Computer Science and Artificial Intelligence Laboratory
Technical Report

MIT-CSAIL-TR-2008-071 December 2, 2008

Resilient Knowledge-Based Mechanisms For 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

CSAIL

Resilient Knowledge-Based Mechanisms For Truly Combinatorial Auctions (And Implementation in Surviving Strategies)

Jing Chen and Silvio Micali

CSAIL, MIT

{jingchen, silvio}@csail.mit.edu

November 17, 2008

Abstract

We put forward a new mechanism achieving a high benchmark for (both revenue and) the sum of revenue and efficiency in truly combinatorial auctions. Notably, our mechanism guarantees its performance

- in a very adversarial collusion model;
- for any profile of strategies surviving the iterated elimination of dominated strategies; and
- by leveraging the knowledge that the players have about each other (in a non-Bayesian setting).

Our mechanism also is computationally efficient, and preserves the players' privacy to an unusual extent.

1 Introduction

The Problems of Traditional Mechanism Design. Traditional mechanism design achieves a desired property $\mathbb{P}$ "at equilibrium" and suffers from two main problems: (1) *Equilibrium Selection* and (2) *Collusion.*

The first problem is connected with the existence of multiple "reasonable" equilibria. Indeed, even if the game is designed so that $\mathbb{P}$ holds at each of the possible equilibria, if some players believe that the equilibrium which will be played out is σ while others believe that it is τ, then the profile of strategies actually selected will not be an equilibrium at all (but rather a mix and match of σ and τ), and $\mathbb{P}$ may not hold.

Equilibrium selection ceases to be a problem when $\mathbb{P}$ is achieved in dominant strategies, that is when $\mathbb{P}$ holds at an equilibrium σ such that, for all player i, σ_i is the best strategy for i no matter what strategies the other players may select. (In this case, in fact, whether or not other equilibria exist, one can confidently predict σ to be the one actually played out by rational players.) Even so, however, $\mathbb{P}$ may not be guaranteed at all in the presence of collusion. Indeed equilibria (even dominant-strategy ones) are very fragile notions: they only imply that no single player has any incentive to deviate from his envisaged strategy, but two or more players may have all the incentives in the world to *jointly deviate* from their equilibrium strategies.

The aim of this paper is to find new solutions to these fundamental problems for the case of truly combinatorial auctions[1].

Truly Combinatorial Auctions and the Power of Collusion. Truly combinatorial auctions are very general: there are multiple goods for sale and any player may value any subset of them for an arbitrary amount. These general auctions are quite difficult to work with. (To keep their complexity "confined", we assume that no resale of the goods is allowed.) As for all auctions, the traditional goals are to maximize revenue or economic efficiency. However, no revenue guarantee for truly combinatorial auctions was known until this conference [MV'08]. As for economic efficiency, it is achieved in dominant strategies by the famous VCG mechanism [V'61, C'71, G'73], but in a way that is totally vulnerable to collusion. Indeed, as shown by Ausubel and Milgrom [AM'06], even two minimally informed collusive players, who minimally value the goods for sale, can coordinate their bids so that the mechanism assigns to them all goods for a price of 0.

Resilient Mechanism Design and Its Limitations. In [MV'08] a mechanism is called *resilient* if it guarantees its desired property *without any equilibrium-selection and/or collusion problems.* In their framework, collusion is unrestricted, but nothing is expected from collusive players: that is, ideally all collusive players disappear by magic, leaving whatever mechanism to be run with just the *independent* players. Thus the designer of a mechanism has no responsibility if all players are collusive, but is fully responsible so long as a single independent player exists. We adopt the same principle. Technically, focusing on the case of a revenue oriented auction, this principle translates into the fact that a mechanism should aim at guaranteeing the highest possible fraction of a revenue benchmark B evaluated on just the bids of the independent players. (Of course, the mechanism should provide such a guarantee without knowing which players are independent, by solely relying on a properly designed incentive structure.)

Although providing a resilient mechanism for truly combinatorial auctions, [MV'08] also provides a very draconian upperbound on the revenue that can be guaranteed in a truly combinatorial auction by any dominant-strategy truthful (DST) mechanism. Essentially, 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)[2], they prove that

[1] In a combinatorial auction context, there are n players and a set G of m goods. Each player i has a *valuation* for the goods —a mapping from subsets of G to nonnegative reals—, denoted as TV_i. The *profile* (i.e., a vector indexed by the players) TV is called the true valuation profile of the auction. An *outcome* 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. A_0 is referred to as the set of unallocated goods, A_i is the set of goods allocated to i, and P_i is the price of i. Relative to an outcome (A, P), the *social welfare* is defined by the function $\text{SW}(A, TV) = \sum_i TV_i(A_i)$, and the *revenue* is defined by $\text{REV}(A, P) = \sum_i P_i$.

[2] For any valuation (sub)profile V, letting $\text{MSW}(V) = \max_A \text{SW}(A, V)$, and letting the "star" player, $\star$, be defined as $\star =$

For any truly combinatorial auction with n players and m goods, and any DST mechanism M, there exists a bid profile BID such that the (expected) revenue of $M(BID)$ is $O(\frac{\mathbb{MSW}_{-\star}(BID)}{\log \min\{n,m\}})$.

This revenue bound is 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 their bound implies that one has only two alternatives:

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

A2. Adopting a solution concept weaker than dominant strategies.

Since maximizing revenue is one of our goals, to bypass the revenue upperbound of [MV'08], it is necessary for us to take at least one of these alternatives. Actually, in this paper we take *both*, but *without* violating the basic principle of mechanism design (i.e., "all knowledge resides with the players") and *without* introducing any equilibrium-selection problem.

1.1 Our Contributions

New Goals. The goal of this paper is to guarantee both high *revenue* and high *total performance* in any truly combinatorial auction. By "total performance" we mean the sum of social welfare and revenue. Since maximizing revenue is well understood, let us clarify how our second goal differs from (and may often be preferable to) the classical goal of maximizing social welfare.

The 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. As already mentioned, in absence of collusion, the VCG mechanism perfectly achieves this classical goal. In doing so, the VCG imposes prices to the players, but such prices are almost an "after thought," or a "necessary evil:" they are just means to maximize social welfare. But what is wrong with revenue? Even a benevolent government could transform it into roads, hospitals and other infrastructure from which everyone benefits. Taking this point of view, in addition to seeking resilient mechanisms for revenue, we also seek resilient mechanisms for the *sum* of social welfare and revenue.

Note that, in light of the cited counterexample of Ausubel and Milgrom, the VCG mechanism is not adequate for total performance, and for two reasons. First, even in absence of collusion, the VCG is not optimal for this less traditional goal: in a sense, it achieves it only *"within a factor of 2,"* because it can maximize social welfare returning 0 revenue. Second, in the presence of collusion, both the VCG's returned social welfare and revenue can be (essentially) 0.

Note too that, although in a rational setting "revenue lowerbounds social welfare," a resilient mechanism aiming at maximizing revenue may not maximize the sum of revenue and social welfare. This is so because, in order to guarantee revenue in the presence of collusive players, the mechanism may have to give up some efficiency. Consequently, the social welfare of a resilient revenue-oriented mechanism M may be exactly equal to M's revenue, so that the total social-welfare-plus-revenue is just twice a modest revenue. However, by directly aiming at maximizing total performance, a resilient mechanism may actually do much better.

A New, Knowledge-Based, Benchmark. Designing mechanisms achieving our goals might be easier with some information about the players' true valuations, but acquiring this information may be too hard.[3] We thus assume that *the designer has no knowledge whatsoever about the players, but that the players have some knowledge about each other.* In our setting each player i not only has *internal knowledge*, that is knowledge of his own true valuation TV_i, but also some *external knowledge*, that is some information about

$\arg\max_i \max_{S\subseteq G} V_i(S)$, then $\mathbb{MSW}_{-\star}(V) = \text{MSW}(V_{-\star})$.

[3] In particular, for auctions of a single good, Cremer and McLean [CM'88] have fully captured the information structure needed to generate the maximum possible revenue, but concluded that acquiring this information would be too difficult for their result to be of practical use.

TV_{-i}, the other players' true valuations. (This is without any loss of generality, since the external knowledge of i may be "empty.") While a player i's external knowledge could be very *general*, and is indeed denoted by GK_i, the *relevant* knowledge, denoted by RK_i, is "how well he can sell the goods to the other players" via take-it-or-leave-it offers. (That is, RK_i is an outcome for an auction whose only players are those in $-i$, where everyone only pays if he receives some goods, and pays no more than his true value for the received goods.) That is, we benchmark the performance of a *totally ignorant* seller with the revenue obtainable by the *best informed* independent player. Again, our benchmark focuses on independent players because we too do not count on collusive ones to achieve our goals, and consider ourselves fortunate if they spontaneously leave the auction.

A New and General Collusive Model. We envisage a very adversarial collusion model. In particular, we allow any number of collusive players, we let them form any number of —disjoint— collusive sets, we do not restrict the cardinality of collusive sets, and we do not restrict the way in which the members of a collusive set coordinate their actions. (If they so want, the members of a collusive set may enter binding agreements on how to act.)

We insist, however, that all players be *rational.* As usual, an *independent player*, one not belonging to any collusive set, is *individually rational*, that is he acts so as to maximize his *individual utility function* u_i, mapping each possible outcome (A, P) to the real value $TV_i(A_i) - P_i$. A collusive set C is *collectively rational,* that is its members coordinate their actions so as to maximize their own *collective utility function* u_C, mapping any outcome (A, P) to a real number.

To maximize meaningfulness, we want the relationship between u_C and the individual utility functions of C's members to be as general as possible, provided that we do not "transform collusive players into irrational ones." (Indeed, what is the difference between saying that (1) C is a set of crazy players and (2) C is a set of rational players acting to maximize a crazy collective utility function u_C?[4]) Accordingly, we demand that u_C be *minimally monotone.* Let us explain. Consider two outcomes that are absolutely identical, as far as C's members are concerned, except for member i who receives no goods and pay nothing in the first outcome, but receives a subset of goods A_i for a price P_i in the second one. Then, minimal monotonicity requires that C prefer the second outcome if $P_i < TV_i(A_i)$.[5]

Minimal monotonicity is of course a restriction on C's collective utility function, but: (a) it is the only restriction to our otherwise general collusion model; and (b) it is a very reasonable restriction.[6]

A New Solution Concept. Although not DST, the mechanism designed in this paper is immune to any equilibrium-selection problem. The reason 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. We call such a solution concept *implementation in surviving strategies.* After weakly dominated strategies are removed iteratively, each player is left with a plurality of surviving strategies, and ultimately he chooses one of them to play. Thus, with an implementation in surviving strategies, it is quite possible that the profile of strategies actually played is not an equilibrium at all. Yet, the desired property is guaranteed just the same: any profile of "not-dumb" strategies will do.

[4]Indeed, irrational players may be modeled as taking arbitrary (i.e., universally quantified) 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.

[5]For example, a minimally monotone u_C may consist of the sum of the individual utilities of C's members. As for a more eccentric example, u_C may 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.

[6]In a sense, since each of them receives exactly the same goods for exactly the same price, the other members of C —if consulted when choosing u_C— should have no reason to object against i's receiving goods that he values more than he pays. If side-payments were possible, then they could request additional compensation from a happier i. And if side-payments were not possible, once they have the same goods and pay the same price, then under traditional economic models "making i unhappier would not make them happier." In any case, mechanism design presupposes the players' rationality, and as we said some formal restriction is needed to prevent collusion from becoming de facto indistinguishable from irrationality.

Therefore, differently from a Nash equilibrium, our solution concept does not rely at all on the players' *beliefs,* but only on their *rationality.* As for another point, in our implementation it is hard to predict precisely which profile of strategies will be ultimately played. But while "strategy predictability" has always been the cornerstone of traditional mechanism design, it has always been a *mean to an end*, not *the end itself.* Ultimately, in an auction,

we do not care about predicting strategies, but we care a lot about predicting revenue and efficiency.

Our solution concept generalizes the one of implementation in undominated strategies, even in absence of collusion, and applies to the presence of collusion as well. (See Section 2 for details.)

Our Mechanism. In sum, we exhibit a *very resilient* mechanism $\mathcal{M}$ for truly combinatorial auctions. As long as the players select strategies surviving iterated elimination of dominated strategies, $\mathcal{M}$ guarantees, under a very general collusion model and without any information about the players, a total performance within a factor of 2 of the revenue obtainable by the best informed independent player.

2 Prior Work

A large body of work has been devoted to protect auctions against collusion. Notably, [JV'01, MS'01, FPS'00] proposed *group strategy-proof* mechanisms, which are robust against collusive players unable to make side-payments to each other. By contrast, we do not envisage such restrictions in our paper. (In particular, we allow collusive players to make payments to one another and/or enter secret and binding contracts.)

The paper of [GH'05] proposes the notion of a *c-truthful* mechanism, for which no collusive set with at most c members can collectively gain more than what they could get by acting individually. However, the only mechanisms satisfying this notion are those that, independent of any bids, offer any subset of goods S to any player i for a fixed price $p_{S,i}$. Such mechanisms, therefore, are far from maximizing revenue if one does not have a proper Bayesian information about the players' true valuations for the goods. In the same paper the authors also propose mechanisms that are free to choose more general outcomes (so as to approximate maximum revenue), but satisfy a weaker notion of collusion resilience and apply to restricted auctions: namely, single good in unlimited supply. By contrast, we do not put any restriction on the cardinality of the collusive sets, nor restrict the types of combinatorial auctions. Other notable resilient mechanisms for a variety of restricted auctions are due to [LOS'02, BLP'06, BBM'07, GHKKKM'05, LS'05, FGHK'02, GHKSW'06, BBHM'05]. Once more, however, we do not rely on any auction restriction.

We also note that all the above mentioned papers aim at guaranteeing high efficiency and/or revenue, but not —as we do— total performance. Further, the benchmarks of all these prior works are expressed in terms of the players' valuations. As far as we know we are the first to work with a benchmark based on the players' knowledge.

Finally, all prior works mentioned above adopt dominant-strategies as a solution concept. A solution concept relevant to ours is the classical one of *implementation in undominated strategies.* In essence, in our language, a mechanism M achieves a property $\mathbb{P}$ in undominated strategies if $\mathbb{P}$ holds for any outcome obtained by running M on a profile of undominated strategies. (See Jackson [J'92] for a formal version.) This notion, however, was never exemplified in any setting of incomplete information, let alone in auctions. Babaioff, Lavi and Pavlov [BLP'06] both proposed a *feasible* variant of this notion (in essence, each player can compute his undominated strategies efficiently) and provided the first (and efficient too) mechanism satisfying it for a restricted type of combinatorial auctions. Namely, their mechanisms applies to auctions in which each player i has only two possible values for any subset of the goods: either 0 or a fixed value v_i. In sum, the solution concept of [BLP'06] requires less rationality than ours, but their mechanism does not address collusion at all, and does not apply to truly combinatorial auctions.

3 Our Knowledge Benchmark

Recall that we denote by GK_i the *general external knowledge* of a player i, and by RK_i i's *relevant external knowledge*, properly deduced from (and thus compatible with) GK_i. Let us consider a few examples.

1. GK_i consists 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."[7] 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 with integer prices such that, for all players $j \in -i$ and all valuations subprofiles $V \in GK_i$, $P_j < V_j(A_j)$. 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.)
2. GK_i consists 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 feasible for the support of GK_i.
3. GK_i consists 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 is 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 feasible for 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.

In sum, alongside with the true-valuation profile TV, we consider the profiles GK and RK to be integral components of the original context of any combinatorial auction.

An Important Clarification. It is crucial to clarify that, although our benchmark is solely based on the relevant external knowledge of the players, our results do not assume that the players' sole knowledge is the one relevant to our benchmark (i.e., that $GK_i = RK_i$ for all i), nor that each GK_i has a specific form. Such assumptions may be very convenient for designing mechanisms, but very unrealistic for running these mechanisms. Indeed, we expect the players to have all kinds of external knowledge in addition to the one relevant to us, and to rationally act relying on *all the knowledge available to them* once a mechanism is chosen. Accordingly, to enhance the meaningfulness of our results, we do not restrict the players' external knowledge at all. That is,

Our mechanism achieves our RK_i-based benchmark for all possible GK_is consistent with the RK_is.

(We are fully aware, of course, that better performance could be guaranteed by assuming some suitable restriction for the players' external knowledge.) Let us now be more precise.

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

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.

[7] 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."

Definition 2. (Original Context) *The* original context *of a combinatorial auction is a triple of profiles, (TV, GK, RK), where for each player i (1) TV_i is i's true valuation; (2) GK_i is the information known to i about TV_{-i}; and (3) 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. (The $\mathbb{MEW}$ Benchmark) *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 $\mathbb{MEW}(RK_I)$.

4 Our Collusion Model

Definition 4. *Given an original context (TV, GK, RK), a* minimally monotone collusive set *consists of a subset C of two or more players and a function u_C from outcomes to real numbers such that, $\forall i \in C$ and $\forall$ outcomes (A, P) and (A', P') for which (1) $(A_j, P_j) = (A'_j, P'_j)$ $\forall j \in C \setminus \{i\}$ and (2) $A'_i = \emptyset$ and $P'_i = 0$:*

$$u_C(A, P) \geq u_C(A', P') \text{ if and only if } TV_i(A_i) - P_i \geq 0.$$

Notice that the original context specifies the knowledge of all players, including the members of C. Yet, we do not specify how this individual knowledge is combined (if at all) to form a possible "collective knowledge" of C. This is so for two reasons. First, such a knowledge may be hard to predict, as in principle there is no way to guarantee that C's members truthfully reveal their private knowledge to each other. (For instance, if C arose from an initial negotiation, then a member i of C might have had incentives to lie about his knowledge in order to enter C and/or influence in his favor the choice of u_C.[8]) Second, we do not need any assumption on such knowledge to achieve our results. The minimal monotonicity of u_C is all we require.

Let us now formally discuss the possibility of having a plurality of collusive sets. Here, our only restriction is the disjointness of the collusive sets. (Else, discussing collective rationality would become more problematic.) For uniformity of presentation, we specify the collusive as well as the independent players via a partition $\mathbb{C}$ of the set of all players: namely, a player set 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 5. (Minimally Monotone Collusive Contexts and Auctions) *A* minimally monotone collusive context *$\mathscr{C}$ is a tuple $(TV, GK, RK, \mathbb{C}, I, \mathbb{U})$ where*

- *(TV, GK, RK) is the original context of the auction.*
- *$\mathbb{C}$ is a partition of the players.*
- *I is the set of all players i such that $\{i\} \in \mathbb{C}$* (explicitly specified for convenience only).
- *$\mathbb{U}$ is a vector of functions, indexed by the subsets in $\mathbb{C}$, such that (1) $\mathbb{U}_C$ is a minimally monotone collective utility function if C's cardinality is > 1, and (2) $\mathbb{U}_{\{i\}}(A, P) = TV_i(A_i) - P_i$ if $i \in I$.*

A minimally monotone collusive auction *is a pair $(\mathscr{C}, M)$, where $\mathscr{C}$ is a minimally monotone collusive auction context and M is an auction mechanism.*

[8] Perhaps better results may be obtained by restricting the collective knowledge (or the process of coalition formation) but these possibilities are not investigated in this paper.

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* and to its corresponding $\mathbb{U}_C$ as C's collective utility function, and $\mathbb{U}$ as *the utility vector of* $\mathbb{C}$. 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 $\mathbb{C}$ to which i belongs.

If $\mathscr{C}$ is a minimally monotone collusive context whose components have not been explicitly specified, then by default we assume that $\mathscr{C} = (TV^{\mathscr{C}}, GK^{\mathscr{C}}, RK^{\mathscr{C}}, \mathbb{C}^{\mathscr{C}}, I^{\mathscr{C}}, \mathbb{U}^{\mathscr{C}})$.

5 Our Solution Concept

Implementation in Σ^1/Σ^2_I Strategies. In practice, there seem to be different levels of rationality. That is, 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. Our mechanism achieves our benchmark for any vector of strategies surviving the following two-round elimination process. First, each agent (i.e., each independent player and each 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. 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.

Difficulties with Collusion. Before formalizing our solution concept, it is worth noticing that such a solution concept is of independent interest, and we expect it to play a larger role in perfect-information and non-collusive settings. In such settings our notion is in fact *significantly easier,* because it is easy to determine which strategies are dominated. In the present setting, instead, whether a strategy is dominated depends on such additional factors as the collusive sets actually present and their collective utility functions, factors about which no information is publicly available.

The presence of collusion significantly complicates both our notion and the analysis of our mechanism, and significantly "increases" the number of surviving strategies. Indeed, since we are also dealing with players who secretly collude and optimize secret collective utility functions, it is hard for a player or a collusive set to "dismiss more than just a handful of strategies at each iteration."

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 mechanisms are probabilistic, and their coin tosses too become of public domain as soon as they are made. Finally, since in this paper we are dealing with secret collusion, our mechanisms specify only the strategies of individual players. Accordingly, denoting the set of all deterministic strategies of a player i by Σ^0_i, the set of all deterministic 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

$$\Sigma^0 = \prod_i \Sigma^0_i, \quad \Sigma^0_C = \prod_{i \in C} \Sigma^0_i, \quad \text{and } \Sigma^0_{\mathscr{C}} = \Sigma^0.^9$$

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 (where an agent is an independent player or a collusive set).

Definition 6. (Dominated, Undominated, and Σ^1 Strategies in Collusive Auctions.) *Let $\mathcal{A}$ be an agent in a collusive auction $(\mathscr{C}, M)$. We say that a deterministic strategy $\sigma_{\mathcal{A}}$ of $\mathcal{A}$ 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. $\forall \tau_{-\mathcal{A}} \in \Sigma'_{-\mathcal{A}}$, $E[u_{\mathcal{A}}(M(\sigma_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))] \leq E[u_{\mathcal{A}}(M(\sigma'_{\mathcal{A}} \sqcup \tau_{-\mathcal{A}}))]$.

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

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'.

We denote by $\Sigma^1_{\mathcal{A},\mathscr{C}}$ *the set of deterministic strategies of* $\mathcal{A}$ *undominated over* $\Sigma^0_{\mathscr{C}}$ $(=\Sigma^0)$*, and by* $\Sigma^1_{\mathscr{C}}$ *the set of strategy vectors surviving the first round of elimination of dominated strategies, that is,* $\Sigma^1_{\mathscr{C}} = \prod_{C\in\mathbb{C}^{\mathscr{C}}} \Sigma^1_{C,\mathscr{C}}$.

Definition 7. (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)* $RK_i = RK_i^{\mathscr{C}}$
- a collusive set C *if (a)* $C \in \mathbb{C}^{\mathscr{C}}$ *and (b)* $\mathbb{U}^{\mathscr{C}}_C$ *is* C*'s collective utility function.*

Remarks. Fixing the mechanism M,

- $\Sigma^1_{\mathcal{A},\mathscr{C}}$ *is the same for any* $\mathscr{C}$ *compatible with* $\mathcal{A}$ *(and is computable by* $\mathcal{A}$*).* In fact, the set Σ^0 is fully determined from M alone, and which strategies of $\mathcal{A}$ are undominated over Σ^0 solely depends on $\mathcal{A}$'s utility function (rather than, say, on the partition of the other players into collusive sets, and their utility functions). Accordingly, we shall more simply write $\Sigma^1_{\mathcal{A}}$ instead of $\Sigma^1_{\mathcal{A},\mathscr{C}}$, and thus $\Sigma^1_{\mathscr{C}} = \prod_{C\in\mathbb{C}^{\mathscr{C}}} \Sigma^1_C$.
- $\Sigma^1_{\mathscr{C}}$ *is crucially dependent on* $\mathscr{C}$. In fact, although each C can compute Σ^1_C "independent of the overall collusive context", if C is a collusive set in one context $\mathscr{C}$, it may not be a collusive set in another context $\mathscr{C}'$. (In $\mathscr{C}$, for example, C may in particular consist of 10 players and Σ^1_C of just one collective strategy. In $\mathscr{C}'$, however, all the players in C may be independent, and for each $i \in C$, Σ^1_i may consist of two individual strategies. Therefore, the players in C contribute a single strategy subvector to $\Sigma^1_{\mathscr{C}}$, but 1024 strategy subvectors to $\Sigma^1_{\mathscr{C}'}$.)

Definition 8. (Σ^2_I Strategies and Σ^1/Σ^2_I Plays for Minimally Monotone Contexts.) *Let* $(\mathscr{C}, M)$ *be a minimally monotone collusive auction.*

If i *is an independent player in* $(\mathscr{C}, M)$*, then we denote by* $\Sigma^2_{i,\mathscr{C}}$ *the set of all strategies* $\sigma_i \in \Sigma^1_i$ *undominated over* $\Sigma^1_{\mathscr{C}}$*; and by* Σ^2_i *the union of* $\Sigma^2_{i,\mathscr{C}'}$ *for all* $\mathscr{C}'$ *compatible with* i*.*

We say that a strategy vector σ *is a* Σ^1/Σ^2_I play *of* $(\mathscr{C}, M)$ *if*

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

Definition 9. (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 minimally monotone collusive contexts* $\mathscr{C}$*, and all* Σ^1/Σ^2_I *plays* σ *of the auction* $(\mathscr{C}, M)$*,* $\mathbb{P}$ *holds for* $M(\sigma)$*.*

Remarks.

- 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.
- Precisely computing $\Sigma^2_{i,\mathscr{C}}$ requires precise knowledge of $\Sigma^1_{\mathscr{C}}$. Accordingly, if i is unaware of the *actual* collusive context $\mathscr{C}$, he may be unable to compute $\Sigma^2_{i,\mathscr{C}}$ (but he could compute the much larger set Σ^2_i by going through all possible collusive contexts compatible with him). In particular, i may not even think of collusion, and thus behave as if all players were independent. However, no matter what i believes about collusion, the strategies of i that —in i's mind!— survive the second round of elimination of dominated strategies will be a *subset* of Σ^2_i. Thus if a property $\mathbb{P}$ is implemented in Σ^1/Σ^2_I strategies, then it holds no matter what strategies each independent player i believes to be his Σ^2_i strategies.

6 Our Mechanism

Although requiring some modifications, the basic idea behind our mechanism is very simple: each player i, simultaneously with the others, announces an canonical external outcome Ω^i for i, and then we run a "second-price auction." Let us explain. Let R' be the second highest of all revenues of the announced outcomes, and let $\star$ be the "star player", that is the one who has announced the outcome with the highest revenue. Then, we would like to (1) sell the goods according to $\Omega^\star$, so as to generate revenue $R^\star = \text{REV}(\Omega^\star)$, and (2) divide this revenue as follows: the seller gets R' and the star player gets $R^\star - R'$. To implement the first step, let $\alpha^\star$ and $\pi^\star$ respectively be the allocation and the price profile of $\Omega^\star$. Then we ask each player $i \in -\star$, receiving some goods in $\alpha^\star$, whether he is willing to buy the subset of goods $\alpha^\star_i$ for price $\pi^\star_i$. If i agrees, the subsale is final. Else, the star player pays a fine equal to $\pi^\star_i$ and the goods in $\alpha^\star_i$ remain unallocated.

With this way of proceeding, it can be shown that our benchmark is achieved *if each independent player i does not "underbid",* that is, if he announces either his relevant-external-knowledge outcome RK_i or an outcome whose revenue is higher than that of RK_i. But: can we be sure that an independent player i does not underbid? The problem is that, depending on his general knowledge GK_i, underbidding may be a quite rational option. For instance GK_i can be compatible with RK_i and yet assure i that (1) he is the best informed player, that is, the revenue of RK_i is higher than that of the relevant-external-knowledge outcome of any other player; (2) the next best informed player is j; (3) the revenue of RK_j is very close to that of RK_i; and yet (4) player j is badly informed about him: that is RK_j allocates to i a subset of goods S_i that i highly values for a ridiculously low price P_i. In this case, i would be better off if (a) j announced RK_j and (b) j were the star player. In sum, underbidding may be far from being a dominated strategy for i.

Our mechanism thus modifies the above basic procedure as follows. Together with RK_i, each player i also announces his favorite subset of the goods. And if i is declared the star player, then a coin toss of the mechanism determines whether the above basic procedure takes place or i receives, for free, his favorite subset. A precise (if tedious) analysis proves that this modification is an incentive sufficient to ensure that underbidding will not survive the iterated elimination of dominated strategies for any *independent* player i. The same cannot be said about *collusive* players, but then we do not rely on them for efficiency nor for revenue. However, we must ensure that they do not hurt the achievement of our benchmark, something that unfortunately adds significant complexity to our analysis.

(In $\mathcal{M}$'s description below, real actions occur in "numbered steps" and public updates in "bulleted steps.")

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 external outcome for i, $\Omega^i = (\alpha^i, \pi^i)$, and (2) a subset S_i of the goods.

- Set: $R_i = \text{REV}(\Omega^i)$ 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. Publicly flip a fair coin.

- (If Heads:) reset $A_\star = S_\star$ and HALT.

3. (If Tails:) Each player i such that $\alpha^\star_i \neq \emptyset$ simultaneously and publicly announces YES or NO.

- Reset: (1) $P_\star = P_\star + R'$; (2) for each player i who announces NO, $P_\star = P_\star + \pi^\star_i$; and (3) for each player i who announced YES, $A_i = \alpha^\star_i$, $P_i = \pi^\star_i$, and $P_\star = P_\star - \pi^\star_i$.

Let us now state our theorem, proved in the appendix.

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

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

7 Privacy and Computational Efficiency

By a proper interpretation of the revelation principle, our mechanism possesses a normal-form counterpart. (In it, however, a collusive set C should be careful in choosing its collective strategy when the star player is not in C; or it would essentially reveal itself as collusive.) But the revelation principle *does not preserve privacy*, and privacy issues may alter the way games are played in practice, threatening the meaningfulness of fundamental solution concepts such as dominant-strategy truthfulness. A general result of Lepinski, Izmalkov, and Micali [ILM'08] guarantees that every normal-form mechanism has an extensive-form version that perfectly implements it, in particular without any privacy loss. However, their construction —although feasible— is practically inconvenient. By contrast, let us point out that mechanism $\mathcal{M}$, although not designed to achieve perfect privacy, (1) does guarantee a lot of privacy and (2) is computationally very efficient.

Putting it playfully, $\mathcal{M}$ is *"tax-free"*. Consider a second-price auction for a single good in which the winner bids \$10M and wins the item for \$1M, and assume an overreaching and tyrannical tax code. Then, the IRS may demand 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. Such a demand cannot arise in an Ascending English Auction (AEA). In the AEA, the players who drop out reveal their true valuations, but are not "taxable" because they have no utility. As for the winner, he could always declare that his true value for the object was (assuming the same valuations above) exactly \$1M (plus \$1 if he really feels to look more "legitimate"). In an ascending English auction, therefore, there is nothing to "tax".

Leaving taxes aside, there is a legitimate privacy issue here. The winner of an auction may not want to inform its competitors of the utility he has received, and thus of his real financial strength. More generally, no player wants to declare his own true valuation. Yet, in a combinatorial auction —even if the mechanism used is not DST— the players themselves may provide some information or evidence about their own valuations. By contrast, in mechanism $\mathcal{M}$, no player reveals more information about himself than that "deducible from the final outcome itself." If $\mathcal{M}$'s 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: whatever the other players say about him is just "hearsay." If $\mathcal{M}$'s 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.

(Note that further privacy could be gained if $\mathcal{M}$ first asked each player i to announce just the revenue of Ω^i in Stage 1, and then asked only the star player to reveal both $\Omega^\star$ and $S_\star$. However, this alternative way of proceeding would enable the star player to announce $\Omega^\star$ depending on the revenues announced by the other players, and thus an independent player i may have incentives to underbid.)

Finally note that, unlike the VCG, $\mathcal{M}$ is computationally trivial. At most, it is the players who have to work hard to deduce RK_i from GK_i. But this is a different matter. When an exponential-time mechanism is (necessarily!) approximated by a computationally efficient one, crucial distortions of incentives may ensue. But this is not the case here. If a player i *imperfectly,* because computationally bounded, deduces RK_i from GK_i, our $\mathcal{M}$ still achieves our total-performance benchmark —not defined on the *perfect* relevant knowledge of the players, but on the relevant knowledge *actually computable* by the players.

8 Extensions and Conclusions

With additional work, our mechanism can be extended so as to achieve a more demanding benchmark: the maximum social welfare known to the independent players. (I.e., the relevant knowledge of a player now includes the possibility of "giving himself some goods.") Achieving this other benchmark however requires mechanisms that also solicit *separate and special* bids from collusive sets. Such mechanisms were pioneered by Micali and Valiant in a yet unpublished manuscript that has much influenced the present paper [MV'07].

In closing, we believe knowledge-based bechmarks and equilibrium-free solution concepts to be very powerful conceptual tools that will enable us to provide meaningful and attractive solutions to a host of other problems in mechanism design. (In particular, similar techniques can be applied to provision of a public good.)

References

[AM'06] L.M. Ausubel and P. Milgrom. The Lovely but Lonely Vickrey Auction. *Combinatorial Auctions*, MIT Press, pages 17-40, 2006.

[BBHM'05] M.F. Balcan, A. Blum, H. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. *Foundations of Computer Science*, pages 605-614, 2005.

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

[BLP'06] 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.

[C'71] E.H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, Vol.11, No.1, pages 17-33, Sep., 1971.

[CM'88] 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.

[FGHK'02] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. *Symposium on Theory of Computing*, pages 72-81, 2002.

[FPS'00] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. *Symposium on Theory of Computing*, pages 218-226, 2000.

[G'73] T. Groves. Incentives in Teams. *Econometrica*, Vol. 41, No. 4, pages 617-631, Jul., 1973.

[GH'05] A. Goldberg and J. Hartline. Collusion-Resistant Mechanisms for Single-Parameter Agents. *Symposium on Discrete Algorithms*, pages 620-629, 2005.

[GHKKKM'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.

[GHKSW'06] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive Auctions. *Games and Economic Behavior*, Vol. 55, Issue 2, pages 242-269, May, 2006.

[ILM'08] 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".)

[J'92] M. Jackson. Implementation in Undominated Strategies: a Look at Bounded Mechanisms. *Review of Economic Studies*, Vol. 59, No. 201, pages 757-775, Oct., 1992.

[JV'01] K. Jain and V. Vazirani. Applications of Approximation Algorithms to Cooperative Games. *Symposium on Theory of Computing*, pages 364-372, 2001.

[LOS'02] D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. *Journal of the ACM*, Vol. 49, Issue 5, pages 577-602, Sep., 2002.

[LS'05] A. Likhodedov and T. Sandholm. Approximating Revenue-Maximizing Combinatorial Auctions. *National Conference on Artificial Intelligence (AAAI)*, pages 267-274, 2005.

[MS'01] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, Vol.18, No.3, pages 511-533, 2001.

[MV'08] S. Micali and P. Valiant. Resilient Mechanisms for Truly Combinatorial Auctions, submitted to *Symposium on Theory of Computing*, 2009.

[MV'07] S. Micali and P. Valiant, Leveraging Collusion in Combinatorial Auctions, *Unpublished Manuscript*, 2007

[V'61] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, Vol. 16, No. 1, pages 8-37, Mar., 1961.

Appendix

Our mechanism and analysis assume that a player's true valuation maps subsets of the goods to non-negative numbers (but we could handle negative valuations as well).[10]

A Analysis of Our Mechanism

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

Lemma 1. *For all independent players i and all $\sigma_i \in \Sigma_i^1$: if $i \neq \star$ and $\alpha_i^\star \neq \emptyset$ after Stage 1, then in Stage 3 (that is, when $\mathcal{M}$'s coin toss comes up Tails)*

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 exists 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 in Stage 3. Consider the following alternative strategy for 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 3. If $\overline{\mathcal{P}}$, run σ_i and answer whatever σ_i does.[11]
If $\mathcal{P}$, answer YES.

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 in Stage 3. (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 whenever the coin toss of $\mathcal{M}$ is the same.) 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 in Stage 3. (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 execution of $\sigma_i' \sqcup \tau_{-i}$.

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

[10] In traditional auctions, valuations *are* bids, and the seller would immediately dismiss bids associating a subset S of the goods to a negative number (since he has no intention to assign S to a player and also pay him to accept S). The "bidding process" of our extensive-form mechanism however asks each player i to announce in Step 1 a subset S_i of the goods without mentioning any value for S_i. In principle, therefore, i may have a negative valuation for S_i. And leaving things as they stand, i may have (subtle) reasons to announce such an S_i.

[11] 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 $\overline{\mathcal{P}}$.

(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

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

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. *For all minimally monotone collusive sets C and all $\sigma_C \in \Sigma_C^1$: if $\star \notin C$ after Stage 1, then in Stage 3, for all players i in C*

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

Proof. We again restrict ourselves to just prove the first implication, and proceed by contradiction. Assume that there exist a minimally monotone collusive set C, a player $i \in C$, and a strategy vector σ such that $\sigma_C \in \Sigma_C^1$, $\sigma_{-C} \in \Sigma_{-C}^0$, and in σ's execution i answers NO in Stage 3 and the following property holds:

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

Then, denoting by $\overline{\mathcal{P}_{i,C}}$ the negation of $\mathcal{P}_{i,C}$, that is,

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

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 3. 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 $\sigma_C \notin \Sigma_C^1$.

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 $\tau_{-C} \in \Sigma_{-C}^0$ such that the execution of $\sigma_C \sqcup \tau_{-C}$ satisfies $\mathcal{P}_{i,C}$ and i answers NO in Stage 3. (Note that by hypothesis, $\tau_{-C} = \sigma_{-C}$ is one such strategy sub-vector.) For each such τ_{-C}, we have that (1) for all $j \in C \setminus \{i\}$, $(\mathcal{M}_a(\sigma_C \sqcup \tau_{-C})_j, \mathcal{M}_p(\sigma_C \sqcup \tau_{-C})_j) = (\mathcal{M}_a(\sigma_C' \sqcup \tau_{-C})_j, \mathcal{M}_p(\sigma_C' \sqcup \tau_{-C})_j)$, and (2) $\mathcal{M}_a(\sigma_C \sqcup \tau_{-C})_i = \emptyset$ and $\mathcal{M}_p(\sigma_C \sqcup \tau_{-C})_i = 0$, no matter what the coin toss of $\mathcal{M}$ is. Thus, due to C's minimal 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 $u_i(\mathcal{M}(\sigma_C' \sqcup \tau_{-C})) = 0$ when the coin toss of $\mathcal{M}$ comes up Heads, and that $u_i(\mathcal{M}(\sigma_C' \sqcup \tau_{-C})) > 0$ 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 $\forall$ $\sigma_i \in \Sigma_i^2$, $\text{REV}(\Omega^i) \geq \text{REV}(RK_i)$ (that is, i does not "underbid").*

Proof. We proceed by contradiction. Assume that there exists an independent player i and a strategy $\sigma_i \in \Sigma_i^2$ 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 the outcome $\widehat{\Omega}^i = (\widehat{\alpha}^i, \widehat{\pi}^i) = RK_i$, and
the subset of goods $\widehat{S}_i = \arg\max_{S \subseteq G} TV_i(S)$.

Stage 3. 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 minimally 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 $\overline{\sigma}_i$ be a strategy such that $\overline{\sigma}_i \neq \widehat{\sigma}_i$ and $\overline{\sigma}_i$ dominates $\widehat{\sigma}_i$ over Σ^0. Assume that $\overline{\sigma}_i$ announces $\overline{\Omega}^i \neq \widehat{\Omega}^i$ or $\overline{S}_i \neq \widehat{S}_i$, 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, announces YES in Stage 3 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, since $\text{REV}(RK_i) > 0$ by hypothesis, $i = \star$ under the profile $\widehat{\sigma}_i \sqcup \sigma_{-i}$ and $\mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \sigma_{-i}))] > \frac{TV_i(\widehat{S}_i)}{2} = \frac{\max_{S \subseteq G} TV_i(S)}{2}$. While $\mathbb{E}[u_i(\mathcal{M}(\overline{\sigma}_i \sqcup \sigma_{-i}))] \leq \frac{\max_{S \subseteq G} TV_i(S)}{2}$. Therefore such a $\overline{\sigma}_i$ can not dominate $\widehat{\sigma}_i$ over Σ^0.

Accordingly, if $\overline{\sigma}_i$ dominates $\widehat{\sigma}_i$, it must be that $\overline{\sigma}_i$ 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. If the coin toss of $\mathcal{M}$ comes up Heads, then the final outcomes under the profiles $\widehat{\sigma}_i \sqcup \sigma_{-i}$ and $\overline{\sigma}_i \sqcup \sigma_{-i}$ are clearly the same, so are $u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \sigma_{-i}))$ and $u_i(\mathcal{M}(\overline{\sigma}_i \sqcup \sigma_{-i}))$. Let us now consider the case when the coin toss of $\mathcal{M}$ comes up Tails and the two executions run into Stage 3. There, Lemma 1 implies that the only possible difference between $\widehat{\sigma}_i$ and a dominating $\overline{\sigma}_i$ consists of what the two strategies announce when $i \neq \star$, $\alpha_i^\star \neq \emptyset$ and $TV_i(\alpha_i^\star) = \pi_i^\star$: namely, $\widehat{\sigma}_i$ answers YES (by definition) and $\overline{\sigma}_i$ 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 $\overline{\sigma}_i \neq \widehat{\sigma}_i$ can dominate $\widehat{\sigma}_i$ over Σ^0. In sum, $\widehat{\sigma}_i \in \Sigma_i^1$ as we wanted to show. $\square$

Claim 2: For all minimally monotone collusive contexts $\mathscr{C}$ compatible with i, $\widehat{\sigma}_i$ dominates σ_i over $\Sigma_{\mathscr{C}}^1$.
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 $\tau_{-i} \in \Sigma_{\mathbb{C}\setminus\{i\}}^1$, 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 in Stage 3, 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 $\Sigma_{\mathscr{C}}^1$, we analyze the following four exhaustive cases, again after arbitrarily fixing $\tau_{-i} \in \Sigma_{\mathbb{C}\setminus\{i\}}^1$.

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, $\alpha_i^\star = \widehat{\alpha}_i^\star$ and $\pi_i^\star = \widehat{\pi}_i^\star$. There are four sub-cases.

(a) $\alpha_i^\star = \emptyset$. 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 Stage 3 of 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 Stage 3 of 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 Stage 3 of both executions, and we have $\mathbb{E}[u_i(\mathcal{M}(\sigma_i \sqcup \tau_{-i}))] = \frac{TV_i(\alpha_i^\star) - \pi_i^\star}{2} = \frac{TV_i(\widehat{\alpha}_i^\star) - \widehat{\pi}_i^\star}{2} = \mathbb{E}[u_i(\mathcal{M}(\widehat{\sigma}_i \sqcup \tau_{-i}))]$.

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 in Stage 3, 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 in Stage 3, 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 in Stage 3, causing himself to be assigned the subset of goods $\alpha_i^\star$ for price $\pi_i^\star$, and thus can have positive utility. 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.[12] 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 3 of the execution of $\widehat{\sigma}_i \sqcup \tau_{-i}$ (that is, $\mathcal{M}$'s coin toss comes up Tails), we have

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

By definition of RK_i and compatibility, $\forall j \in -i$ such that $\widehat{\alpha}_j^i \neq \emptyset$, $\widehat{\pi}_j^i < TV_j(\widehat{\alpha}_j^i)$. Thus by observation O_4 every such player j answers YES in Stage 3: in our notation $\sum_{j:\widehat{YES}} \widehat{\pi}_j^i = \sum_j \widehat{\pi}_j^i$ and $\sum_{j:\widehat{NO}} \widehat{\pi}_j^i = 0$. Accordingly, we have

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

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

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

by "$\sum_{j:YES}$" (respectively,"$\sum_{j:NO}$") the sum taken over every player j who answers YES (respectively, NO) in Stage 3 of the execution of $\sigma_i \sqcup \tau_{-i}$, we have

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

Since $\sum_{j:NO} \pi^i_j \geq 0$, 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) $\text{REV}(\widehat{\Omega}^i) = \text{REV}(RK_i) > \text{REV}(\Omega^i)$; 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 $\Sigma^1_{\mathscr{C}}$. □

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

Comment. Note that, while ruling out the underbidding (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 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 be better off announcing Ω^i such that $\alpha^i_j = S$, $\pi^i_j > v$, 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.

Finally, let us now (recall and) prove our main theorem.

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

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

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) $\text{REV}(\Omega^*) \geq \text{REV}(RK_*) = \mathbb{MEW}(RK_I)$.

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

Indeed, Inequality (a) holds because player $*$ is independent and, by Lemma 3, he does not underbid; and Inequality (b) holds by Inequality (a) and the fact that $R_\star \geq \text{REV}(\Omega^*)$ 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 1 and 2 guarantee that when $\mathcal{M}$'s coin toss comes up Tails, 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:

$$\mathbb{E}[\text{SW}(\mathcal{M}(\sigma), TV)] = \frac{TV_\star(S_\star)}{2} + \frac{\sum_{i:YES} TV_i(\alpha_i^\star)}{2} \geq \frac{\sum_{i:YES} TV_i(\alpha_i^\star)}{2} \geq \frac{\sum_{i:YES} \pi_i^\star}{2}. \tag{1}$$

At the same time,

$$\mathbb{E}[\text{REV}(\mathcal{M}(\sigma))] = \frac{R' + \sum_{i:NO} \pi_i^\star}{2},$$

which immediately translates into the following two inequalities:

$$\mathbb{E}[\text{REV}(\mathcal{M}(\sigma))] \geq \frac{R'}{2} \tag{2}$$

and

$$\mathbb{E}[\text{REV}(\mathcal{M}(\sigma))] \geq \frac{\sum_{i:NO} \pi_i^\star}{2}. \tag{3}$$

Thus, in the case at hand, Inequality (2) coincides with the second part of our thesis, while the first part follows by combining Inequalities (1) and (3) and then using Inequality (b): namely,

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

Case 2: $\star \neq *$.

In this case, we have $* \in -\star$, and thus $R' \geq \text{REV}(\Omega^*)$ by the definition of R'. Therefore, by Inequality (a), we have

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

Of course

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

Thus summing term by term we have

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

Q.E.D.