Computer Science and Artificial Intelligence Laboratory
Technical Report

MIT-CSAIL-TR-2008-059 October 8, 2008

New Resiliency in Truly Combinatorial Auctions (and Implementation in Surviving Strategies)

Jing Chen and Silvio Micali

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

New Resiliency in Truly Combinatorial Auctions (and Implementation in Surviving Strategies)

Jing Chen and Silvio Micali

October 8, 2008

Abstract

Following Micali and Valiant [MV07.a], a mechanism is resilient if it achieves its objective without any problem of (1) equilibrium selection and (2) player collusion. To advance resilient mechanism design,

- We put forward *a new meaningful benchmark* for the *combined* social welfare-revenue performance of any mechanism in truly combinatorial auctions.
- We put forward *a new notion of implementation,* much more general than the ones used so far, which we believe to be of independent interest.
- We put forward *a new resilient mechanism* that, by leveraging the knowledge that the players have about each other, guarantees at least one half of our benchmark under a very general collusion model.

1 Introduction

1.1 Background on Combinatorial Auctions and Mechanism Design

Truly Combinatorial Auction Contexts. The context of a truly combinatorial auction is so described. There are n players (denoted 1 through n) and a set G of m (indivisible) goods for sale. Each player i has a *valuation* for the goods —a mapping from subsets of G to non-negative reals— denoted by TV_i. The *profile* (i.e., a vector indexed by the players) TV is called the true valuation profile of the auction. An *outcome* specifies how the goods are sold, that is, it consists of: (1) an *allocation* A, that is, a partition of G into $n+1$ subsets, $A = (A_0, A_1, \ldots, A_n)$, and (2) a *price profile* P, that is, a profile of real numbers. We refer to A_0 as the set of unallocated goods, and for each player i, we refer to A_i as the set of goods allocated to i and to P_i as the price of i. Relative to an outcome (A, P), the *utility* of an individual player i is taken to be $TV_i(A_i) - P_i$.

Combinatorial contexts may be hard to work with, and one may want to restrict the players' valuations or assume that multiple copies of each good are available.[1] We say that an auction is *truly combinatorial* if the valuations can be arbitrary, that is for any players i and j and any subsets S and T of G such that $(i, S) \neq (j, T)$, the value of $TV_i(S)$ gives no information about the value of $TV_j(T)$.

Traditional Mechanism Design. A mechanism specifies the *strategies* available to the players, and an *outcome function* mapping a profile of strategies to a *final outcome*. In a normal-form auction mechanism M, a player's (pure) strategy consists of a valuation (also referred to as a *bid*): essentially, all players simultaneously submit their bids, and then the outcome function maps the bid profile to an outcome. Together with an auction context, an auction mechanism defines a game, namely the auction itself. As for any game, the players try to maximize their utilities, and an equilibrium consists of a profile of strategies $(\sigma_1, \ldots, \sigma_n)$ such that, for any player i, as long as i believes that the other players stick to their equilibrium strategies —that is, that they play the strategy sub-profile σ_{-i}— then i is better off (or no worse) sticking to σ_i than switching to any other strategy σ'_i available to him.

Auction mechanisms are designed so that, at equilibrium, the final outcome enjoys a desired

[1] Typical restricted models for combinatorial auctions include: *sub-modularity*, namely $TV_i(S \cup T) \leq TV_i(S) + TV_i(T)$ for any subsets S and T of the goods; *additive-valuation*, namely $TV_i(S) = TV_i(g_1) + \ldots + TV_i(g_k)$ whenever $S = \{g_1, \ldots, g_k\}$; *free-disposal*, namely $TV_i(S) \leq TV_i(T)$ whenever $S \subset T$; *single-mindedness*, namely for each i there is a subset of goods S and a value v such that $TV_i(T) = v$ if $T \supseteq S$, and 0 otherwise; and *unlimited-supply*, informally, there is an unbounded number of copies of each good available, and each player values only sets of distinct goods.

property that depends on the players' true valuations. A typical desideratum is maximizing *social welfare,* that is returning an outcome (A, P) maximizing the function $\textsc{sw}((A, P), TV)$ defined as $\sum_i TV_i(A_i)$. Another typical desideratum is maximizing *revenue*, that is maximizing the function $\textsc{rev}((A, P))$ defined as $\sum_i P_i$.

Resilient Mechanism Design. Mechanism design is not robust, as it suffers from two problems: equilibrium selection and collusion. Let us explain. Traditionally, mechanism design guarantees a given property "at equilibrium." But what if the resulting game has multiple equilibria? Even if the mechanism were such that the desired property held at each of the possible equilibria, the property may not be guaranteed at all. Indeed, even in the presence of just two reasonable equilibria, some of the players may choose their strategies believing that a first equilibrium will be played out, while others choose theirs believing that a second one will be played out. The resulting "mix-and-match" strategy profile may not be an equilibrium at all, and thus the property may not hold. Even when one could guarantee the property in question by a dominant strategy equilibrium, things may go wrong. In fact, although in this case one could reasonably predict the equilibrium that will be played out by rational and independent players, any equilibrium is defined solely in terms of single-player deviations. Now, although no single player may have incentives to deviate from his equilibrium strategy, two or more players may have plenty of reasons for colluding and jointly deviating from their equilibrium strategies. Indeed, as shown by [AM06], the famous VCG mechanism, although dominant-strategy truthful (DST for short) in truly combinatorial auctions, is extremely vulnerable to collusion.[2]

Following [MV07.a], we call *resilient* those mechanisms that guarantee their properties without any equilibrium-selection and no matter how collusive players may behave. The performance of a resilient auction mechanism is typically defined relative to a *benchmark* B, that is a function mapping (sub)profiles of true valuations to real numbers: a mechanism's performance simply consists of the "fraction of B it can guarantee."[3] Indeed, relative to B, a mechanism M is preferable to another mechanism M' if, at equilibrium, it guarantees (either in revenue or in social welfare) a higher fraction

[2] Consider the following auction of 2 goods $\{a, b\}$ and three players $\{1, 2, 3\}$. Player 1 values \$2 billion for a and b together, player 2 values \$1 for a, player 3 values \$1 for b, and nothing else. If they bid truthfully, then player 1 gets $\{a, b\}$ and pays \$2, achieving the social welfare \$2 billion and the revenue \$2. But if player 2 and 3 collude and each bids \$2 billion, then 2 gets a, 3 gets b, and both pay nothing, resulting the social welfare \$2 and the revenue 0.

[3] In a Bayesian setting, where it is assumed that the players' true-valuation profile is drawn from a known distribution, it is natural to define the performance of an auction mechanism relative to this distribution —e.g., as its expected revenue, or social welfare. But when, like in our case no such distribution exists or is known, the only object of interest is the true-valuation profile TV. And, although unknown to the designer, it is "indirectly" relative to it that performance must be defined.

of $B(TV)$ than M'. In auctions under a sufficiently powerful collusion model, it is hard to count on collusive players to generate revenue, social welfare or any other desideratum. Accordingly, the mechanism's performance is measured as fraction of $B(TV_I)$, where I represents the set of independent players. That is, one takes the point of view that one must consider himself lucky if all collusive players spontaneously leave the room, leaving one to conduct the auction with just the independent players alone. We thus benchmark our performance relative only to the true-valuation subprofile of the independent players alone.

The Harmonic Revenue Bound. For truly combinatorial auctions, [MV07.a] provides a very tight upperbound on the revenue obtainable in dominant strategies. Namely, letting H_j denote the j-th Harmonic number (that is, $H_j = \sum_{i=1}^{j} \frac{1}{i}$), and letting $\mathbb{MSW}_{-\star}$ denote the benchmark consisting of the maximum social welfare after removing the "star" player (that is the one valuing some subset of the goods more than anyone values any subset)[4],they prove the following:

([MV07.a], Theorem 3) *For any $n, m > 1$, and any DST mechanism M, there exists a valuation profile BID for a truly combinatorial auction with n players and m goods such that the (expected) revenue generated by M on input BID is at most*

$$\frac{\mathbb{MSW}_{-\star}(BID)}{H_{\min\{n,m\}} - 1},$$

(We note that the above revenue upperbound does not affect restricted combinatorial auctions, for which many revenue mechanisms have been considered [FPS00, JV01, MS01, GH05, BBM07, GHK+05, LS05, FGH+02, GHK+06, BBH+05]. We note too that [MV07.a] also prove that their upperbound is essentially tight, since they can achieve it, in truly combinatorial auctions, by means of an extremely resilient mechanism —indeed one that not only withstands collusive players, but also *irrational* ones.)

The Harmonic revenue bound is very relevant to understand the power available to resilient mechanisms. Indeed, although the best way to ban equilibrium-selection problems consists of designing a DST mechanism, if one wants to guarantee more revenue than a logarithmic fraction of $\mathbb{MSW}_{-\star}$, then the Harmonic revenue bound implies that one has only two alternatives available: either

[4]For any valuation (sub)profile V, letting $\text{MSW}(V) = \max_A \text{SW}(A, V)$, and letting the "star" player, $\star$, be a player for whom there exists a subset of goods $S \subseteq G$ such that $V_\star(S) \geq V_j(T)$ for any player j and any $T \subseteq G$, then $\mathbb{MSW}_{-\star}(V) = \text{MSW}(V_{-\star})$.

A1. Assume that more knowledge is available (e.g., that the seller has some Bayesian information about the players' valuations), or

A2. Adopt a solution concept weaker than dominant strategies.

As we shall see, in this paper we take both alternative, but ($1'$) *without* violating the principle of mechanism design in its purest form —namely that all knowledge resides with the players themselves— and ($2'$) *without* suffering from any equilibrium-selection problems.

1.2 Our Contributions

A New Goal: Social Welfare Plus Revenue (and Resiliently Too!)

An allocation $\mathscr{A}$ is economically efficient if its social welfare is at least as high as that of any other allocation: that is, $\mathscr{A} = \arg\max_A \text{SW}(A, TV)$. Economic efficiency is a very old goal of mechanism design, and indeed achieved in dominant strategies and in truly combinatorial auctions by the VCG mechanism (but not in a collusive setting!). The traditional story for motivating economic efficiency is that of a benevolent government, whose sole desire is that of allocating a set of national resources to its citizens (or domestic firms) so that they end up in the hands of those who value them the most. Indeed, no matter who gets the goods, society as a whole is ultimately better off when this is the case. Since (at least benevolent) governments are supposed to be non-profit organizations, revenue is not a desideratum here. Of course, economically efficient auctions also impose prices to the players, but these prices are almost an "after thought," or a "necessary evil:" they are just means to guarantee that efficiency will be reached. But what is wrong with revenue even in this setting? After all, a benevolent government could transform it into roads, hospitals and other infrastructure from which everyone benefits. From this point of view, a benevolent government might be interested in maximizing social welfare *together with* revenue. Thus:

- *We wish to design mechanisms that aim at maximizing the sum of social welfare and revenue.*

Note that the VCG mechanism is no longer optimal for this "higher" goal: the VCG guarantees it only *"within a fraction of 2."* (Indeed, while it is 100% efficient, it has no revenue guarantees at all, and revenue can in principle match social welfare.) Moreover, such a factor of 2 is guaranteed *only assuming that all players are rational and independent.* In fact, a trivial modification of the example of [AM06] shows that VCG's social welfare and revenue can *both* be essentially 0, despite the presence of a large maximum social welfare and fierce competition for the goods. And this fact can be interpreted as VCG achieving an infinitely small factor of the new goal in the collusive setting. Accordingly, since the government may be benevolent but its citizens may not, we refine our goal above as

- *Designing* **resilient** *mechanisms aiming at maximizing the sum of social welfare and revenue.*

Notice that, although in a rational setting "revenue should lowerbound social welfare," a mechanism aiming at maximizing revenue may not maximize revenue plus social welfare. This is so because, in order to guarantee revenue in the presence of sufficiently powerful collusive players, the mechanism may have to give up some efficiency. Consequently, the social welfare of a revenue-oriented mechanism may exactly equal the revenue generated, so that the total social-welfare-plus-revenue is just twice a modest revenue. However, by directly aiming at maximizing their sum, a resilient mechanism may actually perform much better.

Our Knowledge-Leveraging Approach

The maximization of social welfare plus revenue is certainly easy if the mechanism designer is showered with plenty of information about the players.

If the designer had precise knowledge of TV, the true valuation profile of the players, then he could return the mechanism that (without even asking for or looking at bids) (1) sets the allocation $\mathscr{A}$ to be $\arg\max_A \text{SW}(A, TV)$, and (2) offer each player i such that $\mathscr{A}_i \neq \emptyset$ a take-it-or-leave-it price $\mathscr{P}_i$ where $\mathscr{P}_i$ is infinitesimally less than $TV_i(\mathscr{A}_i)$. Receiving such an offer, each player is better off accepting it, as it gives him a positive utility anyway. Thus, a perfectly informed auctioneer can essentially guarantee that the sum of social welfare and revenue equals twice the maximum social welfare, which is clearly optimal in any mechansim in which no player can be asked to "pay more than he bids." However, perfectly informed designers are a rarity if they exist at all.

A less informed auctioneer may have some Bayesian information about the players, that is he may precisely know the distribution from which the players' true valuations are drawn. Assuming that this is the case, the auctioneer may still be capable of producing a mechanism that generates significant social welfare *and* revenue. However, to obtain the right Bayesian information the designer may have to work very hard, trying to extract it from the players *before the auction begins.* This procedure may prove very expensive, and sometimes just impossible. In particular, for auctions of a single good, [CM88] has fully captured the information structure needed for the designer to generate the maximum possible revenue. But, as concluded by the authors themselves, the assumption that the designer can acquire the knowledge required to produce the optimal mechanism is unusually strong, making it difficult to implement their result.

By contrast, we assume that

- *The designer has no knowledge whatsoever about the players, but*
- *The players have some knowledge about each other.*

That is in our setting at least some player not only knows his own valuation, but also has some information about the other players' valuations. Note that this setting is quite compatible with mechanism design in its purest form, in that we do assume that all knowledge resides with the players themselves. Our assumption about the players is quite realistic: people routinely have some information about their competitors. As for the designer, although everyone is bound to have *some* knowledge, more often than not this knowledge may not be sufficient for the goal at hand. Whatever the case may be, however, it is certain that

> *If we can construct mechanisms that (1) are resilient against a broad collusion model and (2) guarantee reasonable social welfare plus revenue without assuming any special knowledge of the designer, then such mechanisms can be employed with confidence in most realistic settings.*

Construct such mechanisms precisely is what we plan to do.

Our Collusion Model

To be on the safe side, we envisage a very adversarial collusion model. In particular, we let the players collude in total secrecy, and with perfect coordination (for instance via secret binding contracts, if they so want). Further, we put no restrictions on the number of collusive players, nor on the number of collusive sets in which they may partition themselves. We also let each collusive set try to maximize its own collective utility function (rather than —say— the sum of the individual utilities of its members). We do, however, impose a minimal restriction on such collective utility functions in order to "prevent collusive players from becoming irrational" —which would be a totally different ball game. (Indeed, the difference between a group of crazy players bidding in a crazy manner and a group of rational players bidding so as to maximize a crazy collective utility function is quite vague.)

Our Solution Concept

As already announced, our collusion-resilient mechanism is not dominant-strategy truthful. Yet, it does not suffer from equilibrium-selection problems. The reason for this is very simple: we rely on an *equilibrium-less* solution concept. In essence, our mechanism guarantees that, as long as each player selects a strategy surviving iterated elimination of weakly dominated strategies, our goal is achieved. Thus, we do not rely on the players' beliefs, only on their rationality. Nor do we rely on the players' joint selection of a profile of strategy forming an equilibrium: any profile of "not-dumb" strategies will do. Thus, while we loose predictability about the profile of strategies ultimately played out, we do not loose meaningfulness: after all such predictability is a means to an end, not the end itself!

We note that in our setting the process of iterated elimination of weakly dominated strategies is more lax (that is, lets more strategies survive) than in the classical setting, where all players are rational and independent. Indeed, since we are also dealing with players who secretly collude and optimize secret collusive utility functions, it is hard for a player to "dismiss more than just a handful of strategies at each iteration."

Our Result

As usual in resilient mechanism design, we are "off the hook" if all players are collusive, but are responsible for significant social welfare and revenue as long as even a single rational and independent player exists. We provide a probabilistic mechanism that, in surviving strategies, guarantees that the sum of social welfare and revenue is within a factor of 2 of the following benchmark: the revenue that the "most informed independent player could guarantee if he were in charge of selling the goods." In essence, although the designer is clueless while the players are quite informed about each other, our mechanism lets a totally ignorant designer to perform as well as the most informed player.

The difficulty for our mechanism comes from the fact that we do not control the knowledge that the players may have. That is, we do not assume that the players only have the knowledge useful to our mechanism. But we let them have *arbitrary* knowledge, and thus —whether independent or collusive— they will use their knowledge for their selfish purposes, which may be quite antagonistic to our goals and the desired mode of operation of our mechanism.

Let us now proceed a bit more formally.

2 Our Knowledge Benchmark

As mentioned, we envisage a setting where the players, but not the designer, have "useful knowledge." In a combinatorial auction in which a government wishes to generate revenue from —say— selling 8 licenses to 8 wireless companies, useful knowledge naturally is the revenue that each of the wireless companies could generate if it were to sell the licenses to its peers in a personalized sale: that is, by offering a separate subset of the licenses to each of the other companies in a take-it-or-leave-it price. We then take the position that

> *Any of the 8 wireless players could do a better job in selling the 8 licenses to its other 7 peers than the government could in selling the licenses himself to all 8 players.*

Accordingly, a reasonable benchmark for the government might be the minimum revenue that one of the 8 wireless players could guarantee. A better benchmark, could be the revenue that a randomly selected company among the 8 ones can guarantee. A better yet benchmark is *the maximum revenue that one of 8 wireless companies could guarantee.*

Our Main result is that the sum of social welfare and revenue can achieve this benchmark within a factor of 2, in a very adversarial collusion model.[5] In so doing, however, we do not want to restrict the players to have just the knowledge relevant to our benchmark, thus facilitating our mechanisms.

We view each player i of a combinatorial auction as having not only precise *internal* knowledge of his own TV_i, but also some *external* knowledge about TV_{-i}. This view does not imply any loss of generality, since a player's external knowledge could be "empty."

The external knowledge we need for our benchmark essentially consists of the "best way known to a player of selling the goods to the other players." But if we base our benchmark on this natural knowledge, we should also refrain from assuming that this is the only type of knowledge the players have. Such an assumption might very well be very convenient for designing mechanisms achieving our benchmark, but risks of being quite unrealistic and diminish the meaningfulness of our results. Indeed, the players may have all kind of knowledge in addition to ways of selling the goods, and whatever mechanism we choose, in the resulting game they will rationally choose their actions based on all the knowledge available to them. Therefore, to enhance the meaningfulness of our results, we should aim at achieving our benchmark, no matter what other additional knowledge the players may have.

Let us now see, first intuitively and then more formally, how our benchmark can be derived from any kind of external knowledge.

[5] As for revenue alone, our benchmark is the second highest revenue known to the 8 wireless players. In a second result, we show that this benchmark too can be achieved within a factor of 2, in the same adversarial model.

2.1 An Intuitive Discussion

General vs. Relevant External Knowledge. We distinguish two types of external knowledge for each player i: (1) his *general external knowledge*, denoted by GK_i, which the player uses to rationally choose his actions, and (2) his *relevant external knowledge*, denoted by RK_i, derived from GK_i and used to define our benchmark. We insist that GK_i be *genuine,* and that RK_i be (deduced from and thus) *consistent* with GK_i. Accordingly, RK_i is somewhat redundant, but we find it convenient to our analysis specifying it explicitly. Given our auction setting, the relevant knowledge RK_i consists of the "best *guaranteed* way known to i of selling the goods to the other players." Let us stress again, however, that in our model no mechanism designer has any idea about any player's general or relevant external knowledge!

Examples. For a first example, GK_i may consist of a subset of $\mathbb{V}_{-i}$, the set of all possible valuation sub-profiles for the players in $-i$, such that $TV_{-i} \in GK_i$. Here GK_i represents the set of possible candidates, in i's opinion, for the other players' true valuations. Such GK_i is genuine in the sense that one of its candidates is the "right one."[6] In this example, RK_i is deduced from GK_i in two conceptual steps. First, one computes all outcomes (A, P) *feasible* for GK_i, that is the outcomes such that, for all players $j \in -i$ and all valuations subprofiles $V \in GK_i$, $P_j \leq V_j(A_j)$. Then, the relevant external knowledge RK_i consistent with GK_i is the outcome with maximum revenue among the outcomes feasible for GK_i. (Thus, if $GK_i = \mathbb{V}_{-i}$ then RK_i is the null outcome.)

As for another example, GK_i may consist of a probabilistic distribution over $\mathbb{V}_{-i}$ that assigns positive probability to the actual TV_{-i}. In this case, RK_i is the outcome with the maximum revenue among all those outcomes consistent with support of GK_i.

As for a third example, GK_i may consist of a "partial" probability distribution over $\mathbb{V}_{-i}$. For instance, starting with a distribution D assigning positive probability to the actual subprofile TV_{-i}, GK_i may be derived from D as follows: when the probability p_V of each subprofile $V \in \mathbb{V}_{-i}$ is positive, then p_V is replaced with a subinterval I_V of $[0, 1]$ that includes p_V. ($I_V = [0, 1]$ is interpreted as i knowing "nothing" about profile V.) In this case, the outcomes consistent with GK_i are those consistent with the set of subprofiles V whose subinterval does not coincide with $[0, 0]$. And among such outcomes, RK_i is the one whose revenue is maximum.

[6] Notice that $GK_i = \mathbb{V}_{-i}$ expresses the fact that i knows "nothing" about TV_{-i}. Also notice that a proper choice of GK_i can precisely express pieces of i's external knowledge such as "player h's valuation for subset S is larger than player j's valuation for subset T."

Deriving Our Benchmark. Having (informally) presented the general and relevant knowledge of each player i, GK_i and RK_i, as intrinsic to i, we have in fact presented two intrinsic profiles for any combinatorial auction: the general-knowledge profile, GK, and the relevant-knowledge profile RK. Without loss of generality, alongside with the true-valuation profile TV, GK and RK can be considered integral components of the original context of any combinatorial auction. It is thus natural to consider a function mapping each RK to non-negative number, and thus a knowledge-based *benchmark*. The function chosen in this paper is the one that, given the subprofile RK_S of any subset of players S, returns the maximum revenue of the outcomes in RK_S.

Following [MV07.a], our corresponding benchmark is then defined to be this function evaluated at the subprofile RK_I, the relevant external knowledge of the *independent players*. Restricting our benchmark to the independent players expresses that we do not count on collusive players to set our goals. Indeed, we would actually consider ourselves fortunate if all collusive players "spontaneously leave the room," letting us conduct our auction with the independent players alone!

2.2 A More Formal Presentation

Let us now define a bit more formally general and relevant knowledge, and then our benchmark.

Definition 1. (Feasible External Outcomes.) *We say that an outcome* (A, P) *is a* feasible external outcome *for a player* i*, relative to a valuation profile* V*, if (1)* $A_i = \emptyset$ *and* $P_i = 0$ *and (2)* $\forall j \neq i$, P_j *is* 0 *if* $A_j = \emptyset$*, and a positive number* $< V_j(A_j)$ *otherwise.*

Notice that a feasible external outcome (A, P) for i, relative to the true-valuation profile TV, corresponds to a simple and *guaranteed* way of selling the goods to the players in $-i$. Namely, offer the subset of goods A_j to player j for price P_j: if j accepts the offer, he will receive the goods in A_j and pay P_j; else j pays nothing and receives no goods. Such a way of selling the goods is guaranteed to succeed if the players are rational. Indeed, since each non-empty subset of goods is offered at an "attractive" price, each player offered some goods should rationally accept the offer.

Definition 2. (Original Context) *The* original context *of a combinatorial auction is a triple of profiles,* (TV, GK, RK)*, where for each player* i*:*

- TV_i *is* i*'s true valuation;*
- GK_i *is the information known to* i *about* TV_{-i}*;*

- *RK_i is the outcome with maximum revenue among all feasible external outcomes for i, relative to all valuation profiles V consistent with GK_i.*

We refer to GK_i as i's general external knowledge, *and to RK_i as i's* relevant external knowledge.

Notice that our relevant external knowledge is non-Bayesian. Indeed, although the general external knowledge of a player i may naturally arise in a Bayesian setting, RK_i always is a way for i to sell the goods to the other players that succeeds with probability 1 when the players are rational.

Definition 3. *We let $\mathbb{MEW}$, the* maximum external welfare*, to be the function so defined: for any relevant-external-knowledge subprofile RK_S,*

$$\mathbb{MEW}(RK_S) = \max_{i \in S} \text{REV}(RK_i).$$

Letting I denote the set of all independent players, we define our benchmark to be $\mathbb{MEW}(RK_I)$.

Note that our benchmark is very demanding. To a designer *totally ignorant* about the players, it imposes the goal of achieving the same efficiency and/or revenue achievable by the *best informed* independent player (if he were in charge of selling of goods to the others). Therefore, the question is not whether our benchmark is meaningful, but whether there exists a mechanism capable of guaranteeing a meaningful fraction of our benchmark.[7] As we shall see, we construct mechanisms guaranteeing a fraction 1/2 of $\mathbb{MEW}(RK_I)$.

3 Our Collusion Model

Without collusion, each player i is assumed to be *individually rational,* that is acting independently so as to maximize his individual utility function u_i, mapping each possible outcome (A, P) to the real value $TV_i(A_i) - P_i$.

With collusion, we allow for the possibility that, starting with an original context (TV, GK, RK), a subset C of two or more players secretly form —for whatever reason— a *collusive set.* Such C is assumed to be *collectively rational:* that is, its members are assumed to coordinate their actions so as to maximize —based on their own "collective" knowledge, k_C— their own utility function, u_C.

[7] In light of our benchmark, a mechanism should ideally "force" all collusive players to bid the null valuation, and all independent players to bid truthfully their relevant external knowledge, so as to guarantee a fraction 1 of $\mathbb{MEW}(RK_I)$. More realistically, a mechanism should succeed (somehow) to "filter out" the bids of the collusive players, and guarantee a reasonable fraction of our benchmark.

As usual, a collusive set C has no incentives to diminish the utilites of other players, and any two outcomes in which each of its members receives the same set of goods and pays the same prices are the same from C's point of view.

Definition 4. *Given an original context* (TV, GK, RK)*, a* collusive *set consists of*

- *A subset* C *of two or more players;*
- *A binary string* k_C*, encoding information about the subprofile* TV_{-C}*; and*
- *A function* u_C *mapping any outcome to a real number such that, for any two outcomes* (A, P) *and* (A', P')*,* $u_C(A, P) = u_C(A', P')$ *whenever* $A_i = A'_i$ *and* $P_i = P'_i$ *for any player* $i \in C$*.*

We refer to u_C *as* C*'s* collective utility function*, and to* k_C *as* C*'s* collective external knowledge.

The above definition of a collusive set C is purposely a minimalist one (so as to make our results more widely applicable), but does provide C with all it needs to maximize its own collective utility. To specify a collusive model lots of questions need to be answered; in particular: how many collusive sets there may be? How are collusive sets formed? How does a collusive set coordinate the actions of its members? How does a collusive set settles on its collective knowledge? How does it settle on its utility function?

Perhaps working under specific answers to the above questions would enable one to obtain stronger results, but in this paper *we genuinely aim at maximizing the meaningfulness and applicability of our results by achieving our benchmark in a very general collusive model.* In particular,

- Our results do not depend on how many collusive sets there may be, but we insist that collusive sets be disjoint. Else, speaking of collective rationality would become quite more problematic.
- Our results do not depend on how or why collusive sets are formed. Perhaps each collusive set C was brought into existence by the Devil, who corrupted C's members and forced them to cooperate in maximizing u_C. Perhaps C was the product of an initial negotiation. Whatever the case, even if it came about in some "irrational" way, C must be collectively rational once formed.
- Our results do not depend on how the members of a collusive set C coordinate their actions. Maybe it is in their interest to follow a common plan, or they have entered a secret binding agreement specifying how to act and how to make side-payments to one another. (In any case, as we shall see, our results —perhaps counterintuitively— hold whether or not collusive sets can guarantee that their members stick to their coordinated strategies.)

- Our results do not depend on how the collective external knowledge of C arises. Indeed, there is no guarantee that k_C is related in a specific way to the individual external knowledge of C's members. For instance, if C came into existence from a negotiating process, a player i might have been welcomed in C only because he —taking some chances!— successfully boasted a general external knowledge GK_i more accurate than that he really had.

- Also, differently from an individual utility functions u_i, we do not constraint a collective utility function u_C to be of any specific form. It might be reasonable to expect that u_C is related to TV_C, that is to the individual utilities of its members, but such relationship may not be explicitly known to C's members. For instance, assume that C was formed via a negotiating process. Then, not knowing what the results of this negotiation might be, C's players might have been reluctant to reveal their true valuations for the good for sale to each other. However, they might have been able to agree on assigning a "formal" collective value to each possible outcome, and then maximizing the resulting utility function. This said, it is of course possible that u_C may be the sum of the individual utilities of C's members, that is, $u_C(A, P) = \sum_{i \in C} TV_i(A_i) - P_i$. As for another possibility, the players may agree to maximize the last u_C, but then make side payment to each other so as to make the result more fair to all. (Needless to say, such an agreement may induce the players to lie about their true valuations, so that, rather than maximizing $\sum_{i \in C} TV_i(A_i) - P_i$, they end up maximizing another function only loosely related to it!) As for a totally different example, u_C maycoincide with the individual utility function of a specific member of C (who might have convinced the others to so cooperate with him in return of a fixed payment).

However, without going as far as demanding that a collective utility function u_C be of a specific form, some general constraints on u_C are necessary to prevent modeling C's members as *irrational*.[8] The restriction on collective utility functions envisaged in our model is "individually monotonicity." By this we mean that, fixing the allocations and the prices of all players in C except for some player i, C's collective utility cannot but increase with i's individual utility. Let us now be more precise.

Definition 5. (Individually Monotone Utilities) *We say that the collective utility function u_C of a collusive set C is* individually monotone *if for all players $i \in C$, and for all pairs of outcomes (A, P) and (A', P') such that $(A_j, P_j) = (A'_j, P'_j)$ whenever $j \in C \setminus \{i\}$, we have:*

$$u_C(A, P) \geq u_C(A', P') \text{ if and only if } TV_i(A_i) - P_i \geq TV_i(A'_i) - P'_i.$$

[8]Indeed, irrational players may be modeled —see [MV07.a]— as taking arbitrary actions, and for any tuple of actions actually taken by the members of C, one might be able to find an *ad hoc* collective utility function u_C so as to rationalize their actions as maximizing that u_C.

A simple example of an individually monotone u_C consists of the sum of the individual utilities of C's members.[9] (For a more eccentric example, let u_C be the sum of: the individual utility of C's first member, half of the utility of C's second member, a third of the individual utility of C's third member, and so on.)

Individual monotonicity is the *only* restriction envisaged in our results. In particular, therefore, our results are independent of the collusive sets' collective knowledge. This perhaps surprising independence is due, as we shall see, both to individual monotonicity and our choice of mechanisms.

Let us now present our collusive contexts more precisely. For uniformity of presentation, we specify the collusive players via a partition $\mathbb{C}$ of the players: namely, a set in $C \in \mathbb{C}$ is collusive if it has cardinality greater than 1, and a player i is independent if his collusive set has cardinality 1 —that is, if $\{i\} \in \mathbb{C}$. (This way each player i, collusive or not, belongs to a single set of $\mathbb{C}$, denoted by C_i.)

Definition 6. (Individually Monotone Collusive Contexts) *In a combinatorial auction, a* collusive context $\mathscr{C}$ *is a tuple* $(TV^{\mathscr{C}}, GK^{\mathscr{C}}, RK^{\mathscr{C}}, \mathbb{C}^{\mathscr{C}}, I^{\mathscr{C}}, \mathbb{K}^{\mathscr{C}}, \mathbb{U}^{\mathscr{C}})$ *where*

- $(TV^{\mathscr{C}}, GK^{\mathscr{C}}, RK^{\mathscr{C}})$ *is the original context of the auction.*
- $\mathbb{C}^{\mathscr{C}}$ *is a partition of the players.*
- $I^{\mathscr{C}}$ *is the set of all players i such that* $\{i\} \in \mathbb{C}^{\mathscr{C}}$. (Set $I^{\mathscr{C}}$ is explicitly specified for convenience only.)
- $\mathbb{K}^{\mathscr{C}}$ *is a vector of strings indexed by the subsets in* $\mathbb{C}^{\mathscr{C}}$*: for each* $C \in \mathbb{C}^{\mathscr{C}}$, $\mathbb{K}^{\mathscr{C}}_C$ *is the collective external knowledge of C, where* $\mathbb{K}^{\mathscr{C}}_{\{i\}} = GK^{\mathscr{C}}_i$ *whenever* $i \in I^{\mathscr{C}}$.
- $\mathbb{U}^{\mathscr{C}}$ *is a vector of functions indexed by the subsets in* $\mathbb{C}^{\mathscr{C}}$*: for each subset of players* $C \in \mathbb{C}^{\mathscr{C}}$, $\mathbb{U}^{\mathscr{C}}_C$ *is the collusive utility function of C, where* $\mathbb{U}^{\mathscr{C}}_{\{i\}}((A,P), TV^{\mathscr{C}}) = TV^{\mathscr{C}}_i(A_i) - P_i$ *whenever* $i \in I^{\mathscr{C}}$.

We say that a collusive context $\mathscr{C}$ is individually monotone *if* $\mathbb{U}^{\mathscr{C}}_C$ *is individually monotone for any* $C \in \mathbb{C}^{\mathscr{C}}$.

We refer to a player in $I^{\mathscr{C}}$ as *independent*, to a player not in $I^{\mathscr{C}}$ as *collusive*, to a subset in $\mathbb{C}^{\mathscr{C}}$ with cardinality > 1 as a *collusive set*, and to $\mathbb{U}^{\mathscr{C}}$ as *the utility vector of* $\mathbb{C}^{\mathscr{C}}$. We use the term *agent* to denote either an independent player or a collusive set. For any player i, we denote by C_i the set in $\mathbb{C}^{\mathscr{C}}$ to which i belongs.

[9] Notice that a collusive set with such collective utility function should not be construed to be an example of "shilling," because in our setting the goods —as for a spectrum auction— are not necessarily transferable.

Definition 7. (Collusive Auctions.) *A* collusive auction *is a pair* $(\mathscr{C}, M)$*, where* $\mathscr{C}$ *is a collusive auction context, and* M *an auction mechanism. Such an auction is* individually monotone *if so is* $\mathscr{C}$*.*

4 Our Solution Concept

4.1 Intuition

Resilient mechanisms in particular demand solution concepts immune to any equilibrium-selection problem. The best such concept, in a sense, is dominant-strategy solvability, since in this case there is "a single equilibrium to be predicted" and no one has to rely on the rationality of the others. However, in light of the upper-bounds of [MV07.a] and [MV07.b] (holding for all dominant-strategy truthful mechanisms), to guarantee better performance we must explore other solution concepts. Short of this, another "safe" way to predict which equilibrium will be played is when the game is dominance solvable, that is, when after the iterative procedure in which at each round all dominated strategies are removed, only a single strategy profile survives. Our mechanisms do not yield such games, but guarantee an essentially equivalent property.

Implementation in Surviving Strategies. Assume that, after iteratively removing dominated strategies, plenty of *surviving strategies* remain for each agent. Then, one cannot predict with certainty which profile of strategies will be actually played. But, to us, in mechanism design predictability of the actually played strategies is *a useful mean to an end, not the goal itself.* To guarantee that a mechanism satisfies a desired property P it suffices to prove that

P holds for any possible profile of surviving strategies.

This is indeed the notion of implementation delivered by our mechanisms. We call it *implementation in surviving strategy.* (Note that, even in a non-collusive setting, our notion is more general than that of implementation in undominated strategies, as proposed by [BLP06].)

Implementation in Σ^1/Σ^2_I Strategies. Experience seems to indicate that, in practice, there are different levels of rationality; that is, that many players are capable of completing the first few iterations of elimination of dominated strategies, but fail to go "all the way." Accordingly, one should prefer mechanisms that guarantee their desired property for any vector of strategies surviving just the first few iterations. This is exactly the case for our mechanisms.

Specifically, we envisage the following two-round elimination process. First, each agent (i.e., an independent player or a collusive set) removes all his weakly dominated strategies. Then, each independent player eliminates all strategies which become weakly dominated after the first round of elimination is completed. Accordingly, our solution concept envisages each collusive player to choose his final strategy based only on his own rationality, and each independent player based on his own rationality as well as a modicum of rationality for the other players. Since the set of strategies surviving the first iteration is often referred to as Σ^1, and the set of those surviving the first two iterations is commonly referred to as Σ^2, we call this refinement of our solution concept *implementation in* Σ^1/Σ^2_I *Strategies.* For simplicity, we formalize just this latter refinement of our solution concept, and only for our auction setting.

The Difficulties with Collusion. We believe implementation in surviving and/or Σ^1/Σ^2_I strategies to be of independent interest, and we expect it to play a larger role in perfect-information and non-collusive settings. In such settings the notion is significantly easier, because it is easy to determine which strategies are dominated. In our case, instead, whether a strategy is dominated depends on such additional factors as the collusive sets actually present and their utility functions, factors about which no information is publicly available. This complicates our notion and the analysis of our mechanism.

4.2 Formalization

Our mechanisms are of a very simple form. At each decision node all players act simultaneously, and their actions become public as soon as they are chosen. Also, our mechanism are probabilistic, and their coin tosses too become of public domain as soon as they are made. Since in this paper we are considering collusion to be illegal (and thus secret), our mechanisms specify only the strategies of individual players. Note that, denoting the set of all deterministic strategies of a player i by Σ^0_i, the set of all strategy profiles by Σ^0, the set of all deterministic collective strategies of a collusive set C by Σ^0_C, the set of all deterministic strategy vectors of a collusive context $\mathscr{C}$ by $\Sigma^0_{\mathscr{C}}$, and the Cartesian product by $\prod$, we have

$$\Sigma^0 = \prod_i \Sigma^0_i, \quad \Sigma^0_C = \prod_{i \in C} \Sigma^0_i, \quad \text{and } \Sigma^0_{\mathscr{C}} = \Sigma^0.$$ [10]

To formalize implemention in Σ^1/Σ^2_I strategies, we start by adapting the standard definition of dominated and undominated strategies to collusive auctions.

[10] Indeed, for all $\mathscr{C}$ we have $\Sigma^0_{\mathscr{C}} = \prod_{C \in \mathbb{C}^{\mathscr{C}}} \Sigma^0_C = \prod_{C \in \mathbb{C}^{\mathscr{C}}} \prod_{i \in C} \Sigma^0_i = \Sigma^0$.

Definition 8. (Dominated and Undominated Strategies.) *In a collusive auction* $(\mathscr{C}, M)$, *we say that a deterministic strategy* $\sigma_{\mathcal{A}}$ *of an agent* $\mathcal{A}$ *is* dominated *over a set of strategy vectors* Σ' *if* $\sigma_{\mathcal{A}} \in \Sigma'_{\mathcal{A}}$ *and there exists* $\sigma'_{\mathcal{A}} \in \Sigma'_{\mathcal{A}}$ *such that*

1. $\forall \tau_{-\mathcal{A}} \in \Sigma'_{-\mathcal{A}}$, $E[u_{\mathcal{A}}(M(\sigma_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))] \leq E[u_{\mathcal{A}}(M(\sigma'_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))]$.
2. $\exists \tau_{-\mathcal{A}} \in \Sigma'_{-\mathcal{A}}$ *such that* $E[u_{\mathcal{A}}(M(\sigma_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))] < E[u_{\mathcal{A}}(M(\sigma'_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))]$.

Else, we say that $\sigma_{\mathcal{A}}$ *is* undominated *over* Σ'.

Remark. *$\mathcal{A}$ has no conceptual difficulty in determining whether $\sigma_{\mathcal{A}}$ is dominated over Σ'.* This is so despite the facts that dominated strategies are defined for collusive auctions $(\mathscr{C}, M)$, and that no agent $\mathcal{A}$ is assumed to have knowledge of the overall collusive context $\mathscr{C}$. It suffices for $\mathcal{A}$ to know its own (collective or individual) utility function and the strategies of all other agents specified by Σ'.[11]

Definition 9. (Compatibility.) *We say that a collusive context* $\mathscr{C}$ *is* compatible

- with an independent player i *if (1)* $i \in I^{\mathscr{C}}$, *(2)* $TV_i = TV_i^{\mathscr{C}}$, *(3)* $GK_i = GK_i^{\mathscr{C}}$, *and (4)* $RK_i = RK_i^{\mathscr{C}}$
- with a collusive set C *if (a)* $C \in \mathbb{C}^{\mathscr{C}}$ *and (b)* $\mathbb{U}_C^{\mathscr{C}}$ *is* C*'s collective utility function.*

Definition 10. (Σ^1 Strategies.) *Fix a mechanism M. Then, if $\mathcal{A}$ is an agent in a collusive auction $(\mathscr{C}, M)$, then $\Sigma^1_{\mathcal{A},\mathscr{C}}$ denotes the set of deterministic strategies of $\mathcal{A}$ undominated over $\Sigma^0_{\mathscr{C}}$.*

Remarks.

- *$\Sigma^1_{\mathcal{A},\mathscr{C}}$ is the same for any $\mathscr{C}$ compatible with $\mathcal{A}$.* In fact, as noted, $\Sigma^0_{\mathscr{C}} = \Sigma^0$ for all $\mathscr{C}$. Accordingly, we shall more simply write $\Sigma^1_{\mathcal{A}}$ instead of $\Sigma^1_{\mathcal{A},\mathscr{C}}$.
- *$\mathcal{A}$ can compute $\Sigma^1_{\mathcal{A}}$.* In fact, $\mathcal{A}$ can determine whether any of its strategies $\sigma_{\mathcal{A}}$ is dominated over any given set of strategy vectors Σ', and the set Σ^0 is publicly known because it solely depends on the publicly known mechanism M. (In a sense, to compute $\Sigma^1_{\mathcal{A}}$, it suffices for $\mathcal{A}$ to assume that all other players are independent.)
- *For any collusive context $\mathscr{C}$, the set of strategy vectors surviving the first round of elimination of dominated strategies coincides with $\prod_{C \in \mathbb{C}^{\mathscr{C}}} \Sigma^1_C$.*

[11] Note that something about the actual collusive sets can be deduced from Σ', but not enough to precisely know the collective utility functions of the other agents. But note too that this knowledge is inessential for $\mathcal{A}$ to determine whether $\sigma_{\mathcal{A}}$ is dominated over Σ'.

- $\prod_{C \in \mathbb{C}^{\mathscr{C}}} \Sigma^1_C$ *is crucially dependent on* $\mathscr{C}$*,* although each Σ^1_C only depends on C but not on the actual collusive context compatible with C.

Definition 11. (Σ^2_I Strategies for individually monotone context.) *Fix a mechanism M. If i is an independent player in an individually monotone collusive auction $(\mathscr{C}, M)$, then we denote by $\Sigma^2_{i,\mathscr{C}}$ the set of all strategies $\sigma_i \in \Sigma^1_i$ undominated over $\prod_{C \in \mathbb{C}^{\mathscr{C}}} \Sigma^1_C$. We define Σ^2_i the union of $\Sigma^2_{i,\mathscr{C}}$ for all $\mathscr{C}$ compatible with i.*

Remarks.

- *Without demanding individual monotonicity, we would have $\Sigma^2_i = \Sigma^1_i$.*
- Σ^2_i *depends solely on i.* Thus, i can determine Σ^2_i whether or not he believes that the context is collusive (indeed, i may have been involved in some preliminary negotiation about colluding), and no matter what he believes about such a possible collusive context.

Definition 12. (Σ^1/Σ^2_I Plays.) *Fix a mechanism M. We say that a strategy vector σ is a Σ^1/Σ^2_I* play *of an individually monotone collusive auction $(\mathscr{C}, M)$ if*

$$\sigma \in \prod_{i \in I^{\mathscr{C}}} \Sigma^2_i \times \prod_{C \in \mathbb{C}^{\mathscr{C}}, |C|>1} \Sigma^1_C.$$

Remark. Although each Σ^2_i depends only on i and each Σ^1_C only on C (and not on the actual collusive context), the Σ_1/Σ^2_I plays of $(\mathscr{C}, M)$ crucially depend on $\mathscr{C}$.

Definition 13. (Implementation in Σ^1/Σ^2_I Strategies) *Let $\mathbb{P}$ be a property over auction outcomes, and M an auction mechanism. We say that M* implements $\mathbb{P}$ in Σ^1/Σ^2_I strategies *if, for all individually monotone collusive contexts $\mathscr{C}$, and all Σ^1/Σ^2_I plays σ of the auction $(\mathscr{C}, M)$, $\mathbb{P}$ holds for $M(\sigma)$.*

Note that, although σ is a vector of deterministic strategies, M may be probabilistic. In this case, $M(\sigma)$ is a distribution over outcomes, and $\mathbb{P}$ a property of outcome distributions.

5 Our Result

Our result essentially states that there exists an auction mechanism that, in Σ^1/Σ^2_I strategies, implements the following property: the sum of the expected social welfare and the expected revenue is at least half of our knowledge benchmark —that is, $\frac{\mathbb{MEW}(RK_I)}{2}$. Let us now formalize and prove this statement. Since our proof is constructive, we start by presenting our mechanism.

5.1 Our Mechanism

Our mechanism is probabilistic and of extensive-form. It consists of three stages: two player stages followed by a final mechanism stage, where the mechanism produces the final outcome (A, P).

In the first stage, each player i publicly (and simultaneously with the others) announces (1) a canonical outcome Ω^i for the players in $-i$; and (2) a subset of goods S_i. (Allegedly, Ω^i is actually feasible, and indeed represents the "best way known to i to sell the goods to the other players." Allegedly too, S_i is i's favorite subset of goods, that is the one i values the most.)

After the first stage, everyone can compute (a) the revenue R_i of Ω^i for each player i, (b) the highest and second highest of such revenues, respectively denoted by $R_\star$ and R', and (c) the player whose announced outcome has the highest revenue —the lexicographically first player in case of "ties". Such player is called the "star player" and is denoted by "$\star$". (Thus, $\star \in N$.)

In the second stage, each player i, envisioned to receive a non-empty set of goods (for a positive price) in $\Omega_\star$, publicly (and simultaneously with the other such players) answers yes or no to the following implicit question: "are you willing to pay your envisioned price for your envisioned goods?" (The players not receiving any goods according to $\Omega_\star$ announce the empty string.)

After the second stage, for each asked player i who answers no, the star player is punished with a fine equal to the price he envisioned for i.

In the third and final stage, the mechanism flips a fair coin. If Heads, $S_\star$ is given to the star player at no *additional* charge (and thus player $\star$ pays nothing altogether if no player says no in the second stage). If Tails, (1) the goods are sold according to $\Omega_\star$ to the players who answered yes in the second stage, (2) all the revenue generated by this sale is given to the star player, and (3) the star player additionally pays R' to the seller/auctioneer. (Thus, the star player pays only R' if he has not been fined.) A more precise description of our mechanism is given below. In it, for convenience, we also include three "variable-update stages" and mark them by the symbol "•". In such stages the contents of some public variables are updated based on the strings announced so far.

Mechanism $\mathcal{M}$

- Set $A_i = \emptyset$ and $P_i = 0$ for each player i.

1. Each player i simultaneously and publicly announces (1) a canonical outcome for $-i$, $\Omega^i = (\alpha^i, \pi^i)$, and (2) a subset S_i of the goods.

- Set: $R_i = \textsc{rev}(\Omega^i)$ for each player i, $\star = \arg\max_i R_i$, and $R' = \max_{i \neq \star} R_i$.

  (We shall refer to player $\star$ as the "star player", and to R' as the "second highest revenue".)

2. Each player i such that $\alpha_i^\star \neq \emptyset$ simultaneously and publicly announces YES or NO.

- For each player i who announces NO, $P_\star = P_\star + \pi_i^\star$.

3. Publicly flip a fair coin.

   – If Heads, reset $A_\star = S_\star$.

   – If Tails: (1) reset $P_\star = P_\star + R'$; and (2) for each player i who announced YES in Stage 2, reset: $A_i = \alpha_i^\star$, $P_i = \pi_i^\star$, and $P_\star = P_\star - P_i$.

Comment. The outcome (A, P) may not be canonical, as the price of the star player may be non-zero even though he may receive nothing.

5.2 Analysis of Our Mechanism

In what follows, all (individual, collective and vectors of) strategies are relative to mechanism $\mathcal{M}$.

Lemma 1. $\forall$ *independent players* i *and* $\forall \sigma_i \in \Sigma_i^1$: *if* $i \neq \star$ *and* $\alpha_i^\star \neq \emptyset$ *after Stage 1, then in Stage 2*

1. *i answers YES whenever $TV_i(\alpha_i^\star) > \pi_i^\star$, and*
2. *i answers NO whenever $TV_i(\alpha_i^\star) < \pi_i^\star$.*

Proof. We restrict ourselves to just prove, by contradiction, the first implication (the proof of the second one is totally symmetric). Define the following properties of an execution of $\mathcal{M}$:

$\mathcal{P}$: $i \neq \star$, $\alpha_i^\star \neq \emptyset$, and $TV_i(\alpha_i^\star) > \pi_i^\star$.

$\overline{\mathcal{P}}$: $i = \star$, or $\alpha_i^\star = \emptyset$, or $TV_i(\alpha_i^\star) \leq \pi_i^\star$.

Assume that there exist an independent player i and a strategy profile $\sigma \in \Sigma^0$ such that (1) $\sigma_i \in \Sigma_i^1$; (2) σ's execution satisfies $\mathcal{P}$; and (3) i answers NO. Then, consider the following alternative strategy for player i:

Strategy σ_i'

Stage 1. Run σ_i (with stage input "1" and private inputs TV_i and GK_i) and announce Ω^i and S_i as σ_i does.

Stage 2. If $\overline{\mathcal{P}}$, run σ_i and answer whatever σ_i does.[12]

If $\mathcal{P}$, answer YES.

[12] The first implication of Lemma 1 specifies that $i \neq \star$ and $TV_i(\alpha_i^\star) > \pi_i^\star$. However, a strategy must be specified in all cases, and thus σ_i' must be specified also when $\overline{\mathcal{P}}$.

We derive a contradiction by proving that σ_i is dominated by σ_i' over Σ^0, which implies that $\sigma_i \notin \Sigma_i^1$.

Notice that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \overline{\tau_{-i}}))] = \mathbb{E}[u_i(\mathcal{M}(\sigma_i' \sqcup \overline{\tau_{-i}}))]$ for all subprofiles $\overline{\tau_{-i}} \in \Sigma_{-i}^0$ such that the execution of $\sigma_i \sqcup \overline{\tau_{-i}}$ either satisfies (1) $\overline{\mathcal{P}}$, or (2) $\mathcal{P}$ and i answers YES. (This is so because for such $\overline{\tau_{-i}}$ the executions of $\sigma_i \sqcup \overline{\tau_{-i}}$ and $\sigma_i' \sqcup \overline{\tau_{-i}}$ coincide, and so do their outcomes before $\mathcal{M}$'s coin toss.) Therefore to prove that σ_i is dominated by σ_i' over Σ^0, it suffices to consider the strategy subprofiles $\tau_{-i} \in \Sigma_{-i}^0$ such that the execution of $\sigma_i \sqcup \tau_{-i}$ satisfies $\mathcal{P}$ and i answers NO. (Notice that, by assumption, $\tau_{-i} = \sigma_{-i}$ is one such subprofile.)

For all such τ_{-i}, observe that, since σ_i' coincides with σ_i in Stage 1, the outcome profile Ω is the same in the executions of $\sigma_i \sqcup \tau_{-i}$ and $\sigma_i' \sqcup \tau_{-i}$. Accordingly, the star player too is the same in both executions. Since (by hypothesis) the execution of $\sigma_i \sqcup \tau_{-i}$ satisfies $\mathcal{P}$, so does the executions of $\sigma_i' \sqcup \tau_{-i}$.

We now distinguish two cases, each occurring with probability 1/2.

(1) $\mathcal{M}$'s coin toss comes up *Heads.*

In this case, because only the star player receives goods, we have

$$u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i})) = u_i(\mathcal{M}(\sigma_i' \sqcup \tau_{-i})) = 0.$$

(2) $\mathcal{M}$'s coin toss comes up *Tails.*

In this case, because by hypothesis, (1) $TV_i(\alpha_i^\star) > \pi_i^\star$, (2) player i answers NO in the execution of $\sigma_i \sqcup \tau_{-i}$ and (3) i answers YES in the execution of $\sigma_i' \sqcup \tau_{-i}$, we have

$$u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i})) = 0 \quad \text{and} \quad u_i(\mathcal{M}(\sigma_i' \sqcup \tau_{-i})) = \mathcal{TV}_i(\alpha_i^\star) - \pi_i^\star > 0.$$

Combining the above two cases yields

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\sigma'_i \sqcup \tau_{-i}))].$$

Therefore σ_i is dominated by σ'_i over Σ^0. ∎

Lemma 2. *$\forall$ individually monotone collusive sets C and $\forall \sigma_C \in \Sigma^1_C$: if $\star \notin C$ after Stage 1, then for all players i in C*

1. *i answers YES whenever $\alpha^\star_i \neq \emptyset$ and $TV_i(\alpha^\star_i) > \pi^\star_i$, and*
2. *i answers NO whenever $\alpha^\star_i \neq \emptyset$ and $TV_i(\alpha^\star_i) < \pi^\star_i$.*

Proof. We again restrict ourselves to just prove the first implication, and proceed by contradiction. Assume that there exist an individually monotone collusive set C, a player $i \in C$, and a strategy vector σ such that $\sigma_C \in \Sigma^1_C$, $\sigma_{-C} \in \Sigma^0_{-C}$, and in σ's execution i answers NO and the following property holds:

$\mathcal{P}_{i,C} : \star \notin C$, $\alpha^\star_i \neq \emptyset$, and $TV_i(\alpha^\star_i) > \pi^\star_i$.

We denote by $\overline{\mathcal{P}_{i,C}}$ the negation of this property. That is,

$\overline{\mathcal{P}_{i,C}} : \star \in C$, or $\alpha^\star_i = \emptyset$, or $TV_i(\alpha^\star_i) \leq \pi^\star_i$.

Consider the following alternative collective strategy for C.

Strategy σ'_C

Stage 1. Run σ_C and announce Ω_j and S_j as σ_C does for all $j \in C$.

Stage 2. If $\overline{\mathcal{P}_{i,C}}$, continue running σ_C and answer whatever σ_C does for all $j \in C$.
If $\mathcal{P}_{i,C}$, continue running σ_C, answer YES for i and whatever σ_C does for all $j \in C \setminus \{i\}$.

We derive a contradiction by proving that σ_C is dominated by σ'_C over Σ^0, which implies $\sigma_C \notin \Sigma^1_C$.

Similar to Claim 1, to prove that σ_C is dominated by σ'_C over Σ^0, it suffices to consider all strategy sub-vectors $\tau_{-C} \in \Sigma^0_{-C}$ such that the execution of $\sigma_C \sqcup \tau_{-C}$ satisfies $\mathcal{P}_{i,C}$ and i answers NO. (Note that by hypothesis, $\tau_{-C} = \sigma_{-C}$ is one such strategy sub-vector.) For each such τ_{-C}, for all $j \in C \setminus \{i\}$, we have $(\mathcal{M}_a(\sigma_C \sqcup \tau_{-C})_j, \mathcal{M}_p(\sigma_C \sqcup \tau_{-C})_j) = (\mathcal{M}_a(\sigma'_C \sqcup \tau_{-C})_j, \mathcal{M}_p(\sigma'_C \sqcup \tau_{-C})_j)$ whenever the final coin toss of $\mathcal{M}$ is the same. Thus, due to C's individual monotonicity, to show that $\mathbb{E}[u_C(\mathcal{M}(\sigma_C \sqcup \tau_{-C}))] < \mathbb{E}[u_C(\mathcal{M}(\sigma'_C \sqcup \tau_{-C}))]$ it suffices to prove that $u_i(\mathcal{M}(\sigma_C \sqcup \tau_{-C})) = u_i(\mathcal{M}(\sigma'_C \sqcup \tau_{-C}))$ when the coin toss of $\mathcal{M}$ comes up Heads, and that $u_i(\mathcal{M}(\sigma_C \sqcup \tau_{-C})) < u_i(\mathcal{M}(\sigma'_C \sqcup \tau_{-C}))$ when the coin toss of $\mathcal{M}$ comes up Tails. This proof is analogous to the corresponding one of Claim 1, and is ignored. ∎

Lemma 3. *$\forall$ independent player i and $\forall \sigma_i \in \Sigma_i^2$, σ_i is such that*

1. *(i does not "under-bid":) i announces Ω^i such that $\text{REV}(\Omega^i) \geq \text{REV}(RK_i)$.*

If $i \neq \star$ and $\alpha_i^\star \neq \emptyset$, then

2. *i answers YES whenever $TV_i(\alpha_i^\star) > \pi_i^\star$; and*
3. *i answers NO whenever $TV_i(\alpha_i^\star) < \pi_i^\star$.*

Proof. Properties 2 and 3 follow directly from Lemma 1 and the fact that $\Sigma_i^2 \subseteq \Sigma_i^1$.

We prove properties 1 by contradiction. Assume that $\exists$ independent player i and $\exists \sigma_i \in \Sigma_i^2$ such that in Stage 1 of σ_i, i announces Ω^i such that $\text{REV}(\Omega^i) < \text{REV}(RK_i)$. Now consider the following alternative strategy for player i.

Strategy $\widehat{\sigma}_i$

Stage 1. Announce an outcome $\widehat{\Omega}_i = (\widehat{\alpha}^i, \widehat{\pi}^i)$ and a subset of goods $\widehat{S}_i$ computed as follows:

- Run σ_i so as to compute its outcome Ω^i.
- Compute $\epsilon = \text{REV}(RK_i) - \text{REV}(\Omega^i)$
- Set $(\widehat{\alpha}^i, \widehat{\pi}^i) = RK_i$;
- $\forall j \in -i$ such that $\widehat{\alpha}_j^i \neq \emptyset$, set $\widehat{\pi}_j^i = \widehat{\pi}_j^i(1 - \frac{\epsilon}{2 \cdot \text{REV}(RK_i)})$.
- Set $\widehat{S}_i = \arg\max_{S \subseteq G} TV_i(S)$.

Stage 2. Announce YES, NO, or the empty string as follows:

- If $\star = i$ or $\alpha_i^\star = \emptyset$, announce the empty string.

  Else, announce YES if $TV_i(\alpha_i^\star) \geq \pi_i^\star$, and announce NO if $TV_i(\alpha_i^\star) < \pi_i^\star$.

We derive a contradiction in two steps, that is by proving two separate claims: namely, (1) $\widehat{\sigma}_i \in \Sigma_i^1$, and (2) σ_i is dominated by $\widehat{\sigma}_i$ over $\Sigma_{\mathscr{C}}^1$ for all individually monotone collusive contexts $\mathscr{C}$ compatible with i. The second fact of course contradicts the assumption that $\sigma_i \in \Sigma_i^2$.

Claim 1: $\widehat{\sigma}_i \in \Sigma_i^1$.

Proof: Proceeding by contradiction, let σ_i' be a strategy such that $\sigma_i' \neq \widehat{\sigma}_i$ and σ_i' dominates $\widehat{\sigma}_i$ over Σ^0. Assume that σ_i' announces

widehatOmega$^{i'} \neq \widehat{\Omega}_i$ or $S_i' \neq \widehat{S}_i$, and let σ_{-i} be the subprofile of strategies in which every player $j \in -i$ announces Ω_j such that $\text{REV}(\Omega_j) = 0$ and $S_j = \emptyset$ in Stage 1, and announces YES if $\Omega^\star = \widehat{\Omega}_i$ and $S^\star = \widehat{S}_i$, and NO otherwise. Notice that σ_{-i} clearly belongs to Σ_{-i}^0. (Indeed Σ^0 consists of what all that the players can do, independent of any rationality consideration.) Notice too however, that

i's expected utility is greater than $TV_i(\widehat{S}_i) = \frac{\max_{S \subseteq G} TV_i(S)}{2}$ under the profile $\widehat{\sigma}_i \sqcup \sigma_{-i}$, while less than or equal to it under the profile $\sigma'_i \sqcup \sigma_{-i}$. Therefore such a σ'_i can not dominate $\widehat{\sigma}_i$ over Σ^0. Accordingly, if σ'_i dominates $\widehat{\sigma}_i$, it must be that σ'_i announces the same outcome and the same subset of goods as $\widehat{\sigma}_i$ does, and thus coincides with $\widehat{\sigma}_i$ in Stage 1. Let us now consider Stage 2. There, Lemma 1 implies that the only possible difference between $\widehat{\sigma}_i$ and a dominating σ'_i consists of what the two strategies announce when $i \neq \star$, $\alpha^\star_i \neq \emptyset$ and $TV_i(\alpha^\star_i) = \pi^\star_i$: namely, $\widehat{\sigma}_i$ answers YES (by definition) and σ'_i answers NO (because it must be different from $\widehat{\sigma}_i$). But this syntactic difference does not translate into any utility difference: indeed, accepting a subset of goods and paying what your true valuation for it or receiving no goods at all and paying nothing is equivalent. Therefore no $\sigma'_i \neq \widehat{\sigma}_i$ can dominate $\widehat{\sigma}_i$ over Σ^0. In sum, $\widehat{\sigma}_i \in \Sigma^1_i$ as we wanted to show. $\square$

Claim 2: $\forall$ individually monotone collusive contexts $\mathscr{C}$ compatible with i, $\widehat{\sigma}_i$ dominates σ_i over $\Sigma^1_{\mathscr{C}}$.

Proof: To prove our claim we need to compare $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))]$ and $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ for all strategy subprofiles $\tau_{-i} \in \Sigma^1_{\mathbb{C}\setminus\{i\}}$, where $\mathbb{C}$ denotes the player partition of $\mathscr{C}$. Arbitrarily fixing such a τ_{-i}, denoting by $\Omega_j = (\alpha^j, \pi^j)$ and $\widehat{\Omega}_j = (\widehat{\alpha}^j, \widehat{\pi}^j)$ the outcomes respectively announced by a player j in the executions of $\sigma_i \sqcup \tau_{-i}$ and $\widehat{\sigma}_i \sqcup \tau_{-i}$, and denoting by R' and $\widehat{R}'$ respectively the second highest revenue in the two executions, the following four simple observations hold.

O_1: $\forall j \in -i,\ \Omega_j = \widehat{\Omega}_j$.

O_2: If $i \neq \star$ in both executions, then the star player is the same in both executions.

O_3: If $i = \star$ in both executions, then $R' = \widehat{R}'$.

O_4: If $i = \star$, then each player j offered some goods in the outcome announced by player i answers YES if his true valuation for these goods is greater than his price in such outcome, and NO if it is less.

(O_1 holds because outcomes are announced in Stage 1 where all players act simultaneously without receiving any information at all from the mechanism $\mathcal{M}$; O_2 and O_3 are immediate implications of O_1; and O_4 follows from Lemmas 1 and 2, and the fact that i does not belong to any collusive set.)

To establish that $\widehat{\sigma}_i$ dominates σ_i over $\Sigma^1_{\mathscr{C}}$, we analyze the following four exhaustive cases, again after arbitrarily fixing $\tau_{-i} \in \Sigma^1_{\mathbb{C}\setminus\{i\}}$.

Case 1: *$i \neq \star$ in the execution of $\sigma_i \sqcup \tau_{-i}$* and *$i \neq \star$ in the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$.*

In this case, by observations O_1 and O_2, $\alpha^\star_i = \widehat{\alpha}^\star_i$ and $\pi^\star_i = \widehat{\pi}^\star_i$. There are four sub-cases.

(a) $\alpha^\star_i = \emptyset$. In this sub-case we have $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = 0$.

(b) $\alpha_i^\star \neq \emptyset$ *and* $TV_i(\alpha_i^\star) = \pi_i^\star$. In this sub-case, no matter whether player i answers YES or NO in σ_i, we have $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = 0$.

(c) $\alpha_i^\star \neq \emptyset$ *and* $TV_i(\alpha_i^\star) < \pi_i^\star$. In this sub-case, by Lemma 1, i answers NO in both executions, and we have $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = 0$.

(d) $\alpha_i^\star \neq \emptyset$ *and* $TV_i(\alpha_i^\star) > \pi_i^\star$. In this sub-case, by Lemma 1, i answers YES in both executions. Thus when $\mathcal{M}$'s coin toss comes up Heads, $u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i})) = u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i})) = 0$; and when $\mathcal{M}$'s coin toss comes up Tails, $u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i})) = TV_i(\alpha_i^\star) - \pi_i^\star = TV_i(\widehat{\alpha}_i^\star) - \widehat{\pi}_i^\star = u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))$.

In sum, no matter which sub-case applies, Case 1 implies $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$.

Case 2: *$i \neq \star$ in the execution of $\sigma_i \sqcup \tau_{-i}$* and *$i = \star$ in the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$.*

In this case, let us first prove that $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \frac{TV_i(\widehat{S}_i)}{2}$. To this end, we consider the same four sub-cases as above. Namely,

(a) $\alpha_i^\star = \emptyset$. In this sub-case, $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = 0$. Therefore, since $TV_i(\widehat{S}_i) \geq 0$ by definition, we have $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \frac{TV_i(\widehat{S}_i)}{2}$ as desired.

(b) $\alpha_i^\star \neq \emptyset$ *and* $TV_i(\alpha_i^\star) = \pi_i^\star$. In this sub-case, no matter whether player i answers YES or NO, we also have $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = 0$, and thus $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \frac{TV_i(\widehat{S}_i)}{2}$.

(c) $\alpha_i^\star \neq \emptyset$ *and* $TV_i(\alpha_i^\star) < \pi_i^\star$. In this sub-case, player i answers NO, and thus $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = 0 \leq \frac{TV_i(\widehat{S}_i)}{2}$.

(d) $\alpha_i^\star \neq \emptyset$ *and* $TV_i(\alpha_i^\star) > \pi_i^\star$. In this sub-case, player i answers YES, and thus can have positive utility only when the mechanism's coin toss comes up Tails, causing player i to be assigned the subset of goods $\alpha_i^\star$ for price $\pi_i^\star$. Accordingly $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \frac{TV_i(\alpha_i^\star) - \pi_i^\star}{2} \leq \frac{TV_i(\alpha_i^\star)}{2} \leq \frac{TV_i(\widehat{S}_i)}{2}$. In fact, $\pi_i^\star$ is always non-negative, and $TV_i(\widehat{S}_i) = \max_{S \subseteq G} TV_i(S)$.

In sum, no matter which sub-case applies, we have $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \frac{TV_i(\widehat{S}_i)}{2}$.

Let us now prove that $\frac{TV_i(\widehat{S}_i)}{2} \leq \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$. In this case, i's expected utility in the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$ is the weighted sum of his utility when $\mathcal{M}$'s coin toss is Heads and his utility when $\mathcal{M}$'s coin toss is Tails.[13] Therefore, denoting by "$\sum_{j:\widehat{YES}}$" (respectively,"$\sum_{j:\widehat{NO}}$") the sum taken over every player j who answers YES (respectively, NO) in Stage 2 of the execution of

[13]Both individual utilities are expected, if the strategies of the other players are probabilistic.

$\widehat{\sigma}_i \sqcup \tau_{-i}$, we have

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = \frac{TV_i(\widehat{S}_i) - \sum_{j:\widehat{NO}} \widehat{\pi}^i_j}{2} + \frac{\sum_{j:\widehat{YES}} \widehat{\pi}^i_j - \sum_{j:\widehat{NO}} \widehat{\pi}^i_j - \widehat{R}'}{2}.$$

By construction, $\forall j \in -i$ such that $\widehat{\alpha}^i_j \neq \emptyset$, $\widehat{\pi}^i_j < TV_j(\widehat{\alpha}^i_j)$. Thus by observation O_4 every such player j answers YES in Stage 2: in our notation $\sum_{j:\widehat{YES}} \widehat{\pi}^i_j = \sum_j \widehat{\pi}^i_j$ and $\sum_{j:\widehat{NO}} \widehat{\pi}^i_j = 0$. Accordingly, we have

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = \frac{TV_i(\widehat{S}_i)}{2} + \frac{\sum_j \widehat{\pi}^i_j - \widehat{R}'}{2} = \frac{TV_i(\widehat{S}_i) + \text{REV}(\widehat{\Omega}_i) - \widehat{R}'}{2}.$$

Since $\text{REV}(\widehat{\Omega}_i) \geq \widehat{R}'$, we have $\frac{TV_i(\widehat{S}_i)}{2} \leq \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$ as desired.

Therefore Case 2 implies $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$.

Case 3: *$i = \star$ in the execution of $\sigma_i \sqcup \tau_{-i}$* and *$i = \star$ in the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$.*

In this case, similar to Case 2, i's expected utility in the execution of $\sigma_i \sqcup \tau_{-i}$ is the weighted sum of his utility when $\mathcal{M}$'s coin toss is Heads and his utility when $\mathcal{M}$'s coin toss is Tails. Therefore, denoting by "$\sum_{j:YES}$" (respectively,"$\sum_{j:NO}$") the sum taken over every player j who answers YES (respectively, NO) in Stage 2 of the execution of $\sigma_i \sqcup \tau_{-i}$, we have

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \frac{TV_i(S_i) - \sum_{j:NO} \pi^i_j}{2} + \frac{\sum_{j:YES} \pi^i_j - \sum_{j:NO} \pi^i_j - R'}{2}.$$

Since $\sum_{j:NO} \pi^i_j \geq 0$ and $\sum_{j:YES} \pi^i_j \geq 0$, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] \leq \frac{TV_i(S_i)}{2} + \frac{\sum_j \pi^i_j - R'}{2} = \frac{TV_i(S_i) + \text{REV}(\Omega^i) - R'}{2}.$$

Let us now analyze i's expected utility in the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$. Same as in Case 2, and by observation O_3, we have that

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] = \frac{TV_i(\widehat{S}_i) + \text{REV}(\widehat{\Omega}_i) - \widehat{R}'}{2} = \frac{TV_i(\widehat{S}_i) + \text{REV}(\widehat{\Omega}_i) - R'}{2}.$$

According to our construction of $\widehat{\sigma}_i$, we have that: (1) $\text{REV}(\widehat{\Omega}_i) = \text{REV}(RK_i) - \epsilon/2 > \text{REV}(RK_i) -$

$\epsilon = \text{REV}(\Omega^i)$; and (2) $TV_i(\widehat{S}_i) = \max_{S\subseteq G} TV_i(S)$. Therefore

$$\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))] > \frac{TV_i(S_i) + \text{REV}(\Omega^i) - R'}{2}.$$

In sum, Case 3 implies $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] < \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$.

Case 4: *$i = \star$ in the execution of $\sigma_i \sqcup \tau_{-i}$* and *$i \neq \star$ in the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$.*

Fortunately, this case can never happen. Since $\text{REV}(\widehat{\Omega}_i) > \text{REV}(\Omega^i)$ (by construction) and $\forall j \in -i$ $\Omega_j = \widehat{\Omega}_j$ (by observation O_1), we have that if $i = \star$ in the execution of $\sigma_i \sqcup \tau_{-i}$, it must be true that $i = \star$ also in the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$.

Having finished to analyze all possible cases, we conclude that σ_i is dominated by $\widehat{\sigma}_i$ over $\Sigma^1_{\mathscr{C}}$. $\square$

Since both Claims 1 and 2 hold, so does Lemma 3. ■

Comment. Note that, while ruling out the under-bidding (relative to RK_i) of independent players, our analysis says nothing about the possibility of "over-bidding." In fact, assume that player i's general knowledge GK_i includes some Bayesian information about about the true valuation of another player j that enables i to compute the probability that $TV_j(S) > v$ for some subset S of goods and a particular value v. Then, depending on such probability and v, rather then announcing the outcome $\Omega^i = RK_i$, i may expect be better off announcing Ω^i such that $\alpha^i_j = S$ and $\text{REV}(\Omega^i) > \text{REV}(RK_i)$ (taking into account the probability that j may reject this offer.) Therefore over-bidding may not be a dominated strategy for player i over $\Sigma^1_{\mathscr{C}}$. But as shown in the following proof, if a player over-bids, our result still holds, and thus we do not care whether over-bidding is dominated or not.

We are finally ready to formally state and prove our main theorem .

Theorem 1. *$\forall$ individually monotone collusive contexts $\mathscr{C}$ and $\forall$ Σ^1/Σ^2_I plays σ of $(\mathscr{C}, \mathcal{M})$,*

$$\mathbb{E}[\text{REV}(\mathcal{M}(\sigma))] + \mathbb{E}[\text{SW}(\mathcal{M}(\sigma), TV)] \geq \frac{\mathbb{MEW}(RK_I)}{2}.$$

Proof. Denote by $*$ the *independent* player "realizing" our benchmark: that is,

$$* = \arg\max_{i\in I} \text{REV}(RK_i).$$

Notice that the players $*$ and $\star$ need not to coincide, and notice that the following two inequalities hold in any Σ^1/Σ^2_I play of $(\mathscr{C}, \mathcal{M})$:

(a) $\textsc{rev}(\Omega^*) \geq \textsc{rev}(RK_*)$.

(b) $R_\star \geq \mathbb{MEW}(RK_I)$.

Indeed, inequality (a) holds because that $*$ is independent and by Lemma 3 it does not under-bid; and inequality (b) holds by inequality (a) and the fact that$R_\star \geq \textsc{rev}(\Omega_*)$ by the very definition of the star player.

To prove our theorem, we distinguish two cases.

Case 1: $\star = *$.

In this case, as player $*$ is independent, so is player $\star$, and thus $\star \notin C$ for all collusive sets C in the player partition of $\mathscr{C}$. Therefore Lemma 2 and 3 guarantees that every $i \neq \star$ answers YES only if $TV_i(\alpha_i^\star) \geq \pi_i^\star$. Accordingly, the following inequality holds for $\mathcal{M}$'s expected social welfare:

$$\mathbb{E}[\textsc{sw}(\mathcal{M}(\sigma), TV)] = \frac{TV_\star(S_\star)}{2} + \frac{\sum_{i:YES} TV_i(\alpha_i^\star)}{2} \geq \frac{\sum_{i:YES} TV_i(\alpha_i^\star)}{2} \geq \frac{\sum_{i:YES} \pi_i^\star}{2}.$$

At the same time,

$$\mathbb{E}[\textsc{rev}(\mathcal{M}(\sigma))] = \frac{\sum_{i:NO} \pi_i^\star}{2} + \frac{R' + \sum_{i:NO} \pi_i^\star}{2} \geq \frac{\sum_{i:NO} \pi_i^\star}{2}.$$

Thus

$$\mathbb{E}[\textsc{sw}(\mathcal{M}(\sigma), TV)] + \mathbb{E}[\textsc{rev}(\mathcal{M}(\sigma))] \geq \frac{\sum_{i:YES} \pi_i^\star + \sum_{i:NO} \pi_i^\star}{2} = \frac{\sum_i \pi_i^\star}{2} = \frac{R_\star}{2} \geq \frac{\mathbb{MEW}(RK_I)}{2}.$$

Case 2: $\star \neq *$.

In this case, $* \in -\star$. Thus, in virtue of inequality (a) and the fact that $R' \geq \textsc{rev}(\Omega_*)$, $\mathcal{M}$'s expected revenue is

$$\mathbb{E}[\textsc{rev}(\mathcal{M}(\sigma))] = \frac{\sum_{i:NO} \pi_i^\star}{2} + \frac{R' + \sum_{i:NO} \pi_i^\star}{2} \geq \frac{R'}{2} \geq \frac{\mathbb{MEW}(RK_I)}{2}.$$

Of course

$$\mathbb{E}[\textsc{sw}(\mathcal{M}(\sigma), TV)] = \frac{TV_\star(S_\star)}{2} + \frac{\sum_{i:YES} TV_i(\alpha_i^\star)}{2} \geq 0.$$

Thus summing term by term we have

$$\mathbb{E}[\textsc{sw}(\mathcal{M}(\sigma), TV)] + \mathbb{E}[\textsc{rev}(\mathcal{M}(\sigma))] \geq \frac{\mathbb{MEW}(RK_I)}{2}.$$

Q.E.D.

6 Discussions and Extensions

- Why extensive-form is preferable to normal-form? Because in the corresponding normal-form mechanism the players will lose a lot of privacy and efficiency. Moreover, in the normal-form mechanism, the collusive sets may be revealed. If the collusive players do not want to be punished, they may want to deviate and choose a collective strategy which has been eliminated in the extensive-form mechanism.

- Another point, related to the privacy mentioned above is that our mechanism is *"tax-free"*. Consider a second-price (single-item) auction in which the winner bids \$10M and wins the item for \$1M, because this was the second-highest bid. Then, "uncle Sam" may want to collect taxes on \$9M, reasoning that the winner himself, being rational in a dominant-strategy truthful mechanism, freely admitted that he is receiving a \$10M value. Notice too that in an English ascending auction this cannot happen. Indeed, the players who drop out reveal their true valuations, but are not "taxable" because they have no utility. As for the winner of this alternative auction, he could also declare that his value for the object was (assuming the same valuations above) exactly \$1M (plus \$1 if he really feels to look more "legitimate"). In this case, therefore, there is nothing to "tax".

  In a combinatorial auction, in principle "taxation" could be a problem, as something about the true valuations and thus about the individual utilities could transpire for all the complex bidding of the players. Notice, however that this is NOT the case for our mechanism. In essence if the coin ends up Heads then the star player receives for free his favorite subset S, but he never said anything HIMSELF about his own valuation for S. If the coin ends up Tails, then every player who answers YES receives goods that he may always claim to value for exactly what he was offered to pay and indeed paid.

  The hypothetical point brings home the value of privacy in mechanism design. Indeed, if auctions were "taxable" then DST mechanisms may not be truthful after all. That is, ignoring privacy altogether may distort incentives so as to cause the mechanisms to fail to achieve their desired properties.

  For a full treatment of privacy and mechanism design, the reader is addressed to [ILM08].

- In $\mathcal{M}$, player i publicly announce Ω^i and S_i in Stage 1. An alternative way is to ask the players to announce the revenue of Ω^i in Stage 1, and only ask the star player to announce his $\Omega^\star$ and $S_\star$. This alternative preserves players privacy even better than $\mathcal{M}$ does, but in this way, the star player may announce $\Omega^\star$ according to the revenues announced by the other players, and an independent player i may have incentive to under-bid, that is, to announce a revenue less than $\text{REV}(RK_i)$.
- In our current benchmark, the relevant knowledge of a player is the revenue that he knows he can guarantee if he were in charge to sell the goods to *the other players.* With additional work, our mechanism can be extended to work with a more demanding benchmark: in essence the maximum social welfare known to the rational and independent players. (In other words, the relevant knowledge of a player "includes the possibility of giving himself some of the goods.")
- We are currently investigating similar solutions to other mechanism-design problems, including provision of a public good.

References

[AM06] L.M. Ausubel and P. Milgrom, The Lovely but Lonely Vickrey Auction. *Combinatorial Auctions*, pages 17-40, 2006.

[BBH+05] M.F. Balcan, A. Blum, H. Hartline, and Y. Mansour, Mechanism Design via Machine Learning. *Foundations of Computer Science*, 2005.

[BBM07] M.F. Balcan, A. Blum, and Y. Mansour, Single Price Mechanisms for Revenue Maximization in Unlimited Supply Combinatorial Auctions. *CMU-CS-07-111*, 2007.

[BLP06] M. Babaioff, R. Lavi and E. Pavlov, Single-Value Combinatorial Auctions and Implementation in Undominated Strategies, SODA'06, pages 1054-1063, 2006.

[CM88] J. Cremer and R.P. McLean, Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. Econometrica, Vol.56, No.6, pages 1247-1257, Nov., 1988.

[FGH+02] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin, Competitive Generalized Auctions. *Symposium on Theory of Computing* 2002.

[FPS00] J. Feigenbaum, C. Papadimitriou, and S. Shenker, Sharing the Cost of Multicast Transmissions. *Symposium on Theory of Computing*, pages 218-226, 2000.

[GH05] A. Goldberg and J. Hartline, Collusion-Resistant Mechanisms for Single-Parameter Agents. *Symposium on Discrete Algorithms*, 2005.

[GHK+05] V. Guruswami, J.D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, On Profit-Maximizing Envy-free Pricing. *Symposium on Discrete Algorithms*, pages 1164-1173, 2005.

[GHK+06] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright, Competitive Auctions. *Games and Economic Behavior*, 2006.

[ILM08] S. Izmalkov, M. Lepinski, and S. Micali, Perfect Implementation of Normal-Form Mechanisms. Technical Report, MIT-CSAIL-TR-2008-028, May 2008. (A preliminary version appears in FOCS'05, pages 585-595, with title "Rational secure computation and ideal mechanism design".)

[JV01] K. Jain and V. Vazirani, Applications of Approximation Algorithms to Cooperative Games. *Symposium on Theory of Computing* 2001.

[LS05] A. Likhodedov and T. Sandholm, Approximating Revenue-Maximizing Combinatorial Auctions. *AAAI Conference on Artificial Intelligence*, pages 267-274, 2005.

[MS01] H. Moulin and S. Shenker, Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, Vol.18, No.3, pages 511-533, 2001.

[MV07.a] S. Micali and P. Valiant, Revenue in Truly Combinatorial Auctions and Adversarial Mechanism Design, Technical Report, MIT-CSAIL- TR-2008-039, 2008.

(Previous versions include: MIT-CSAIL-TR-2007-052, 2007; submission to FOCS 2007, submission to SODA 2008; and the versions deposited with the Library of Congress in 2007)

[MV07.b] S. Micali and P. Valiant, Leveraging Collusion in Combinatorial Auctions, Unpublished Manuscript. 2007

[OR94] M. J. Osborne and A. Rubinstein, A Course in Game Theory, MIT Press, 1994.