Leveraging Collusion In Combinatorial Auctions

Silvio Micali and Paul Valiant

April 11, 2008

Abstract

We investigate dominant-strategy auction mechanisms that, should a sufficiently informed coalition of players be present, exploit it so as to guarantee more efficiency and revenue than is otherwise possible.

(Coming from a cryptographic tradition and prizing extreme settings, we refrain from relying on weaker notions of equilibrium; the availability of Bayesian information; any restrictions on the number of players, goods, and colluders; or any restrictions on the type of valuations and the ability of the colluders.)

1 Introduction

1.1 Collusion Resiliency in Combinatorial Auctions

Combinatorial Auctions. In an auction with n players and m goods, an allocation A is a partition of the goods, $A = A_0, A_1, \ldots, A_n$, where A_i are the goods given to player i, and A_0 are the goods unallocated; a *valuation* is a function associating a non-negative number to each possible subset of the goods. If V is a (sub)profile —i.e., a vector indexed by (a subset of) the players— of valuations, then: V's *social welfare* relative to allocation A, denoted by $SW(V, A)$, is $\sum_i V_i(A_i)$; V's *best allocation*, $BA(V)$, is the allocation maximizing the social welfare of V; and V's *maximum social welfare*, $MSW(V)$, is the social welfare of V relative to its best allocation.

The context of an auction consists of a valuation profile TV, where TV_i —i's *true valuation*— represents i's true value for each possible subset of the goods. An auction is called *combinatorial* if the players' true valuations are totally unrestricted.

A *mechanism* is a (possibly probabilistic) function M mapping a profile of valuations V to an allocation A and a profile P of *prices* (i.e., non-negative numbers) such that A_i is empty and P_i is 0 whenever V_i is the null valuation. (Thus any player can "opt out" —i.e., win and pay nothing— by bidding the null valuation.) We view each mechanism M as a pair of functions: the allocation function, M_a, and the price function, M_p.

An auction with mechanism M is *played* as follows. First, each player i —simultaneously with the others— submits a *bid*, that is, a valuation BID_i. Then M is evaluated on the bid profile BID to return an allocation A and a price profile P. The (economic) *efficiency* of such a play is $\sum_i TV_i(A_i)$, and its revenue is $\sum_i P_i$. The *utility* of player i is $TV_i(A_i) - P_i$.

Mechanism Design and Dominant-Strategy Truthfulness. The designer of an auction mechanism typically does not know the players' true valuations, but he knows that, once a mechanism M is announced, rational players wish to maximize their utilities, which very much depend on their true valuations. Thus, he endeavors to design M so that the allocation and prices it returns on the players' rationally chosen bids satisfy a desired property of their true valuations. The classic desiderata for an auction are (1) maximizing efficiency, and/or (2) maximizing revenue.

A rational play of an auction is an *equilibrium*: a profile of distributions over bids $BID = (BID_1, \ldots, BID_n)$ such that no player i has an incentive to deviate from selecting his bid from BID_i, if he believes that every other player j will stick to selecting his bid from BID_j. Accordingly, the designer endeavors to find a mechanism M such that the resulting auction possesses at least one equilibrium BID such that the desired property holds for the allocations and prices returned by $M(BID)$.

However, as for any game, an auction may have several equilibria, and the desired property may hold only for some of them. Moreover, if a player believes that everyone is playing following some equilibrium BID and some other player believes that everyone is playing following a different equilibrium BID', there may be "no real equilibrium at all."

The meaningfulness of mechanism design is thus maximized when M is *dominant-strategy truthful* (DST), that is, such that bidding his true valuation is any player's best choice, *no matter the others' bids may be.* In this case, in fact, there is little or no uncertainty on how rational players will bid.

- *As cryptographers, we focus only on solutions guaranteed to work under extreme uncertainty, and in the most adversarial circumstances. Accordingly, we shall only consider DST mechanisms.*

In combinatorial auctions, the famous VCG mechanism [5, 14, 23] is indeed DST and achieves maximum efficiency, that is, $MSW(TV)$. In single-good auctions, this famous mechanism is much simpler and coincides with the *second-price* mechanism, 2P, that allocates the good to the highest bidder, setting his price to be equal to the second-highest bid, and everyone else's price to be 0.

When the players' true valuations for the single good are known to be drawn from independent (although essentially arbitrary) distributions, Myerson [20] shows how to construct DST mechanisms that also generate optimal revenue from this information.

The Problem of Collusion in Combinatorial Auctions. The classical game-theoretic notion of an equilibrium relates to *single-player deviations.* Accordingly, even in a game with a dominant-strategy equilibrium, it is possible for two or more players to deviate so as to increase their (collective) utility. Auctions are no exception.

Single-good auctions, however, are so simple that the 2P mechanism is automatically quite "resilient to collusion." By this we mean that, as long as all players are rational, no matter how the collusive players may bid, 2P continues to return maximum efficiency and revenue at least as high as the second-highest valuation of the independent players, if any.

The situation is dramatically different for combinatorial auctions, where the vulnerability of the VCG mechanism to collusion is only too well known [1, 22, 6]. Even when there are just two goods and two collusive players (with moderate knowledge about the true valuations of the others), VCG has no useful efficiency or revenue guarantees.

The problem of collusion has inspired many papers, summarized in Section 2. All such papers, however, work in weaker models. Specifically,

- They work with weaker notions of equilibria; or
- They assume some restriction on the collusive players (i.e., that they cannot exchange money with each other, or that they cannot enter binding agreements with each other, etc.); or
- They assume that the auction is of a restricted type (i.e., that there is a single good, or multiple copies of the same good, or that the players' true valuations are sub-additives, etc.)

In sum, they offer no DST protection against a knowledgeable and coordinated coalition in a truly combinatorial auction. The only exception (relevant to our analysis as it provides the benchmark against which we need to measure the results of this paper) is a recent and yet unpublished result of ours.

The MV Result. We have put forward a DST auction mechanism, $\mathbb{MV}$, whose expected revenue, for any bid profile BID in an auction with n players and m goods, is at least $\frac{MSW(BID_{-*})}{c_{n,m}}$, where

"$*$" (the "star" player) denotes the player with the highest valuation for the goods[1], and

$c_{n,m}$ is the only solution greater than 2 to the equation $e^{x-2} = x \min\{n, m\}$.

For instance, when $n, m < 299$ (perhaps a reasonable case), $\mathbb{MV}$'s revenue is always $\geq 10\%$ of $MSW(BID_{-*})$. Furthermore, $\mathbb{MV}$ is naturally collusion-resilient: no matter how the collusive players may bid, if the set of independent players I is non-empty, $\mathbb{MV}$'s revenue is at least $\frac{MSW(BID_{I-*})}{c_{n,m}}$.

Although $c_{n,m}$ increases very slowly and is actually quite reasonable for small values of n and m, it is only human to demand higher revenue. Unfortunately, this time we cannot get what we demand or what we desire: Indeed, *$\mathbb{MV}$'s revenue performance is (essentially) DST optimal.*[2]

Options. In light of $\mathbb{MV}$'s optimality, if we are dissatisfied with its revenue performance we could:

(1) blame Nature; (2) trick ourselves by restating $\mathbb{MV}$'s benchmark so that it "looks better;" (3) work with an equilibrium notion weaker than DST and pray that the auction ends in "an equilibrium good for us;" (4) pray that we are given extra knowledge about the players' true valuations; or (5) lose interest and shop around for some other problem that we can solve with better-looking bounds.

There is, however, another alternative:

- *When you cannot defeat powerful foes: ally yourselves with them.*

This is indeed the alternative we investigate in this paper.

[1] I.e., $MSW(TV_*) \geq MSW(TV_i)$ for all i

[2] I.e., no DST mechanism can guarantee asymptotically greater expected revenue when n and m grow to infinity. And it is optimal within a constant factor "along the way." That is, for any auction with n player and m goods, and any DST mechanism $\mathcal{M}$, there is a profile of bids BID for which $\mathcal{M}$'s expected revenue is $\leq 5\frac{MSW_{-\star}(BID)}{c_{n,m}}$.

1.2 Our Notion (From Collusion Resiliency to Collusion Leveraging)

The High-Level Approach. Constructing DST mechanisms guaranteeing good efficiency and revenue in the presence of collusion would be easier if some useful knowledge about the players' true valuations were available to the auction designer. Less conveniently, we instead assume that such knowledge is only available to the coalition C, if present. Accordingly, rather than "fighting" C we aim to reward it in a way that is mutually beneficial. That is, our auction mechanisms are designed so as to incentivize both C and the independent players (although in quite different ways) so that, in a rational play, we succeed in leveraging C's knowledge to guarantee a reasonable fraction of the following benchmarks:

- $MSW(TV)$ for efficiency; and
- $MSW(TV_{-C})$ for revenue.

(Guaranteeing efficiency might be simpler if the goods could be reallocated, after the auction is over, by the players who win them. But, as for "spectrum auctions," we assume that reallocation is not an option.)

Reasonableness of the Approach. At the highest level, our approach is reasonable in that, as in mainstream mechanism design, we envisage that "all knowledge about the players resides with the players themselves." Focusing on coalitions having some privileged knowledge about the independent players is reasonable too, in the sense that a coalition has fewer reasons to form in absence of such knowledge. Finally, also our benchmarks are quite reasonable. Indeed, our efficiency benchmark is the highest possible. And our revenue benchmark actually coincides with the maximum social welfare of the independent players, and thus with the maximum revenue any DST can produce when "all collusive players have the courtesy of staying away from the auction" —a not bad, although improbable, case!

Coalition and Knowledge Models. To detail our approach fully we must clarify a few things, in particular, the way C operates and the kind of knowledge it has. Informally, we work with the following models.

For C's operational aspects, we assume that its members can make side-payments to each other and coordinate their actions. Indeed, we *incentivize the collusive players to perfectly coordinate themselves.*

For C's knowledge about the independent valuations, we consider two models:

- *Approximate Knowledge.* In this model, we assume that, for some $k \geq 1$, the coalition knows an acceptable way to allocate all goods to the independent players in which they pay at least a fraction $\frac{1}{k}$ of their maximum social welfare.
- *Bayesian Knowledge.* In this model, we assume that the coalition knows the distribution(s) from which the independent valuations are actually drawn.

(Note that one way for the collusive players to have "$\frac{1}{k}$-knowledge" is to know the valuations of the independent players within a factor of k. That is, to know a valuation V_i for each player i such that $V_i \leq TV_i(S) \leq kV_i$. Note too, however, that the latter requirement is much more stringent than $\frac{1}{k}$-knowledge.)

Rationality. Our solution concept is very strong. For the independent players, we require that our collusion-leveraging mechanism be DST, for the collusive players we require that they are *minimally rational.* Let us explain.

1. Our mechanism guarantees that, via a standard strategy s "consisting of playing their own knowledge", the collusive players secure to themselves a reasonable utility —but *does not* exclude that they might have better strategies; further,
2. Our mechanism also guarantees that, no matter what strategy S at least as good as s the collusive players may choose, our efficiency and revenue goals will be achieved.

In a sense, ours is an *"implementation in non-dumb strategies."*

(Note that —if successful— our mechanism seems to contradict $\mathbb{MV}$'s optimality. We address this apparent contradiction in the next sub-section.)

In Sum. In a combinatorial auction whose coalition C enjoys some special and proprietary knowledge about the players' true valuations, *"a mechanism L leveraging C is a DST mechanism approximating the performance of a DST mechanism M optimally designed given the same knowledge as C."*

1.3 Our Results

Our "Positive" Result. We prove that, making use of probabilism, it is possible to meaningfully leverage a coalition with approximate knowledge. A bit more precisely, we show that for every $k \geq 1$ there exists a *probabilistic* DST mechanism $\mathcal{M}_k$ that in any auction with a coalition C with $\frac{1}{k}$-knowledge:

- $\mathcal{M}_k$'s expected efficiency is essentially a fraction $\frac{1}{2k}$ of the maximum social welfare of all players; and
- $\mathcal{M}_k$'s expected revenue is essentially a fraction $\frac{1}{2k}$ of the maximum social welfare of the independent players.

(Note that each $\mathcal{M}_k$ is constructively specified: essentially, by specifying the knowledge factor k to a single mechanism $\mathcal{M}$.) For instance, if the collusive players can guess the independent valuations within a factor of 2, then $\mathcal{M}_2$'s efficiency is $\frac{1}{4}$ of $MSW(TV)$, and its revenue $\frac{1}{4}$ of $MSW(TV_{-C})$. This performance is substantially better than what was previously achievable in the presence of the same coalition C: in fact, $\mathbb{MV}$ only guaranteed revenue (and thus efficiency) that is essentially a fraction $\frac{1}{\log \min\{n,m\}}$ of $MSW_{-\star}(TV_{-C})$.

How can $\mathcal{M}_k$ outperform $\mathbb{MV}$ if the latter mechanism was proved optimal? The answer, of course, is that in the first case we are guaranteed that the bids have been produced by players including a coalition with knowledge factor k, while this needs not to be true in the second case. Of course too, because the existence of a sufficiently knowledgeable coalition among a given set of players is hardly a public fact, an auctioneer may be reluctant to announce his using $\mathcal{M}_k$ in his auction if, at least personally, he is not sure that a coalition with knowledge factor k is present among his players. Indeed, no guarantees have been given so far for the performance of a collusion-leveraging mechanism when no collusion happens to be present. Accordingly, rather than restricting our mechanisms to contexts in which such a coalition is guaranteed to be present, we find it useful to design $\mathcal{M}_k$ so as to budget for the case in which there is no such coalition. Specifically, we give the auctioneer the ability to specify to $\mathcal{M}_k$, as a special input, a DST "sub-mechanism" $\mathbb{S}$ (for "safety"). Such an $\mathbb{S}$ represents the mechanism the auctioneer would like to use if he were guaranteed that the envisaged coalition did not exist. For instance, if the auctioneer were interested only in efficiency and believed that no coalition could exist unless one with knowledge factor k were present, then he might choose $\mathbb{S} = VCG$. As for another example, if the auctioneer were interested in revenue and believed that, in absence of a coalition with knowledge factor k, some other coalition would for sure be present (possibly with some other type of knowledge), then he might choose $\mathbb{S} = \mathbb{MV}$.

Accordingly, to better evaluate the quality of a collusion-leveraging mechanism, we find it important to also quantify its performance when the envisaged coalition fails to exist. In our case:

- In absence of collusion, for any choice of $\mathbb{S}$ and knowledge factor k, $\mathcal{M}_k$'s efficiency and revenue are respectively guaranteed to be within a factor of 2 of $\mathbb{S}$'s efficiency and revenue.

Therefore, $\mathcal{M}$ does not contradict $\mathbb{MV}$'s performance. For instance, for all bid profiles BID, if there is no coalition, then $\mathcal{M}_2$ with $\mathbb{S} = \mathbb{MV}$ returns half of the revenue of $\mathbb{MV}$. (Perhaps this is a small "insurance premium" if the auctioneer suspects that a coalition with knowledge factor 2 is highly likely.)

Our "Negative" Result. Notice that a collusion-leveraging mechanism is *meaningful* only if its performance is guaranteed to be non-trivially better than that of an ordinary DST mechanism "knowing nothing" (neither about the players' valuations, nor about whether a coalition is present and what knowledge it has). Thus a natural question arises: *what type of collusive knowledge can be meaningfully leveraged?*

On the positive side, having proved that approximate knowledge can be meaningfully leveraged, we conjecture that this is also the case for a large class of knowledge types. But we also conjecture that there is another large class of knowledge types for which no meaningful leveraging is possible. In support of this conjecture, we prove that, in general, *Bayesian knowledge* cannot be meaningfully leveraged, neither for efficiency nor for revenue.

For the case of *revenue*, this impossibility is a corollary of the cited MV result. There, in fact, we essentially showed the existence (for any auction with n players and m goods) of a distribution of bid profiles for which no DST mechanism could guarantee revenue higher than $\mathbb{MV}$. Thus, because a collusion-leveraging mechanism is at best a DST mechanism that "knows what the coalition knows," if all true valuations were chosen according to that distribution, and if the coalition only had that Bayesian knowledge, then no collusion leveraging mechanism could outperform $\mathbb{MV}$ in revenue, despite the fact that $\mathbb{MV}$ indeed is a DST mechanism that "knows nothing."

For *efficiency,* things are a bit more complex: in Section 5, we prove that Bayesian knowledge cannot be "too meaningfully" leveraged. Informally, by this mean that we prove the existence of: (1) a DST mechanism M that "knows nothing"; (2) a simple Bayesian setting for auctions with n players and m goods, and (3) a "slow-growing" function $f(n, m)$ such that

> *No mechanism leveraging a coalition with this Bayesian information can guarantee expected revenue greater than $f(n, m)$ times that of M.*

In our prof:

1. We again choose $M = \mathbb{MV}$ as our "point of comparison."
2. We construct a very simple Bayesian setting; namely, a subprofile of distributions D_{-C} for the true valuations of the independent players such that: (a) each TV_i is *independently* selected according to D_i, (b) each TV_i is single-minded (i.e., assigns the same positive value only to the supersets of a single set S_i of goods), and (c) each D_i has a very small support: indeed there are $\leq \min\{n, m\}$ possible values that D_i assigns to all superset of S_i. Finally,
3. We set $f(n, m) = O(\log \log \min\{n, m\})$.

As computer scientists, we indeed consider the function $\log \log \min\{n, m\}$ negligible, and thus interpret this result as *de facto* saying that no mechanism leveraging Bayesian knowledge can guarantee performance better than the non-leveraging mechanism $\mathbb{MV}$. We expect that economists may have a different view of what is negligible. In any case, we find it theoretically important to prove that the possibilities of leveraging Bayesian knowledge are more limited.

Our Structural Result. We also prove that *meaningful collusion-leveraging is impossible for any deterministic mechanism.* This structural result is very strong. In essence, we prove that even in the presence of an *omniscient coalition*, that knows exactly the true valuations of all players and thus has precisely 1-knowledge, there is no way for a DST mechanism to incentivize such a coalition so as to guarantee outperforming the efficiency and or the revenue of the "knowing-nothing" (about valuations and coalitions) mechanism $\mathbb{MV}$.

Understanding the power of probabilism is fundamental in computation and other settings. Game theory is one of the few settings in which this power can be been actually proven. Although impossibility results can be notoriously hard, in some cases they can be quite simple, although still insightful. (E.g., one can prove by "eye-inspection" that the matching-penny game has no deterministic Nash equilibrium.) This does not appear to be the case for the impossibility of deterministic collusion leveraging. Indeed, almost oxymoronly, we must apply a *probabilistic* method to prove the existence of "bad" auction contexts for every *determinist* collusion-leveraging mechanism.

1.4 Road Map

After reviewing some prior work, we more formally define collusion leveraging for the case of coalitions with approximate knowledge. (Extending the definition to Bayesian knowledge case is not hard, and will in any case be done in the final paper.) Prizing simplicity we then present our results in order of increasing complexity: our "positive" result, then our structural one, and finally our "negative" one. Only the first one fits in the required 10-page body of the paper, the other two are in our Appendix.

2 Prior Work on Collusion (and Approximation)

Collusion in auctions is a real and serious concern. Famously, it has occurred in the case of US Treasury [10, 13] and FCC spectrum auctions [7]. This concern has generated both practical and theoretical responses. Practical responses include the use of monitoring systems in auction rooms, and different bid formats (such as those limiting bids to 3 significant digits). Theoretical responses include *group strategyproof* mechanisms, such as those of [16, 19, 9]. Essentially, such mechanisms discourage collusion in that any gain for a collusive player is accompanied by a loss for another collusive player. Such a response, however, is solely applicable to settings in which collusive players cannot make side payments to one another. This restriction is bypassed by the c-truthful mechanisms of [11]. Essentially, such mechanisms guarantee that fewer than c collusive players cannot "collectively gain more than they could by bidding individually." However, such robustness applies only to a very restricted class of auctions: essentially, when there is a single item for sale, although in unlimited supply. In sum, "collusion is under control" for auctions of a single item. Finally,although outside the scope of our paper, it worth recalling the classical notion of the *core.*

We stress that the beautiful and huge literature about approximating efficiency and revenue in absence of collusion is *orthogonal* to our work. We refer to the MV result only because it is DST and approximates revenue (relative to the independent valuations) in a collusive setting, without any Bayesian information, and in (truly general) *combinatorial auctions*— thus providing a benchmark for collusion-leveraging to beat, in order to be meaningful. The MV result, in particular, should not be confused with other ones dealing with simpler settings. For instance, the results of [18] relate to auctions with (1) restricted valuations —specifically, sub-additive— or (2) goods with unlimited supply —which simplifies things by taking away "competition for the goods," since it is now possible to give good g to some player and *also* to another player— and (3) optimizes small-parameter cases —e.g., auctions with 2 or 3 players and goods— *with the help assuming Bayesian information* about the players true valuations.

3 The Notion of Collusion Leveraging for Approximate Knowledge

Definition 1. *A* $\frac{1}{k}$-approximation *of a subprofile of valuations* V_X *is an outcome* (A', P') *such that (a)* $A'_i = \emptyset$ $\forall i \in -C$*, (b)* $P'_i \le TV_i(A'_i)$ $\forall i \in C$*, and (c)* $\sum_i P'_i \ge \frac{1}{k} MSW(TV)$*.*

Definition 2. *Let* $k \ge 1$*. A* collusive action context $\mathtt{C}$ *with* $\frac{1}{k}$*-approximate knowledge (*$\frac{1}{k}$*-knowledge for short) is a traditional auction context augmented with a subset* C *of the players having as common knowledge a string* K *including four distinguished pieces of information: (1) the subset* C *itself; (2) the subprofile of true valuations* TV_C*; and (3) an outcome* (A', P') *that* $\frac{1}{k}$*-approximates* TV_{-C}*.*

We refer to C as the *coalition;* to a player in C as a *collusive;* to a player in $-C$ as *independent;* to $\frac{1}{k}$ as the *knowledge factor;* to K as the *collusive knowledge;* to K's triple of distinguished pieces of information as the *minimal collusive knowledge* or the *minimal type* of the coalition; and to TV_i, $i \in -C$, as i's type.

Together with an auction mechanism $\mathcal{M}$, $\mathtt{C}$ defines a *collusive* auction $\mathcal{G} = (\mathtt{C}, \mathcal{M})$, which is played as usual, except that bids are now arbitrary strings (rather than valuations) and each player in C chooses his bid based on K. From the resulting allocation A and price profile P, the utility of an independent player i is defined as usual, and that of the coalition C as the sum of the (usual) utilities of its members; that is,

$$u_i = TV_i(A_i) - P_i \quad \text{and} \quad u_C = \sum_{j \in C} TV_j(A_j) - P_j.$$

Definition 3. *Let* M *be a DST mechanism;* $\mathtt{C}$ *a collusive auction context with true-valuation profile* TV *and coalition* C*; and* $\mathfrak{B}_C$ *a bid subprofile. A bid profile* BID *is* $\mathfrak{B}_C$-rational *if*

- $BID_i = TV_i$ *for all* $i \in -C$*; and*
- $U(TV_C, M(BID)) \ge U(TV_C, M(\mathfrak{B}_C \sqcup BID_{-C}))$*.*

Definition 4. *Let M and $\mathbb{S}$ be DST auction mechanisms; and x, y and z real numbers in $[0,1]$. We say that M is* collusion-leveraging with safety $\mathbb{S}$, knowledge factor $\frac{1}{k}$, *and* quality factors (x, y, z) *if the following properties hold:*

1. (Safety) *For any traditional context with true-valuation profile TV:*
   2.1 (Efficiency Safety) $E[SW(TV, M_a(TV))] \geq xE[SW(TV, \mathbb{S}_a(TV))]$*; and*
   2.2 (Revenue Safety) $E\left[\sum_{i\in N} M_p(TV)_i\right] \geq xE\left[\sum_{i\in N} \mathbb{S}_p(TV)_i\right]$.
2. (Leveraging) *For any collusive context with true-valuation profile TV, knowledge factor k, coalition C, and minimal collusive knowledge K', if $\mathfrak{B}_i = K'$ $\forall i \in C$, then for any $\mathfrak{B}_C$-rational bid profile BID:*
   2.1 (Efficiency Leveraging) $E[SW(TV, M_a(BID_C \sqcup TV_{-C})] \geq y\frac{MSW(TV)}{k}$*; and*
   2.2 (Revenue Leveraging) $E\left[\sum_{i\in N} M_p(BID_C \sqcup TV_{-C})_i\right] \geq z\frac{MSW(TV_{-C})}{k}$.

We respectively refer to x, y, and z as M's *safety factor, efficiency factor,* and *revenue factor;* and to $\mathbb{S}$ as M's *safety mechanism.* (Note that $\mathfrak{B}_C$ needs not guarantee a reasonable utility to C, although this will be case for our mechanism $\mathcal{M}$.)

4 Our Positive Result

4.1 The Mechanism $\mathcal{M}$

Our mechanism $\mathcal{M}$ essentially is a uniform way to construct, given a DST mechanism $\mathbb{S}$, a constant $\epsilon \in [0,1]$, and a knowledge factor $k > 1$, an auction mechanism $\mathcal{M}_{\mathbb{S},\epsilon,k}$ whose bids are envisaged to be of two different types.

$\mathcal{M}$'s Bids. No matter what $\mathbb{S}$, ϵ, and k may be, in an auction context with player set N and good set G, a bid consists of either

1. A *traditional bid*, that is, a valuation of G; or
2. A *collusive bid*, that is, a quadrule (X, Y, A', P') such that
   - X is a subset of the players with cardinality greater than 1;
   - Y is subprofile of valuations indexed by the players in X: $Y = Y_X$,
   - (A', P') is an outcome; or
3. An *ill-formed bid*, that is a string other than (one encoding) a traditional bid or a collusive bid.

We call a player "independent-looking" if he submits a traditional bid, and "collusive-looking" otherwise.

$\mathcal{M}_{\mathbb{S},\epsilon,k}$'s Code. On input a bid profile BID:

1. If all players are independent-looking, flip a coin whose probability of Heads is $\frac{1}{2} + \epsilon$.
   a. If Heads, ABORT (i.e., return the empty allocation and a price profile such that i's price is 0 if i is indepedent-looking, and > 0 —i.e., an arbitrary fine— otherwise.)
   b. If Tails, run $\mathbb{S}(BID)$ to return its allocation and price profile, and HALT.
2. If there exists an ill-formed bid, ABORT.
3. If there are two different collusive bids, ABORT. Else, let (X, Y, A', P') the value of all collusive bids.
4. If i is an independent-looking player in X, and every player in $X - \{i\}$ is collusive-looking, then flip a coin whose probability of Heads is $\frac{1}{2} + \epsilon$.
   a. If Tails, ABORT.
   b. If Heads, return the (best) allocation $BA(BID_i)$ and the price profile $(0, \ldots, 0)$, and HALT.
5. If X does not coincide with the set of all collusive-looking players, ABORT.

6. Compute the allocation A and the price profile P as follows: if $i \in -C$, $A'_i \neq \emptyset$, and $P'_i \leq BID_i(A'_i)$; then set $A_i = A'_i$ and $P_i = P'_i$; else set $A_i = \emptyset$ and $P_i = 0$.
   a. If Tails, return A and P and HALT.
   b. If Heads: if $\sum_i P_i \geq \frac{1}{k} MSW(BID_{-X})$ then return the allocation $BA(Y)$ and the price profile $(0, \ldots, 0)$, and HALT; otherwise, ABORT charging each member of X a fine of $MSW(BID_{-X})$.

4.2 Proof that $\mathcal{M}$ is Collusion Leveraging

Theorem 1. *For each DST mechanism $\mathbb{S}$, $\epsilon > 0$, and $k > 1$, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ is a collusion-leveraging mechanism with safety $\mathbb{S}$, knowledge factor $\frac{1}{k}$, and quality factors $(\frac{1}{2} - \epsilon, \frac{1}{2} - \epsilon, \frac{1}{2} - \epsilon)$.*

Proof. We begin by showing that $\mathcal{M}_{\mathbb{S},\epsilon,k}$ is DST. That is, no matter what the bid subprofile BID_{-i} of other players may be, each independent player i is better off bidding his own true valuation TV_i (a traditional bid). We distinguish several cases.

First, assume that BID_{-i} does not include any collusive bids. Here there are three sub-cases to consider: (a) i submits an ill-formed bid, (b) i submits a collusive bid, (b) i submits a traditional bid. In subcase (a), $\mathcal{M}_{\mathbb{S},\epsilon,k}$ aborts in Step 2, causing i's utility to be 0 (and thus no better than the utility he would have gotten by bidding TV_i). Consider now subcase (b), and let (X, Y, A', P') be i's bid. Because this bid is well-formed, X necessarily is a subset of 2 or more players. But, then, since we are now assuming that all players other than i bid traditionally, X does not coincide with the set of collusive-looking players. Accordingly, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ aborts in Step 5 fining i, so that i receives a negative utility. In subcase (c), i's utility is ≥ 0. In this subcase, in fact, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ either aborts in Sub-Step 1a, in which case i's utility is 0, or halts in Sub-Step 1b, in which case i receives his utility "according to $\mathbb{S}$" and thus a non-negative utility because $\mathbb{S}$ is by hypothesis DST. Thus, no matter what BID_{-i} may be, bidding a valuation of G provides i with an utility greater or equal to the utility he receives with any non-traditional bid. Moreover, by bidding TV_i i guarantees himself a utility greater or equal to that of that of any valuation of G: in Step 1a because i' utility is 0 in any case, and in Step 1b because $\mathbb{S}$ is DST. Thus TV_i is i's "best bid" whenever BID_{-i} does not include any collusive bids.

Second, assume that BID_{-i} includes an ill-formed bid, or two different collusive bids. In this case, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ aborts either in Step 2 or in Step 3. In which case, TV_i is as good as any other bid, since i's utility is 0 anyway.

Third, assume that BID_{-i} includes one or more well-formed collusive bids, all coinciding with a quadruple (X, Y, A', P') such that $i \in X$. Here we distinguish two subcases: (e) $X - \{i\}$ includes an independent player; and (d) $X - \{i\}$ solely consists of collusive-looking players. In subcase (e), $\mathcal{M}_{\mathbb{S},\epsilon,k}$ aborts in Step 5 no matter what i may bid, but i's utility is negative if he bids collusively, and 0 if he bids traditionally. Thus bidding his own valuation at least guarantees i utility 0. In subcase (d), any possible bid of i is either (d.1) a collusive bid different from (X, Y, A', P'); (d.2) the collusive bid (X, Y, A', P'); and (d.3) a traditional bid. In case D.1, i's utility is negative, because $\mathcal{M}_{\mathbb{S},\epsilon,k}$ aborts in Step 3 imposing a fine on i. In case D.2, all collusive bids are equal and well-formed, so that $\mathcal{M}_{\mathbb{S},\epsilon,k}$ halts in Step 6. Accordingly, because i's utility is 0 when the mechanism halts in Step 6a; because the mechanism halts in Step 6b with probability $(\frac{1}{2} - \epsilon)$; and because the maximum utility that i can receive in Step 6b when his bid is (X, Y, A', P') consists of receiving for free the subset, S', assigned to him by the VCG mechanism run on the bid subprofile Y; i's expected utility in case d.2 is $(\frac{1}{2} + \epsilon)$ times i's value for S'. In case d.3, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ halts in Step 4 so that, denoting by S'' the subset of goods to which BID_i associates the maximum value, i's expected utility is $(\frac{1}{2} + \epsilon)$ times i's true value for S''. However, if i chooses $BID_i = TV_i$, then in this case his expected utility is $(\frac{1}{2} + \epsilon)$ times his true value for his most valued set. Thus bidding truthfully continues to be i's best course of action also in subcase D.

Finally, assume that BID_{-i} includes one or more collusive bids, all coinciding with a quadruple (X, Y, A', P') such that $i \notin X$. Here it is immediately seen that i's utility is never positive, while bidding TV_i guarantees i utility 0.

Having considered all possible cases, we have proved that $\mathcal{M}_{\mathbb{S},\epsilon,k}$ indeed is DST. This property immediately implies that $\mathcal{M}_{\mathbb{S},\epsilon,k}$ has safety $\mathbb{S}$ with safe quality $\frac{1}{2} - \epsilon$ is trivial. Indeed, when there is no coalition, all players are independent, and thus the rational thing for them to do is to bid their true valuations for the goods G. But when they do so, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ halts in Step 1, aborting $\frac{1}{2}+\epsilon$ of the time, but running $\mathbb{S}$ the rest of the time. Thus when no coalition exists, $\mathcal{M}_{\mathbb{S},\epsilon,k}$'s expected efficiency (respectively, revenue) is a fraction $\frac{1}{2}-\epsilon$ of that of $\mathbb{S}$, as desired.

Let us now analyze, via a sequence of simpler claims, the collusion leveraging properties of $\mathcal{M}_{\mathbb{S},\epsilon,k}$. Let C be a $\frac{1}{k}$-knowledge collusive context with minimal collusive knowledge $K' = (C, TV_C, (X, Y, A', P'))$.

Claim 1: In any execution of $\mathcal{M}_{\mathbb{S},\epsilon,k}(BID^*)$ where $BID^*_{-C} = TV_{-C}$ and $BID^*_j = (C, TV_C, A', P')$:
$\sum_i P_i \geq \frac{1}{k} MSW(TV_C)$.

This follows trivially from the hypothesis that (A', P') is a $\frac{1}{k}$-approximation of TV_{-C}, and the fact that, under such BID^*, our mechanism does not halt in any of the first 5 steps.

Claim 2: $U\,(TV_C, \mathcal{M}_{\mathbb{S},\epsilon,k}(BID^*)) \geq (\frac{1}{2}+\epsilon) MSW(TV_C)$.

First notice that, because by Claim 1 the branching condition of Step 6b is satisfied, in any execution of Step 6b our mechanism, rather than aborting, returns the best allocation for the valuation subprofile TV_C and the identically 0 price profile; thus causing C's utility to be $MSW(TV_C)$. Now notice that, again because it does not halt in the first 5 steps under such the envisaged bid profile, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ executes Step 6b with probability $\geq \frac{1}{2}+\epsilon$. Therefore C's expected utility is as claimed. This has the following immediate consequence.

Claim 3: For any rational bid profile BID', $U\,(TV_C, \mathcal{M}_{\mathbb{S},\epsilon,k}(BID^*)) \geq (\frac{1}{2}+\epsilon) MSW(TV_C)$.

Indeed in any rational play the independent players truthfully bid their own valuations (since $\mathcal{M}_{\mathbb{S},\epsilon,k}$ is DST), and the collusive players place a bid subprofile giving them utility $\geq (\frac{1}{2}+\epsilon) MSW(TV_C)$, since they can trivially guarantee themselves such utility by bidding their own minimal type (C, TV, A', P'). Let us now prove that the collusive component of a rational bid must be a a particular form. Namely,

Claim 4: If BID' is a rational bid profile, then there exists a subset X of C and a collusive bid (X, Y, A', P') such that: $BID_j = (X, Y, A', P')$ for all $j \in X$, and BID'_j is a traditional bid for all $j \in C - X$.[3]

Indeed, first note that $\mathcal{M}_{\mathbb{S},\epsilon,k}$ never returns negative prices (i.e., never gives any money). Accordingly, C's utility can only proceed from the allocation returned by $\mathcal{M}_{\mathbb{S},\epsilon,k}$. This implies that a rational BID'_C cannot solely consists of traditional bids. Indeed, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ returns a non-empty allocation only in Step 1b, and only with probability $\leq \frac{1}{2}-\epsilon$. Thus, even if the returned allocation were the best one for TV_C, C's utility would be strictly less than $(\frac{1}{2}+\epsilon)\frac{MSW(TV_C)}{k}$. At the same time, BID'_C cannot include two different collusive bids, nor a collusive bid "claiming an independent player to be part of the coalition." Else, since rational independent players bid truthfully, $\mathcal{M}_{\mathbb{S},\epsilon,k}$ would halt in a step between 2 and 5, where the coalition's utility is ≤ 0. Let us now proceed to state and prove the collusion-leveraging properties of our mechanism. Namely,

Claim 5: For any rational bid profile BID':

1. *The expected revenue of $\mathcal{M}_{\mathbb{S},\epsilon,k}(BID') \geq (\frac{1}{2}-\epsilon)\frac{MSW(TV_{-C})}{k}$ and*
2. *The expected efficiency of $\mathcal{M}_{\mathbb{S},\epsilon,k}(BID') \geq (\frac{1}{2}+\epsilon) MSW(TV_{-C})$.*

Let BID' be a rational bid profile, and let X and (X, Y, A', P') be as in Claim 4 relative to BID'. Then, under BID', $\mathcal{M}_{\mathbb{S},\epsilon,k}$ halts only in either Step 6a or Step 6b. But C's expected utility from Step 6a alone cannot exceed $(\frac{1}{2}+\epsilon) MSW(TV_C)$, as required from a rational bid, because our mechanism executes Step 6a only with probability $\leq \frac{1}{2}-\epsilon$, and rewards C —or any player for that matter— only via the allocation it returns. Thus the rationality of BID'_C implies that C's expected utility from Step 6b is positive. In turn, this implies that Step 6b's branching condition is satisfied, that is, for the values P_i and the allocation A

[3] Indeed the coalition may, undetectably, "declare only part of its members, hiding the others among the independent players, in an attempt to derive utility in two ways."

computed at the beginning of Step 6

$$\sum_{i\in -X} P_i \geq \frac{MSW(BID_{-X})}{k}$$

But then this inequality implies that, in any execution in which $\mathcal{M}_{\mathbb{S},\epsilon,k}(BID')$ halts in Step 6a, the revenue generated is $\geq \frac{MSW(BID_{-X})}{k}$. Therefore, because $\mathcal{M}_{\mathbb{S},\epsilon,k}(BID')$ cannot halt in any of the first 5 steps and thus executes Step 6a with probability $\geq (\frac{1}{2}-\epsilon)$, its total expected revenue is as claimed.

Notice now that the total utility of $\mathcal{M}_{\mathbb{S},\epsilon,k}(BID')$ is $\geq (\frac{1}{2}-\epsilon)MSW(TV_{-C})$. (In fact, by Claim 3, the utility of C alone is so large, while that of each $i \in -C$ is ≥ 0, because $\mathcal{M}_{\mathbb{S},\epsilon,k}$ is DST). Thus, because total expected utility equals total expected efficiency minus the total expected payment, and because we have just proved that the latter amount (i.e., the total revenue as "seen from the other side") is $(\frac{1}{2}-\epsilon)MSW(TV_{-X})$, we lowerbound total efficiency as follows:

$U(TV, \mathcal{M}_{\mathbb{S},\epsilon,k}(BID')) \geq (\frac{1}{2}+\epsilon)MSW(TV_C) + (\frac{1}{2}-\epsilon)\frac{MSW(TV_{-X})}{k} \geq$
$(\frac{1}{2}-\epsilon)\frac{MSW(TV_{-C})}{k} + (\frac{1}{2}-\epsilon)\frac{MSW(TV_{-X})}{k} \geq (\frac{1}{2}-\epsilon)\frac{MSW(TV)}{k}$.

In fact, the second inequality holds because $\epsilon > 0$ and $k > 1$, and the third because MWS is sub-additive and $C \cup -X$ is total set of players. Thus also the total expected efficiency is as claimed. *Q.E.D.*

5 Acknowledgements

We are grateful to Amit Sahai for suggesting the addition of the parameter k to model approximate knowledge.

References

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

[2] Babaioff, M., Lavi, R., and Pavlov, E. Single-Value Combinatorial Auctions and Implementation in Undominated Strategies. *Symposium on Discrete Algorithms*, pp. 1054-1063, 2006.

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

[4] Chung, K.-S., and Ely, J.C. Ex-Post Incentive Compatible Mechanism Design. mimeo. 2003.

[5] Clarke, E.H. Multipart Pricing of Public Goods. *Public Choice* 11:17-33, 1971.

[6] Conitzer, V. and Sandholm, T. Failures of the VCG Mechanism in Combinatorial Auctions and Exchanges. *Proc. of the 5th International Joint Conference on Autonomous Agents and Multi Agent Systems*, pp. 521-528, 2006.

[7] Cramton, P. and Schwartz, J. Collusion in Auctions. *Annales d'Economie et de Statistique*, 15/16:217-230, 1989.

[8] Day, R. and Milgrom, P. Core-selecting Package Auctions. *International Journal of Game Theory*, 36:393407, 2008.

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

[10] Friedman, M. Comment on 'Collusion in the Auction Market for Treasury Bills'. *J. of Political Economy* 9:757-785, 1996.

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

[12] Goldberg, A., Hartline, J., and Wright, A. Competitive Auctions and Digital Goods. *SODA* 2001: 735-744.

[13] Goswami, G., Noe, T.H., and Rebello, M.J. Collusion in Uniform-Price Auctions: Experimental Evidence and Implications for Treasury Auctions. *Review of Financial Studies*, 72:513-514, 1964.

[14] Groves, T. Incentives in Teams. *Econometrica*, 41:617-631, 1973.

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

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

[17] Lehmann, D., O'Callaghan, L.I., and Shoham, Y. Truth Revelation in Approximately Efficient Combinatorial Auctions. *ACM Conf. on E-Commerce*, pp. 96-102, 1999.

[18] Likhodedov, A. and Sandholm, T. Approximating Revenue-Maximizing Combinatorial Auctions. *The Twentieth National Conference on Artificial Intelligence*, pp. 267-274, 2005.

[19] Moulin, H. and Shenker, S. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, 18:511-533, 2001.

[20] Myerson, R. Optimal Auction Design. *Mathematics of Operations Research*, 6: 58-73, 1981.

[21] Porter, R. Detecting Collusion. *Review of Industrial Organization*, 26:2:147-167, 2005.

[22] Rastegari, B., Condon, A., and Leyton-Brown, K. Revenue Monotonicity in Combinatorial Auctions. *AAAI*, 2007.

[23] Vickrey, W. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8-37, 1961.

Appendix

Two Main Properties of DST Mechanisms. Let us state two straightforward properties of DST mechanisms *in a way that applies to our collusive settings too.* (The first is an immediate consequence of the *opt-out* condition —that is, that by submitting the null valuation a player can guarantee that he wins nothing and pays nothing.)

DST-1: For all (probabilistic or not) DST mechanisms M, players i, and bid profile BID, we have:
$0 \leq E[M_p(BID)_i]$ and
$E[M_p(BID)_i] \leq E[BID_i(M_a(BID)_i)]$ whenever BID_i is a traditional bid (i.e., a valuation).

DST-2: For all deterministic DST mechanisms M, players i, and bid profiles BID and BID' such that $BID_{-i} = BID'_{-i}$, we have: $M_a(BID)_i = M_a(BID')_a$ implies $M_p(BID)_i = M_p(BID')_i$.

6 Our Structural Result

Theorem 2. *For any n, m and k, if M is a deterministic collusion-leveraging mechanism with knowledge factor $\frac{1}{k}$, efficiency factor y, and revenue factor z, then $\frac{y}{k}, \frac{z}{k} \leq \frac{8}{\log \min\{n,m\}-6}$.*

Proof. We distinguish two case:

Case 1. There is a bid profile BID, containing a non-null valuation, such that $BID_i(M_a(BID)_i) = 0\ \forall i$, that is, M does not allocate any player a set he declares he values, and

Case 2. There is no such a profile.

In Case 1, the efficiency and revenue factors of M are both 0. To see this, let C be a non-trivial subset of the players such that the subprofile BID_{-C} contains a non-null valuation, and let $\mathtt{C}$ be the collusive context in which the coalition is C, C has $\frac{1}{k}$-knowledge, TV_C consist of null valuations and $TV_{-C} = BID_{-C}$. Then, because the coalition does not value anything, and because by property DST-1 M cannot pay C, C's utility is ≤ 0 for any bid subprofile. Accordingly (no matter what C's utility may be when they bid their minimal collusive knowledge and all other players their true valuations), the bid profile BID is indeed $\mathfrak{B}_C$-rational, and yields 0 efficiency and (by DST-1 again) 0 revenue, although $MSW(TV) \geq MSW(TV_{-C}) > 0$. Below we thus assume that we are in Case 2.

Let $c = \lfloor \sqrt{\min\{n, m\}} \rfloor$. Consider the following distribution of single-minded valuation subprofiles for the players $\{1, \ldots, 2c\}$ and the first c^2 goods.

Recall that a valuation is single-minded if there exists a set of goods S and a positive real v such that its value for each superset of S is v, and 0 for any other set. We simply denote such a valuation by (S, v), and call S the *valued set* and v the *bid value* of the single-minded valuation.

The Distribution $\mathcal{BID}$. Considering the goods as being indexed by pairs (i, j) where $i, j \in [c]$, for each player $i \in \{1, \ldots, c\}$ let his valued set S_i be the set of all goods (i, ℓ), and for each player $i \in \{c + 1, \ldots, 2c\}$ let his valued set S_i be the set of all goods $(\ell, i - c)$. For each of the 2c players let his value for his desired set be drawn independently from the distribution R, defined as the uniform distribution over the set $\{1/\ell\ :\ \ell = 1, 2, \ldots, c\}$.

We note that, considering the c^2 goods as being arranged lexicographically in a $c \times c$ square, the valued set for a player $i \in \{1, \ldots, c\}$ is the ith row, and the valued set for the player $i + c$ is the ith column.

Consider the operation of a DST mechanism M on the distribution $\mathcal{BID}$. For notational convenience let $C = \{1, \ldots, c\}$, and denote by $-C$ the set $\{c + 1, \ldots, 2c\}$. For the sake of symmetry, we assume that the labels of the goods and players have been removed (or randomly permuted), so that M treats players in C and players in $-C$ equally.

We denote by $\mathcal{A}$ and $\mathcal{P}$ the distributions of allocations and price profiles induced by running M on $\mathcal{BID}$. We claim that either there is a bid profile in $\mathcal{BID}$'s support for which M does not allocate any of the players' valued sets to them, or the following three properties hold.

1. With probability $< \frac{1}{4}$, $MSW(\mathcal{BID}_{-C}) \leq \log_e c - \frac{2\pi}{\sqrt{6}}$;
2. With probability $\leq \frac{1}{4}$, $\sum_{i \in C} \mathcal{P}_i \geq 4$; and
3. With probability $\frac{1}{2}$, $SW(\mathcal{BID}_C, \mathcal{A}) = 0$.

Proof of Property 1. We note that the expected value of a player in $\mathcal{BID}$ is $E[R] = \sum_{\ell \leq c} \frac{1}{\ell c}$. Since the valued sets of players in $-C$ are disjoint, the expected "total value" of the c players in $-C$ is $E[MSW(\mathcal{BID}_{-C})] = \sum_{\ell \leq c} \frac{1}{\ell} \geq \log_e c$. Also, the variance of R is $E[R^2] - E[R]^2 \leq E[R^2] = \sum_{\ell \leq c} \frac{1}{\ell^2 c} < \frac{1}{c} \sum_{\ell} \frac{1}{\ell^2} = \frac{\pi^2}{6c}$. Thus, the variance of $MSW(\mathcal{BID}_{-C})$ is less than $\frac{\pi^2}{6}$, and thus $MSW(\mathcal{BID}_{-C})$ has standard deviation less than $\frac{\pi}{\sqrt{6}}$. By Chebyshev's inequality, the probability that $MSW(\mathcal{BID}_{-C})$ is less than its mean, $\log_e c$, minus $\frac{2\pi}{\sqrt{6}}$ is less than $\frac{1}{4}$. ■

Proof of Property 2. Fix a player $i \in C$ and a sub-profile BID'_{-i}, and let $\mathcal{B}$ be the distribution $BID'_{-i} \sqcup \mathcal{BID}_i$. We show that $E_{\mathcal{B}}[\mathcal{P}_i] \leq \frac{1}{c}$. Denoting by $\overline{\mathcal{B}}$ the support of $\mathcal{B}$, we have two mutually exclusive possibilities:

- $\forall BID \in \overline{\mathcal{B}},\ M_p(BID)_i = 0$, and
- $\exists BID \in \overline{\mathcal{B}},\ M_p(BID)_i > 0$.

The first possibility trivially implies that $E_{\mathcal{B}}[\mathcal{P}_i] = 0$. The second one, since M is truthful, $P_i > 0$ implies $A_i = S_i$ (by property DST1). Therefore, consider the set of bid values X in the domain of R such that, whenever $x \in X$ and $(A, P) = M(BID'_{-i} \sqcup (S_i, x))$ we have $P_i > 0$. (By the assumption of this case the bid value of player i is in X, so thus $X \neq \emptyset$.) Since M is truthful, (by property DST2) M must assign the same price to player i for each bid (S_i, x) for $x \in X$; we denote this price by p. Further, since M is truthful, (by property DST1) i's bid value, x, for S_i must be at least p when i wins; namely, $\forall x \in X, x \geq p$. Since elements of X are inverses of integers, $1/\ell$, we see that there are $\lfloor \frac{1}{p} \rfloor$ possible integers whose inverses are at least p, and thus $|X| \leq \frac{1}{p}$. Since $\mathcal{B}$ draws a bid value for S_i from $\{1, \ldots, \frac{1}{c}\}$ uniformly, the probability that i wins (in $\mathcal{B}$) is at most $\frac{|X|}{c} \leq \frac{1}{cp}$. Since the price to player i is p when i wins and 0 otherwise, $E_{\mathcal{B}}[\mathcal{P}_i] \leq \frac{1}{c}$. This concludes the analysis of Case 2, and thus the proof that $E_{\mathcal{B}}[\mathcal{P}_i] \leq \frac{1}{c}$ for all possible bids BID'_{-i}.

Let us now finish the proof of Property 2. The fact that $E_{\mathcal{B}}[\mathcal{P}_i] \leq \frac{1}{c}$ for all possible bids BID'_{-i} implies that, over the whole distribution $\mathcal{BID}$, we have $E[\mathcal{P}_i] \leq \frac{1}{c}$. Summing over $i \in C$ we have $E[\sum_{i \in C} \mathcal{P}_i] \leq 1$. We now invoke Markov's inequality to conclude that with probability $\leq \frac{1}{4}$, $\sum_{i \in C} \mathcal{P}_i \geq 4$. ■

Proof of Property 3. Because we are assuming that we are in Case 2, and thus that M always allocates some player his valued set, we have that $SW(\mathcal{BID}, \mathcal{A}) > 0$. Since the valued set of any player from C intersects the valued set of any player from $-C$, we must have exactly one of $SW(\mathcal{BID}_C, \mathcal{A})$ or $SW(\mathcal{BID}_{-C}, \mathcal{A})$ positive, and thus by symmetry either occurs with probability $\frac{1}{2}$. ■

Applying the union bound, we see that the probability that none of the above 3 properties hold in our experiment is > 0. Therefore,

$\exists$ *a bid profile BID' in the support of $\mathcal{BID}$ such that, letting $(A', P') = M(BID')$, we have:*
$(1')$ $MSW(BID'_{-C}) > \log_e c - \frac{2\pi}{\sqrt{6}}$, $(2')$ $\sum_{i \in C} P'_i < 4$, *and* $(3')$ $SW(BID'_C, A') > 0$.

Fix now $\delta \in (0, 1)$, and let $\mathbb{C}^\delta$ be the collusive context where the coalition is C, and the true-valuation profile for the players is TV^δ, defined as: $TV^\delta_i = BID'_i$ for all $i \in -C$, and $TV^\delta_j = (S_j, P'_j + \delta)$ for all $j \in C$.

Consider the collusive auction $\gamma^\delta = (\mathbb{C}^\delta, M)$, and, further, the result of applying mechanism M to the bid profile BID': $(A', P') = M(BID')$. By construction, C's utility in this case is

$$u'_C = \sum_{j \in C} TV^\delta_j(A'_j) - P'_j = \sum_{j \in C : A'_j \supseteq S_j} TV^\delta_j(S_j) - P'_j = \sum_{j \in C : A'_j \supseteq S_j} \delta$$

where the first equality holds by definition of utility, the second one by DST1, and the third one because, by definition, $TV^\delta_j(S_j) = P'_j + \delta$ whenever $j \in C$. Because $\delta > 0$, we show that u'_C is positive by noting that the summation of the last term is taken over a nonempty set: by Property 3', A' satisfies $SW(BID'_C, A') > 0$, so there exists $j \in C$ for which $A'_j = S_j$.

Consider now the set of $\mathfrak{B}_C$-rational bid profiles for the collusive auction γ^δ. Either BID' is $\mathfrak{B}_C$-rational, or it is not $\mathfrak{B}_C$-rational because M and BID'_{-C} are such that when all coalition members bid their minimal collusive knowledge C obtains higher utility. In either case, there exists a $\mathfrak{B}_C$-rational bid profile with strictly positive utility for the coalition; denote such a bid profile by BID^*. Let $(A^*, P^*) = M(BID^*)$ and let u^*_C be the corresponding utility for the coalition. Then, since M is truthful, (by property DST1) $u^*_C > 0$ implies that there exists $j \in C$ such that $A^*_j \supseteq S_j$. In turn, since for each $i \notin C$, $S_i \cap S_j \neq \emptyset$, this implies that, since by definition $BID'_{-C} = TV^\delta_{-C}$ we have

$$SW(TV^\delta_{-C}, A^*) = 0. \tag{1}$$

Recall that, by definition, for each $j \in C$, $TV^\delta_j(S_j) = \delta + P'_j$, and thus we have

$$SW(TV^\delta, A^*) = SW(TV^\delta_C, A^*) \leq \sum_{j \in C} TV^\delta_j(S_j) \leq c\delta + \sum_{j \in C} P'_j < c\delta + 4,$$

where the first equality is from Equation 1, and the last inequality is from Property $(2')$.

Also, from property $(1')$, $MSW(TV^\delta) \geq MSW(TV^\delta_{-C}) > \log_e c - \frac{2\pi}{\sqrt{6}}$.

Thus, since $\frac{2\pi}{\sqrt{6}} < 3$, for δ sufficiently small we have $\frac{SW(TV^\delta, A^*)}{MSW(TV^\delta)} \leq \frac{4}{(\log_e c) - 3} = \frac{8}{(\log_e \min\{m,n\}) - 6}$, which yields the desired result that $\frac{y}{k} \leq \frac{8}{(\log_e \min\{m,n\}) - 6}$.

Further, since the coalition has positive utility and this utility is the difference between the efficiency allocated to them and the revenue they pay, we conclude that the revenue collected from the coalition is bounded exactly as efficiency; by property DST1 the independent players, who do not receive any goods, cannot be charged anything. Thus the revenue factor, as the efficiency factor, is bounded as $\frac{z}{k} \leq \frac{8}{(\log_e \min\{m,n\})-6}$.

Q.E.D.

7 Our "Negative" Result

This asymptotic result (in our computer science tradition) is of theoretical importance only, as is often the case for impossibility results. Accordingly, we have made no attempt to improve any of our constants.

We begin with a probabilistic bound that will be useful in the following analysis.

Lemma 1. *Given a distribution D over the positive reals of mean m where every element occurs with frequency at least $\frac{1}{x}$, then the expected value of the maximum over z trials of the sum of y samples from D is at most $my(1+2\sqrt{(x/y)\log_e z})$, provided $\sqrt{(x/y)\log_e z} \leq 1$.*

Proof. Our proof makes use of two techniques: the first is the observation that the maximum of several samples from a distribution is bounded by the sum of those samples; the second is the transformation that takes a distribution D to the distribution $e^{\lambda D}$, as is done to prove the Chernoff bounds.

Let $\lambda = \frac{1}{m}\sqrt{(y/x)\log_e z}$, and consider the distribution $e^{\lambda \sum_{i=1}^{y} D_i}$, where we use D_i to denote independent samples from D. The maximum of z samples from this distribution is distributed exactly as $e^{\lambda M}$ where M denotes the maximum of z samples from $\sum_{i=1}^{y} D_i$, the distribution whose expectation we wish to bound. Further, since the logarithm function is concave, we have that the desired quantity is at most the logarithm (base e^{λ}) of the maximum of z samples from $e^{\lambda \sum_{i=1}^{y} D_i}$; as noted above, the maximum of z samples is bounded by the sum of z samples, so we may instead use the bound

$$E[\frac{1}{\lambda}\log_e \sum_{j=1}^{z} e^{\lambda \sum_{i=1}^{y} D_{ij}}]. \tag{2}$$

We evaluate this expression by noting that $e^{\lambda \sum_{i=1}^{y} D_i} = \prod_{i=1}^{y} e^{\lambda D_i}$, and further, the expectation of the product of independent terms is the product of the expectations, and likewise for the sum. Thus Equation 2 equals $\frac{y}{\lambda}\log_e E[e^{\lambda D}] + \frac{\log_e z}{\lambda}$.

Since exponentiation is a convex function, $E[e^{\lambda D}]$ decreases for distributions that are distributed further from their mean m than D is; since D is a distribution on the positive reals with every element appearing with probability at least $\frac{1}{x}$, the value of $E[e^{\lambda D}]$ is less than or equal to the corresponding value for the distribution $\widetilde{D}$ that takes value mx with probability $\frac{1}{x}$ and value 0 the remainder of the time. Thus we have $E[e^{\lambda \widetilde{D}}] = \frac{1}{x}e^{\lambda mx} + (1-\frac{1}{x})e^0$. Since by assumption $\lambda mx \leq 1$, and $e^{\gamma} \leq 1+\gamma+\gamma^2$ for $\gamma \leq 1$, we bound this by $1+\lambda m+\lambda^2 m^2 x = 1+\lambda m(1+\sqrt{\frac{x\log_e z}{y}})$. Using this as our bound for $E[e^{\lambda D}]$, we bound its log by noting that $\log_e \gamma \leq 1+\gamma$, and thus bound Equation 2 by $\frac{y}{\lambda}\lambda m(1+\sqrt{\frac{x\log_e z}{y}}) + \frac{\log_e z}{\lambda} = m(y+2\sqrt{xy\log_e z})$, as desired.

Q.E.D.

Theorem 3. *For Bayesian auction contexts with n players, m goods, and a secret coalition with knowledge of the Bayesian, letting $\mu = \min\{n,m\}$, and $\gamma = \frac{\log \mu}{\log\log \mu}$, no mechanism can guarantee efficiency $\Omega(\frac{1}{\gamma})$.*

Proof. We prove that for large enough m,n, it suffices to take $\gamma = \frac{\log \mu}{6\log\log \mu}$.

We begin by describing a class of distributions for bid profiles, which we call the *colors distributions*.

Let k be the greatest prime less than or equal to $\sqrt{\mu}$, and let ℓ be an integer less than or equal to $k^{6/7}$. We define a distribution of single-minded bid profiles on $k\ell$ players and k^2 goods in two steps: first we define each player's valued set; then we define the distribution of values each player has for his valued set.

- Let the players be indexed by pairs (i, j) where $i \in [\ell]$ and $j \in [k]$, and where players with first index i are said to be of the "ith color". Let the goods be indexed by pairs (a, b) where $a, b \in [k]$. We define player (i, j)'s valued set to be the set of (a, b) such that $b \equiv ai + j (\mod k)$.
- The value of a player for his valued set is chosen independently from the following distribution: draw an integer r uniformly at random from 1 to k; if r is 1, return value $\frac{9}{7}\gamma + 1$, and otherwise return value $1/r^{1-\frac{7}{9\gamma}}$.

Lemma 2. *Bid profiles drawn from any colors distribution are such that every pair of players of the same color has disjoint valued sets, while every pair of players of different colors has intersecting valued sets.*

Proof. To prove the first claim, consider two players (i, j) and (i, j') with $j \neq j'$ both of color i, and suppose for the sake of contradiction that both players' desired sets contain the good (a, b). Then from the definition of the colors distributions we have $ai + j \equiv b \equiv ai + j' (\mod k)$, which implies $j = j'$, the desired contradiction.

To prove the second claim, we consider two players (i, j) and (i', j') where $i \neq i'$, and consider the condition for when good (a, b) is in both players' valued sets: $ai + j \equiv b \equiv ai' + j' (\mod k)$. We solve for a as $a = \frac{j-j'}{i'-i}$, where the division is done in the field $\mathbb{Z}_k$ and is well-defined because $i \neq i'$; letting $b = ai + j$ yields the intersection point for the two valued sets, implying that they do intersect, as claimed. ∎

We now analyze the revenue obtainable by an auction when the bidders bid from a colors distribution.

Lemma 3. *If an auction mechanism M is run on bids from a colors distribution and yields expected revenue R and expected efficiency E (calculated assuming the bids are truthful), then $R \leq 2k^{2/9}\gamma + \frac{7E}{9\gamma}$.*

Proof. We begin by noting a standard fact about DST mechanisms for single-minded auctions from the perspective of a player i. Fixing the bids of the other players and fixing the random coins of the mechanism, the mechanism is specified by a number $\theta \in \mathbb{R}$ such that the mechanism yields allocation and price to player i as: if i bids more than θ for his valued set then allocate him his valued set for price θ; otherwise allocate him nothing for price 0. (The case when i bids exactly θ can be grouped with either of the above cases and this decision generally does not affect the analysis.)

With this characterization of DST mechanisms for single-minded bidders in mind, we analyze mechanism M as the "sum" of two mechanisms: let M^+ be the mechanism that simulates M but then only allocates goods and collects payments if the price to be paid is ≥ 1; and let M^- be the corresponding mechanism that only charges < 1. The expected revenue from M is clearly the sum of the expected revenues from M^+ and M^-. We bound these respectively as $2k^{2/9}\gamma$ and $\frac{7E}{9\gamma}$ to yield the theorem.

To bound the revenue from M^+ we note that M^+ only collects money from some player i if player i bids at least 1, which only happens when player i bids $\frac{9}{7}\gamma + 1$; we will call this "bidding high". Further, M^+ cannot collect more than $\frac{9}{7}\gamma + 1$ from any player, since no matter what the threshold θ is, the mechanism cannot collect more revenue from a player than that player's bid. Thus the revenue of M^+ is at most $\frac{9}{7}\gamma + 1$ times the largest set of players who bid high for disjoint sets, namely (by Lemma 2) $\frac{9}{7}\gamma + 1$ times the largest number of high bids from any of the $\ell \leq \sqrt{k}$ colors. Applying Lemma 1 to the distribution that takes value 1 if and only if a given player bids high yields that the expected revenue of M^+ is at most $\frac{9}{7}\gamma + 1$ times $k^{-7/9}k(1 + 2\sqrt{(k^{7/9}/k)\log_e k})$ yielding a bound of $2k^{2/9}\gamma$.

We now bound the revenue of M^-, as a fraction of its efficiency. Consider the set of bid values a colors distribution may take: let $s(r)$ be the sequence where $s(1) = \frac{9}{7}\gamma + 1$ and $s(r) = 1/r^{1-\frac{7}{9\gamma}}$ for $r > 1$. We bound the partial sums of this sequence as follows: $\sum_{r \leq j} s(r) \geq \frac{9}{7}\gamma + 1 + \int_2^{j+1} 1/x^{1-\frac{7}{9\gamma}} dx \geq \frac{9}{7}\gamma + \int_1^{j+1} 1/x^{1-\frac{7}{9\gamma}} dx \geq \frac{9}{7}\gamma + \int_1^j 1/x^{1-\frac{7}{9\gamma}} dx = \frac{9}{7}\gamma + \left(\frac{9}{7}\gamma x^{\frac{7}{9\gamma}} \Big|_1^j \right) = \frac{9}{7}\gamma j^{\frac{7}{9\gamma}} = \frac{9}{7}\gamma j / j^{1-\frac{7}{9\gamma}}$. Thus for any $j > 1$ we have that the average of the first j terms of $s(r)$ is at least $\frac{9}{7}\gamma s(r)$. Further, we note that the average of the first j elements is exactly the expected efficiency contributed by some player i when the price threshold θ is between $s(j)$ and $s(j+1)$, conditioned on the fact that player i bids high enough to win his set; simultaneously, $s(j)$ is an upper bound on θ, the price he will pay in these cases. Thus $\frac{9}{7}\gamma$ bounds the ratio between efficiency and revenue for M^-. Since the efficiency of M^- is at most the efficiency of M, namely E, we have that M^- yields revenue at most $\frac{7E}{9\gamma}$.

Thus the expected revenue produced by M is at most $2k^{2/9}\gamma + \frac{7E}{9\gamma}$, as desired. ▮

We are now equipped to prove the theorem.

Assume for the sake of contradiction that there is a mechanism M that creates an incentive structure such that it is the dominant strategy for the independent players to bid truthfully, and where these truthful bids along with the optimal bid subprofile for a coalition with Bayesian knowledge induce an allocation that in expectation yields efficiency greater than $\frac{1}{\gamma}$.

Our proof involves constructing a sequence of coalitional auction contexts, showing how each coalition has utility substantially greater than utility of the previous one in the sequence, and deriving a contradiction by noting that the coalition of the final context would have an impossibly high utility.

Let $\ell = 2^{\lfloor \log k \rfloor}$; each auction will have $k\ell$ players, constituting ℓ colors of players, in the sense of the definition of the colors distributions. We construct $1 + \log \ell$ contexts, labeled from 0 up to $\log \ell$, where in auction context j a random fraction 2^{-j} of the groups are independent, with the rest constituting the coalition. Each group is assigned a color, where the true values for the independent players are taken from the corresponding entries of the colors distribution, and the true values for the coalition players are $\frac{8}{9\gamma}$ times the corresponding values from the colors distribution.

We note that in context 0 there are no coalitional players so hence the coalition derives 0 utility. Let $BID^{(j)}$ be an optimal bid for the coalition in context $j \geq 1$. The independent players in this context will bid truthfully their values drawn from the (k, ℓ) colors distribution. Consider context $j+1$, and specifically, consider the case where the coalition bids so as to emulate the jth context: it instructs $\ell(1-2^{-j})$ colors of players to bid from $BID^{(j)}$, with its remaining $\ell 2^{-(j+1)}$ colors of players bidding from the colors distribution. The utility of the coalition here is the sum of two terms, one from the portion of the coalition pretending to be the coalition from the jth context, the other from the $\ell 2^{-(j+1)}$ "hidden" coalition colors, where the first of these terms exactly equals the utility the coalition in context j received. Thus the amount by which the coalition's utility in context $j+1$ exceeds its utility in context j is exactly the second term, the utility of the "hidden" coalition members. We analyze this now.

Consider the performance of M on context j; for convenience we denote the expected value of the total bids of one color by msw_c. By assumption, M yields expected efficiency at least $\frac{1}{\gamma}$ times the expected MSW of the true values of the players, which is at least $\frac{1}{\gamma} msw_c$. Further, we assumed the members of the coalition have true values drawn from $\frac{8}{9\gamma}$ times values drawn from the colors distribution, and by Lemma 1 the expectation of their MSW is thus at most $\frac{8}{9\gamma} msw_c(1 + 2\sqrt{\frac{\log_e k}{k^{2/9}}})$, which we bound as $\frac{17}{18\gamma} msw_c$ for large enough k. Thus, for the mechanism to reach its efficiency target, it needs to extract at least $\frac{1}{18\gamma} msw_c$ efficiency from the independent players. Let E be the efficiency extracted, whereby $E \geq \frac{1}{18\gamma} msw_c$. Lemma 3 yields that the revenue from the independent players is at most $2k^{2/9}\gamma + \frac{7E}{9\gamma}$.

We now consider this situation in the light of context $j+1$, where the coalition merely simulates the bids of context j: half of the simulated independent players are actually coalitional players; since this is a random half, and the bid distributions are identical, half of the goods will be allocated to the coalition, and half of the revenue will come from them. This revenue is thus at most half the bound above, $k^{2/9}\gamma + \frac{7E}{18\gamma}$. The efficiency yielded to the coalition is not half of the total efficiency E, however —since the coalition values the goods at only $\frac{8}{9\gamma}$ fraction of their bids for them— but rather $\frac{4E}{9\gamma}$. Thus the utility to the "hidden" portion of the coalition is at least $\frac{4E}{9\gamma} - [k^{2/9}\gamma + \frac{7E}{18\gamma}] = \frac{E}{18\gamma} - k^{2/9}\gamma \geq \frac{msw_c}{364\gamma^2} - k^{2/9}\gamma$. We bound msw_c (as above) as $k\text{mean}_r[s(r)] = k/k^{7/9}\int_0^{k^{7/9}} x^{1-\frac{7}{9\gamma}} dx = k^{\frac{2}{9}} \frac{9}{7}\gamma k^{-\frac{7}{9\gamma}} = \frac{9}{7}\gamma k^{\frac{2}{9}} k^{-\frac{7 \log\log k}{3 \log k}} = \frac{9}{7}\gamma k^{\frac{2}{9}} (\log k)^{-\frac{7}{3}}$. Thus for large enough k, the utility of the "hidden" portion of the coalition is at least $\frac{msw_c}{600\gamma^2} = \frac{msw_c}{\gamma \log k}\frac{\log\log k}{200}$.

As argued above, this utility lowerbounds the incremental improvement the coalition gains each time they "take over half the remaining players". Thus the utility of the coalition in the final context, context $\log k$, the utility of the coalition is at least $\frac{msw_c}{\gamma}\frac{\log\log k}{200}$. However, the highest possible expected utility the coalition could receive is their expected MSW (since by DST1 they cannot be paid by the auction), which we bounded as $\frac{17}{18\gamma}msw_c$; comparing these two expressions yields the desired contradiction. *Q.E.D.*