathcal{A}}(M(\sigma_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))] < \mathbb{E}[u_{\mathcal{A}}(M(\sigma'_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))]$ for all strategy sub-vectors $\tau_{-\mathcal{A}}$ distinguishing $\sigma_{\mathcal{A}}$ and $\sigma'_{\mathcal{A}}$ over S .

Definition 9. (Compatible Contexts.) We say that a generalized context \mathcal{C}' is compatible with agent \mathcal{A} in a generalized auction $G = (\mathcal{C}, M)$ if: \mathcal{C}' and \mathcal{C} have the same set of players and the same set of goods, $\mathcal{A} \in \mathbb{C}^{\mathcal{C}'}$, $TV_{\mathcal{A}}^{\mathcal{C}'} = TV_{\mathcal{A}}^{\mathcal{C}}$, and $EK_{\mathcal{A}}^{\mathcal{C}'} = EK_{\mathcal{A}}^{\mathcal{C}}$.

Notice that \mathcal{C}' being compatible with \mathcal{A} implies that $RK_{\mathcal{A}}^{\mathcal{C}'} = RK_{\mathcal{A}}^{\mathcal{C}}$ also, since \mathcal{A} 's relevant knowledge is deduced from its internal and external knowledge.

Definition 10. (L_1 -Rationally Robust Plays) Let $G = (\mathcal{C}, M)$ be a generalized auction, i a player and \mathcal{A} an agent in G . Let $\Sigma^0 = \prod \Sigma_i^0$ be a profile of strategy sets, such that Σ_i^0 is the set of all possible strategies of i according to M .

- We define $\Sigma_{\mathcal{C}, \mathcal{A}}^1$ to be the set of strategies in $\Sigma_{\mathcal{A}}^0$ that are not distinguishably dominated over Σ^0 in G , and $\Sigma_{\mathcal{C}}^1$ to be $\prod_{\mathcal{A} \in \mathbb{C}^{\mathcal{C}}} \Sigma_{\mathcal{C}, \mathcal{A}}^1$.
- We say that a strategy $\sigma_{\mathcal{A}} \in \Sigma_{\mathcal{C}, \mathcal{A}}^1$ is globally distinguishably dominated if there exists a strategy $\sigma'_{\mathcal{A}} \in \Sigma_{\mathcal{C}, \mathcal{A}}^1$, such that for all contexts \mathcal{C}' compatible with \mathcal{A} , $\sigma'_{\mathcal{A}}$ distinguishably dominates $\sigma_{\mathcal{A}}$ over $\Sigma_{\mathcal{C}'}^1$, where $\Sigma_{\mathcal{C}'}^1$ is defined as $\Sigma_{\mathcal{C}}^1$ but for auction (\mathcal{C}', M) .
- We denote by $\Sigma_{\mathcal{C}, \mathcal{A}}^2$ the set of all strategies in $\Sigma_{\mathcal{C}, \mathcal{A}}^1$ that are not globally distinguishably dominated.
- We say that a strategy vector σ is an L_1 -rationally robust play of auction G if $\sigma_{\mathcal{A}} \in \Sigma_{\mathcal{C}, \mathcal{A}}^2$ for all agent \mathcal{A} .

Definition 11. (L_1 -Rationally Robust Implementation.) Let \mathcal{C} be a class of generalized auction contexts, \mathbb{P} be a property over (distributions of) outcomes of contexts in \mathcal{C} , and M an extensive-form auction mechanism with simultaneous and public actions. We say that M L_1 -rationally robustly implements \mathbb{P} if, for all contexts $\mathcal{C} \in \mathcal{C}$

1. for each player i , $\text{OUT}_i \notin \Sigma_{\mathcal{C}, i}^2$, where $\Sigma_{\mathcal{C}, i}^2$ is the subset of strategies for player i , obtained by taking the i -th component of each $\sigma_{\mathbb{C}_i} \in \Sigma_{\mathcal{C}, \mathbb{C}_i}^2$.
2. for all L_1 -rationally robust plays σ of the auction (\mathcal{C}, M) , \mathbb{P} holds for $M(\sigma)$.

Remarks.

- Different from [7, 8] where $\Sigma_{\mathcal{C}, \mathcal{A}}^2$ are explicitly defined for independent players only, here we define $\Sigma_{\mathcal{C}, \mathcal{A}}^2$ for both independent players and collusive sets.
- In the definition above, the compatibility of a generalized context with an agent \mathcal{A} is defined with respect to \mathcal{A} 's internal and external knowledge only. However, \mathcal{A} may have all other types of knowledge, and when this is the case then \mathcal{A} is entitled to use it for pruning the compatible contexts it should consider. In particular, an agent may have knowledge about the collusion structure, as well as knowledge about other players' external knowledge. Our L_1 label in definitions 10 and 11 highlights the fact that we only rely on the simplest, level-one knowledge (about players true valuations).

Of course, as per footnote 3, if a mechanism implements a property \mathbb{P} L_1 -rationally robustly, then \mathbb{P} holds whatever additional information each \mathcal{A} may have.

4 Our Mechanism

In the description of our mechanism,

- $\{1, \dots, n\}$ is assumed to be the set of players;
- ϵ , ϵ_1 , and ϵ_2 are three —arbitrarily small— constants in $(0, 1)$ such that $2n\epsilon_2 < \epsilon_1$.
- an outcome (A, P) is called *reasonable* if each P_j is non-negative;
- an allocation A is said to be for a set C of players if $A_j = \emptyset$ whenever $j \notin C$;
- numbered steps refer to steps taken by the players, “bulleted” ones to steps taken by the mechanism.

Mechanism \mathcal{M}

- Set $A_i = \emptyset$ and $P_i = 0$ for each player i . (Outcome (A, P) will be the *final outcome* of the mechanism.)
- 1. Each player i , simultaneously with the others, publicly announces three things:
 - (1) a subset of players including i , C_i (allegedly the collusive set to which i belongs);
 - (2) an allocation for C_i , S^i (allegedly the allocation desired by C_i); and

(3) a reasonable outcome, $\Omega^i = (\alpha^i, \pi^i)$ (allegedly the relevant knowledge of C_i).

- Set: $R_i = \text{REV}(\Omega^i)$ for each player i , $\star = \text{argmax}_i R_i$ (ties broken lexicographically), and $R' = \max_{i \notin C_\star} R_i$.
(We shall refer to player \star as the “star player”, and to R' as the “second highest —announced— revenue”.)
 - For each player i for which C_i includes a player j such that $i \notin C_j$, do:
 - (1) reset $P_i := P_i + R_\star + \epsilon_1$ (i.e., impose to i a fine of $R_\star + \epsilon_1$ payable to the mechanism/seller)
 - (2) for each $j \in C_i$ such that $i \notin C_j$, reset $P_i := P_i + R_\star + \epsilon_1$ and $P_j := P_j - R_\star - \epsilon_1$ (i.e., have i pay $R_\star + \epsilon_1$ to j)
 - If there is a player i such that $P_i > 0$ (i.e., if i has been fined), ABORT the auction (i.e., no further money exchanges hands, and all goods remain unallocated for ever).
 - Publicly flip a biased coin c_1 which comes up Heads with probability ϵ . If Heads: uniformly and randomly choose a player i , reset $\star := i$ and $R' := 0$. (In this case, R' does not quite correspond to the second highest announced revenue, but this “mismatch” only happens rarely.)
 - Publicly flip a fair coin c_2 . If Heads: reset $A := S^\star$ and HALT.
2. (If Tails:) Each player i such that $i \notin C_\star$ and $\pi_i^\star \geq 1$ publicly, and simultaneously with the others, announces YES or NO (i.e., declares whether he wants to receive the subset of goods α_i^\star for a price $\pi_i^\star - \epsilon_2$)
- Reset allocation and prices as follows:
 - (1) $P_\star := R' - n\epsilon_2$;
 - (2) for each player i such that either $i \in C_\star$ or $\pi_i^\star = 0$, reset $A_i := \alpha_i^\star$; and
 - (3) for each player i such that $i \notin C_\star$ and $\pi_i^\star \geq 1$: if i announced NO, then $P_\star := P_\star + \pi_i^\star$ (i.e., \star is punished due to i announcing NO); else, $A_i := \alpha_i^\star$, $P_i := P_i + \pi_i^\star - \epsilon_2$, and $P_\star := P_\star - (\pi_i^\star - \epsilon_2)$ (i.e., \star is rewarded due to i announcing YES).
 - Finally, reset $P_i := P_i - \epsilon_2(1 - \frac{1}{1+R_i})$ for each player i (i.e., to break “utility ties”, a small reward is added to each player, increasing with his announced revenue).

Remarks

- *Consistency Check.* Notice that our mechanism checks consistency among collusive players in the second mechanism step after Step 1. But this consistency check is quite elementary. In particular, if (a) i declares that he belongs to the same collusive set as j and k , while (b) j declares to collude only with i and (c) k declares to col-

lude only with i , then our mechanism continues unperturbed, despite the obvious discrepancies of these declarations. Nonetheless, our elementary consistency check suffices to guarantee that our benchmark is achieved in any rational play of our mechanism, that is, for any profile of Σ^2 strategies.

- *Small Constants.* The mechanism makes use of 3 arbitrarily small constants only for “properly breaking utility ties.”

5 Our Analysis

5.1 Notation

To state our main theorem and lemmas we utilize the following notation

- *Social Welfare of an Allocation.* If A is an allocation, then $\text{sw}(A)$ denotes the social welfare of A : that is

$$\text{sw}(A) = \sum_k TV_k(A_k).$$

- *Revenue of an Outcome.* If $\Omega = (A, P)$ is an outcome, then $\text{REV}(\Omega)$ denotes the revenue of Ω : that is

$$\text{REV}(\Omega) = \sum_k P_k.$$

- *Hidden Value of an Outcome.* If C is an agent and $\Omega = (A, P)$ is an outcome, then the hidden value of Ω for C , $\text{HiddenV}_C(\Omega)$, is

$$\text{HiddenV}_C(\Omega) = \sum_{k \in C} TV_k(A_k) + \sum_{k \notin C} P_k.$$

(Notice that in an execution of \mathcal{M} , if \star 's announced outcome is Ω , then when coin c_2 comes up Tails, the maximum utility that C_\star can possibly get by selling the goods according to Ω is $\text{HiddenV}_{C_\star}(\Omega)$, disregarding small constants. In fact, this utility can be substantially decreased if some players reject “their offers.”)

5.2 Statement of Our Lemmas

Our main theorem is based on five lemmas, stated below but proven in the full version of our paper (available at [http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design/](http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design/)).

Lemma 2 actually is an immediate corollary of the first two lemmas of [7], but the others are new.

The statements of our lemmas refer to a play σ of a game $(\mathcal{C}, \mathcal{M})$, where $\mathcal{C} = (TV, (C, I), EK)$ is a generalized context, and \mathcal{M} our mechanism of Section 4. The relevant knowledge of \mathcal{C} is denoted by RK . For short, we redefine $\Sigma^1 = \Sigma_{\mathcal{C}}^1$ and $\Sigma^2 = \Sigma_{\mathcal{C}}^2$.

Lemma 1. For all agents C and all $\sigma_C \in \Sigma_C^1$, the following two properties hold in Step 1:

- P1. for all $i \in C$, $C_i \subseteq C$ (that is, i never includes a player outside C in his announced collusive set); and
- P2. for any two different players $i_1, i_2 \in C$, $i_2 \in C_{i_1}$ if and only if $i_1 \in C_{i_2}$ (that is, C 's members declare their collusive sets consistently with each other).

Lemma 2. For all agents C and all $\sigma_C \in \Sigma_C^1$, if $\star \notin C$, then in Step 2, for all players i in $C \setminus C_\star$ such that $\pi_i^\star \geq 1$:

- i announces YES whenever $TV_i(\alpha_i^\star) \geq \pi_i^\star$, and
- i announces NO whenever $TV_i(\alpha_i^\star) < \pi_i^\star$.

Lemma 3. For all agents C and all $\sigma_C \in \Sigma_C^1$, if $\star \in C$, then in Step 2, for all players i in $C \setminus C_\star$ such that $\pi_i^\star \geq 1$, i always announces YES.

Lemma 4. For all agents C , all $\sigma_C \in \Sigma_C^2$, and player $j \in C$ such that j is the lexicographical first player among all players $i \in C$ with $\text{REV}(\Omega^i) = \max_{k \in C} \text{REV}(\Omega^k)$, we have that $\text{HiddenV}_C(\Omega^j) \geq \text{REV}(RK_C)$ (that is, C 's members do not “underbid” on the hidden value of their announced outcomes).

Lemma 5. For all agents C and all $\sigma_C \in \Sigma_C^2$, we have that $\max_{k \in C} \text{REV}(\Omega^k) \geq \frac{\text{REV}(RK_C)}{3}$ (that is, C 's members do not “underbid too much” on the revenue of their announced outcomes).

5.3 Statement and Proof of Our Theorem

Theorem 1. For all generalized contexts \mathcal{C} and all L_1 -rationally robust plays σ of $(\mathcal{C}, \mathcal{M})$, we have that

$$\mathbb{E}[\text{REV}(\mathcal{M}(\sigma))] + \mathbb{E}[\text{sw}(\mathcal{M}(\sigma))] \geq \frac{(1 - \epsilon)\text{MIKW}}{6} - \epsilon_1.$$

Proof. First of all, it should be obvious from our lemmas that, for each player i , $\text{OUT}_i \notin \Sigma_i^2$. Now let's proceed with the rest of the proof.

Let C be the agent such that $\text{REV}(RK_C) = \text{MIKW}$. Notice that by Lemma 1, in execution σ , the mechanism does not abort before Step 2.

When $c_1 = \text{Heads}$, no matter whom the star player is, we have that: (1) the expected social welfare is $\mathbb{E}[\text{sw}(\mathcal{M}(\sigma)) | c_1 = \text{Heads}] \geq 0$, because $TV_i(S) \geq 0$ for any player i and any subset S of the goods; and (2) the expected revenue is $\mathbb{E}[\text{REV}(\mathcal{M}(\sigma)) | c_1 = \text{Heads}] > -\frac{\epsilon_1}{2}$, because when and only when $c_2 = \text{Tails}$, the star player pays at least $R' - n\epsilon_2 = -n\epsilon_2$ to the mechanism and the mechanism gives back total reward less than $n\epsilon_2$ to the players. Therefore $\mathbb{E}[\text{sw}(\mathcal{M}(\sigma)) | c_1 = \text{Heads}] + \mathbb{E}[\text{REV}(\mathcal{M}(\sigma)) | c_1 = \text{Heads}] > -\frac{2n\epsilon_2}{2} > -\frac{\epsilon_1}{2}$.

When $c_1 = \text{Tails}$, the mechanism does not reset the value of \star and R' . By the way that ties are broken, the star player is the lexicographically first player in his collusive set among the players who have announced the maximum revenue in Step 1. We distinguish two cases.

Case 1: $\star \in C$.

In this case we have the following observations:

- (1) By Lemma 4, $\text{HiddenV}_C(\Omega^\star) \geq \text{REV}(RK_C)$.
- (2) When $c_2 = \text{Tails}$, the revenue that \star pays to the mechanism is at least $R' - n\epsilon_2 \geq -n\epsilon_2$.
- (3) By Lemma 3, when $c_2 = \text{Tails}$, every $k \in C \setminus C_\star$ with $\pi_k^\star \geq 1$ announces YES, and thus the social welfare generated from players in C is $\sum_{k \in C} TV_k(\alpha_k^\star)$, and the revenue generated from them is 0, because for each player $k \in C \setminus C_\star$ with $\pi_k^\star \geq 1$, the mechanism charges k with price $\pi_k^\star - \epsilon_2$, but rewards the star player the same amount.
- (4) By Lemma 1, $k \notin C_\star$ for all players $k \notin C$, and thus when $c_2 = \text{Tails}$, any such k with $\pi_k^\star \geq 1$ gets to announce YES or NO in Step 2. By Lemma 2, for each such k , if k announces YES, then we have that $TV_k(\alpha_k^\star) \geq \pi_k^\star$, therefore the social welfare generated due to this announcement is at least π_k^\star , and the revenue generated is 0 (again the money paid by k goes to \star); if k announces NO, then the social welfare generated due to this announcement is 0, but the revenue generated is π_k^\star , because the star player is punished by π_k^\star . Therefore the sum of the social welfare and revenue generated due to the announcements made by the players outside C is at least $\sum_{k \notin C} \pi_k^\star$.
- (5) When $c_2 = \text{Tails}$, the reward given to each player i in the last step is $\epsilon_2(1 - \frac{1}{1 + \text{REV}(\Omega^i)}) < \epsilon_2$.

Accordingly, when $c_2 = \text{Tails}$, we have that

$$\begin{aligned} & \text{REV}(\mathcal{M}(\sigma)) + \text{sw}(\mathcal{M}(\sigma)) \\ & > \sum_{k \in C} TV_k(\alpha_k^\star) + \sum_{k \notin C} \pi_k^\star - 2n\epsilon_2 \\ & = \text{HiddenV}_C(\Omega^\star) - 2n\epsilon_2 > \text{MIKW} - \epsilon_1. \end{aligned}$$

Because $\text{REV}(\mathcal{M}(\sigma)) = 0$ and $\text{sw}(\mathcal{M}(\sigma)) \geq 0$

when $c_2 = \text{Heads}$, we have that

$$\begin{aligned} & \mathbb{E}[\text{REV}(\mathcal{M}(\sigma))|c_1 = \text{Tails}] \\ & + \mathbb{E}[\text{SW}(\mathcal{M}(\sigma))|c_1 = \text{Tails}] \\ \geq & \frac{\text{MKW} - \epsilon_1}{2}, \end{aligned}$$

and thus

$$\begin{aligned} & \mathbb{E}[\text{REV}(\mathcal{M}(\sigma))] + \mathbb{E}[\text{SW}(\mathcal{M}(\sigma))] \\ \geq & -\frac{\epsilon\epsilon_1}{2} + (1 - \epsilon) \cdot \frac{\text{MKW} - \epsilon_1}{2} \\ \geq & \frac{(1 - \epsilon)\text{MKW}}{6} - \epsilon_1. \end{aligned}$$

Case 2: $\star \notin C$.

In this case, by Lemma 1, $C_\star \cap C = \emptyset$; and by Lemma 5, $\max_{k \in C} \text{REV}(\Omega^k) \geq \frac{\text{REV}(RK_C)}{3}$. Therefore $R' \geq \frac{\text{REV}(RK_C)}{3}$, by definition of R' . When $c_2 = \text{Tails}$, the star player pays at least $R' - n\epsilon_2$ to the mechanism, and the reward given back by the mechanism to the players is at most $n\epsilon_2$. Thus $\text{REV}(\mathcal{M}(\sigma)) \geq R' - 2n\epsilon_2 > \frac{\text{REV}(RK_C)}{3} - \epsilon_1 = \frac{\text{MKW}}{3} - \epsilon_1$. Because $\text{sw}(\mathcal{M}(\sigma)) \geq 0$ always, we have that

$$\begin{aligned} & \mathbb{E}[\text{REV}(\mathcal{M}(\sigma))|c_1 = \text{Tails}] \\ & + \mathbb{E}[\text{SW}(\mathcal{M}(\sigma))|c_1 = \text{Tails}] \\ \geq & \frac{\text{MKW}}{6} - \frac{\epsilon_1}{2}. \end{aligned}$$

Therefore

$$\begin{aligned} & \mathbb{E}[\text{REV}(\mathcal{M}(\sigma))] + \mathbb{E}[\text{SW}(\mathcal{M}(\sigma))] \\ \geq & -\frac{\epsilon\epsilon_1}{2} + (1 - \epsilon) \cdot \left(\frac{\text{MKW}}{6} - \frac{\epsilon_1}{2} \right) \\ \geq & \frac{(1 - \epsilon)\text{MKW}}{6} - \epsilon_1. \end{aligned}$$

Q.E.D.

References

- [1] L.M. Ausubel and P. Milgrom. The Lovely but Lonely Vickrey Auction. *Combinatorial Auctions*, MIT Press, pages 17-40, 2006.
- [2] S. Barbera and M.O. Jackson. Strategy-Proof Exchange. *Econometrica*, Vol.63, No.1, pages 51-87, Jan., 1995.
- [3] M. Babaioff, R. Lavi and E. Pavlov. Single-Value Combinatorial Auctions and Implementation in Undominated Strategies. (Full version of) *Symposium on Discrete Algorithms*, pages 1054-1063, 2006.
- [4] J. Chen, A. Hassidim, and S. Micali. Guaranteeing Perfect Revenue From Perfectly Informed Players. To appear at *The First Symposium on Innovations in Computer Science (ICS2010)*.
- [5] Y. Che and J. Kim. Robustly Collusion-Proof Implementation. *Econometrica*, Vol.74, No.4, pages 1063-1107, Jul., 2006.
- [6] Y. Che and J. Kim. Optimal collusion-proof auctions. *Journal of Economic Theory*, Vol.144, Issue 2, pages 565-603, Mar. 2009.
- [7] J. Chen and S. Micali. A New Approach to Auctions and Resilient Mechanism Design. Full version available at [http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design/](http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design/). (A preliminary version appeared at *Symposium on Theory of computing*, pages 503-512, 2009.)
- [8] J. Chen and S. Micali. Rational Robustness for Mechanism Design. Submitted to STOC'10, full version available at [http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design/](http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design/).
- [9] E.H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, Vol.11, No.1, pages 17-33, Sep., 1971.
- [10] A. Goldberg and J. Hartline. Collusion-Resistant Mechanisms for Single-Parameter Agents. *Symposium on Discrete Algorithms*, pages 620-629, 2005.
- [11] T. Groves. Incentives in Teams. *Econometrica*, Vol. 41, No. 4, pages 617-631, Jul., 1973.
- [12] S. Izmalkov, M. Lepinski, and S. Micali. Perfect Implementation of Normal-Form Mechanisms. *MIT-CSAIL-TR-2008-028*, May 2008. (A preliminary version appeared in FOCS'05, pages 585-595, with title "Rational Secure Computation and Ideal Mechanism Design".)
- [13] M. Jackson. Implementation in Undominated Strategies: a Look at Bounded Mechanisms. *Review of Economic Studies*, Vol. 59, No. 201, pages 757-775, Oct., 1992.
- [14] K. Jain and V. Vazirani. Applications of Approximation Algorithms to Cooperative Games. *Symposium on Theory of Computing*, pages 364-372, 2001.
- [15] J.J. Laffont and D. Martimort. Collusion under Asymmetric Information. *Econometrica*, Vol.65, No.4, pages 875-911, Jul. 1997.
- [16] J.J. Laffont and D. Martimort. Mechanism Design with Collusion and Correlation. *Econometrica*, Vol. 68, Issue 2, pages 309-342, Mar. 2000.
- [17] S. Micali and P. Valiant. Resilient Mechanisms for Unrestricted Combinatorial Auctions, 2007. Revised version available at [http://people.csail.mit.edu/silvio/Selected Scientific Papers/Mechanism Design](http://people.csail.mit.edu/silvio/Selected_Scientific_Papers/Mechanism_Design)
- [18] S. Micali and P. Valiant. Leveraging Collusion in Combinatorial Auctions, 2008. Available at <http://web.mit.edu/pvaliant/www/MV2.pdf>
- [19] H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice Welfare*, Vol.16, No.2, pages 279-320, Feb. 1999.
- [20] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, Vol.18, No.3, pages 511-533, Nov. 2001.

- [21] M. Pal and E. Tardo, Group Strategyproof Mechanisms via Primal-Dual Algorithms. *44th Annual IEEE Symposium on Foundations of Computer Science (FOCS'03)*, pages 584-593, 2003.
- [22] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, Vol. 16, No. 1, pages 8-37, Mar., 1961.