

RESILIENT MECHANISMS FOR TRULY COMBINATORIAL AUCTIONS

Silvio Micali and Paul Valiant

November 17, 2008

Abstract

Dominant-strategy truthfulness is traditionally considered the best possible solution concept in mechanism design, as it enables one to predict with confidence which strategies *independent* players will actually choose. Yet, as with any other form of equilibrium, it too can be extremely vulnerable to *collusion*. The problem of collusion is particularly evident for *unrestricted combinatorial auctions*, arguably the hardest type of auctions to design mechanisms for.

We thus investigate how much revenue can be guaranteed, in unrestricted combinatorial auctions, by dominant-strategy-truthful mechanisms that are *collusion-resilient* in a very strong sense; and obtain almost matching upper- and lower-bounds.

1 Introduction

The context of an auction of a single good is very simple: each player i has a secret value v_i for the good, and any possible outcome of the auction will specify which player wins the good and which price each player will pay. In this context, the problem of designing an economically efficient auction consists of specifying a mechanism—for each player a set of actions he may choose among, and an outcome function mapping the vector of chosen actions to an outcome—so as to guarantee that, when in the resulting game everyone plays rationally (i.e., acts so as to maximize his own utility), the good is won by the player who values it the most. The problem is nontrivial because “the player who values it the most” depends on the players’ secret values v_i , which by definition are not available to the auctioneer. This particular problem is however simply solved by the famous second-price mechanism, where the players’ actions consist of simultaneously “bidding” a real number, and the outcome function allocates the good to the player with the highest bid at a price equal to the second highest bid. (All other players pay nothing.)

Generalizing, consider a context in which there is a set of possible outcomes which could result from the players’ actions, each player has a secret *type*, and the players’ utilities depend in a given fashion upon the resulting outcome and their types. Given a desired property P defined on the outcome and the players’ types, mechanism design (at its highest level) aims at finding a mechanism such that, in the resulting game, P holds whenever the players’ chosen strategies constitute some equilibrium σ . A player’s strategy (in a normal-form game) is just a distribution over the actions available to him, and an equilibrium σ is a profile (i.e., a vector indexed by the players) of strategies, such that no player i wants to deviate from σ_i as long as he believes that every other player j sticks to σ_j . That is, σ is an equilibrium if, for all players i , denoting the subprofile of strategies of the other players by σ_{-i} , the utility of player i under the strategy profile (σ_i, σ_{-i}) is greater than or equal to his utility under the profile (σ'_i, σ_{-i}) for any possible alternative strategy σ'_i available to him.

1.1 Resilient Mechanism Design

Mechanism design is important and appealing, but not robust. It is vulnerable to two well known problems: (1) *Equilibrium selection* and (2) *Collusion*. Let us explain.

A mechanism M is designed to achieve its desired property P at equilibrium. But what if the resulting game G has multiple equilibria? In this case, one cannot be sure of which equilibrium will be played out. Accordingly, a stronger version of mechanism design calls for designing M so that P holds at each of the possible equilibria of G . This, however, is still a rather weak guarantee, even if the number of possible equilibria is just 2. Indeed, if σ and σ' are two equilibria of G , and a player i —believing that σ will be the selected equilibrium—chooses his strategy to be σ_i , while another player j —believing that the selected equilibrium will be σ' —chooses his strategy to be σ'_j , then G ’s play does not end up in any equilibrium at all, so that—putting it provocatively—whether P holds is “more a question of luck than of rationality.” The equilibrium selection problem goes away if M is *dominant-strategy truthful* (DST for short). This means that in the resulting G , no matter what strategies the other players may choose, the best strategy for each player i consists of the action coinciding with his own type, t_i . When such a mechanism M can be found, no matter how many other equilibria there may be in the resulting game, we can reasonably predict that rational players will choose the strategy profile (t_1, \dots, t_n) .

However, as for any equilibrium notion, dominant-strategy truthfulness is *solely defined relative to single-player deviations*. Accordingly, although a single player may not have any incentive to choose a strategy different from t_i , two players i and j may have all the incentive in the world to *collude*, in essence to deviate from their equilibrium strategies in a joint and coordinated way.

We advocate that a desired property should be achieved by resilient mechanisms: that is, by means of mechanisms that are immune to the problems of both equilibrium selection and collusion.

To be really meaningful, resilient mechanism should be able to withstand *arbitrarily many colluders*, belonging to arbitrarily many *collusive sets*, each capable of *perfectly coordinating* the actions of its members (e.g., via secret binding contracts) in order to maximize its own *collective utility*.

1.2 Truly Combinatorial Auctions and Their Difficulties

Advocating is cheap, and we advocate raising the bar in mechanism design only because we believe it is possible to jump over higher bars. We demonstrate this possibility by exemplifying resilient mechanism design *not* for some new applications —as such applications may always be suspected to be uniquely easy— but rather for a “champion application,” one recognized to be hard even before bringing the problems of collusion and equilibrium selection into the picture. Indeed, we exemplify resilient mechanism design for the champion application of *generating revenue in truly combinatorial auctions*.

In a combinatorial auction there are multiple goods for sale, and each player i has a private *true valuation* for the goods—that is a function, denoted by TV_i , mapping each possible non-empty subset S of the goods to a non-negative number (representing i 's value for S). The difficulty of combinatorial auctions is perhaps evidenced by the many possible restrictions of them that are studied; in particular (1) sub-modularity, (2) additive valuations, (3) free disposal, (4) single mindedness, (5) unit demand, and (6) unlimited supply.¹ In this paper, however, *we assume no restrictions whatsoever for combinatorial auctions*: whenever S and T are distinct subsets of goods, nothing can be inferred about $TV_i(S)$ from $TV_j(T)$. To emphasize that the players' true valuations can indeed be arbitrary, we may use the adjective **truly** in front of combinatorial.

The only general mechanism known for truly combinatorial auctions is the famous VCG mechanism [7, 16, 25]. Given any profile of bids (valuations) for a combinatorial auction, the VCG returns an outcome as follows; the allocation is the one maximizing the social welfare of the bids (i.e., the sum of the value that each player declares for the goods assigned to him); and each player i winning goods in this allocation pays a price equal to the difference of (1) the social welfare of the other winners—computed according to their bids—in the returned allocation and (2) the maximum social welfare of all players but i —again computed according to their bids. It can be proved that these rules make the VCG dominant-strategy truthful (DST); for each player, bidding his own true valuation is the best strategy. Accordingly, in any combinatorial auction, the VCG mechanism *maximizes social welfare in dominant strategies*. This is the strongest possible traditional guarantee in mechanism design. Yet, the following example of Ausubel and Milgrom [1] proves that the VCG mechanism (1) *is totally vulnerable to collusion*, and (2) *has no revenue guarantees*.

Let there be two goods for sale, a and b , and three players, P_1 , P_2 and P_3 . Player P_1 values only good a , for a negligible amount ϵ ; P_2 values only good b , also for ϵ ; while P_3 values only the pair of goods $\{a, b\}$, and for a high amount z . Assume now that that P_1 and P_2 are minimally informed about P_3 's valuation. Namely, they know that he value just the pair $\{a, b\}$ for at most w . Then P_1 and P_2 can get for free the goods they value as follows: P_1 bids w for a , and P_2 bids w for b . Since the VCG is dominant strategy truthful, P_3 bids z for $\{a, b\}$. With such bids, the VCG allocates a to P_1 for a price 0 and b to P_2 also for price 0.

Since the same bids may have arisen from honest play—i.e., when the true valuations of P_1 for a and of P_2 for b really were w —the example also demonstrates that the VCG revenue can very well be 0 even when all players are independent and have high valuations for the goods, and in the presence of lots of “competition” for the goods.

This example suggests a disturbing possibility: in principle, every possible dominant-strategy mechanism may be vulnerable to collusion and return 0 revenue even when the social welfare of the players' true valuations is extremely high. Accordingly,

We want to investigate how much revenue can be resiliently guaranteed in a truly combinatorial auction.

¹That is, (1) $TV_i(S \cup T) \leq TV_i(S) + TV_i(T)$ for any subsets S and T of the goods; (2) $TV_i(S) = TV_i(g_1) + \dots + TV_i(g_k)$ whenever $S = \{g_1, \dots, g_k\}$; (3) $TV_i(S) \leq TV_i(T)$ whenever $S \subset T$; (4) for each i there is a subset of goods S and a value v such that $TV_i(T) = v$ if $T \supset S$, and 0 otherwise; (5) each player has nonzero valuations only for sets of size one; and (6) informally, an unbounded number of copies of each good are available, and each player values only sets of distinct goods.

1.3 Our Notion of Revenue Resiliency for Truly Combinatorial Auctions

We evaluate the performance of a revenue mechanism relative to a *benchmark*, that is a function from valuation profiles to non-negative numbers. We think of a benchmark \mathbb{B} as expressing the revenue we are “content with” for every possible true-valuation profile TV of the players. Let us now provide two definitions for what it means for a revenue mechanism to achieve a fraction c of a benchmark \mathbb{B} in a truly combinatorial auction. The first one simply expresses in the language of truly combinatorial auctions the definition previously used by other computer scientists for restricted auctions (in particular by [4, 10, 14, 13, 17]).

Definition 1. *We say that a DST mechanism M achieves a fraction c of a benchmark \mathbb{B} if, for all numbers of players n , numbers of goods m , and valuation profiles BID :*

The (expected) revenue of $M(BID)$ is at least $c \cdot \mathbb{B}(BID)$.

Notice that Definition 1 captures the “first half” of our intuitive notion of resiliency: indeed, the fact that M is DST implies that rational and independent bidders will bid truthfully so that (1) $BID = TV$, dispelling any problems of equilibrium selection, and (2) M achieves revenue at least $c \cdot \mathbb{B}(TV)$, realizing the intuition of the previous paragraph. Accordingly, it would be interesting to provide the first mechanisms for truly combinatorial auctions satisfying Definition 1 (for a non-trivial choice of \mathbb{B} and c). We wish however to construct mechanisms satisfying our second definition, a stronger variant of Definition 1 capturing “the remaining half of resiliency:” that is, collusion resiliency.

This work’s notion of collusion resiliency. In this work, we take the position that an auctioneer could be happy if, by magic, *all collusive players spontaneously disappear, leaving him to run the auction with just the independent players alone*. Our mechanism effectively “simulates” this presumably unlikely event. Essentially, it guarantees that, in any auction with x independent players and y collusive players, its revenue performance will be at least as good as that of an auction with the same x independent players alone.

This work’s model of collusion. We want to neutralize any y collusive players in a very general sense. For instance, the collusive players may be partitioned into a multiplicity of secret coalitions. These coalitions may or may not be aware of each other’s existence, may or may not have precise information of the true valuations of their members, may or may not have information about the true valuations of the other players, may coordinate their actions in a variety of ways (e.g., by entering their own secret and binding contracts), may or may not make side payments to one another, may try to maximize special utility functions, etc. etc. The literature lacks such an intricate model, and we are not going to formalize one. Rather, we avoid this issue with a very simple yet completely general approach: the universal quantifier. Since the collusive players influence the auction through their bids, and only through their bids, denoting by I the set of independent players, we model (or rather, avoid modeling) all possible behaviors of the collusive players by inserting the phrase “for all bid sub-profiles BID_{-I} ” in every result about collusion-resiliency.

Definition 2. (Resilient Revenue Mechanisms) *We say that a DST mechanism M resiliently achieves a fraction c of a benchmark \mathbb{B} if, for all numbers of players n , for all numbers of goods m , for all true-valuation profiles TV , for all independent-player subsets I , and for all bid sub-profiles BID_{-I} :*

The (expected) revenue of $M(TV_I \sqcup BID_{-I})$ is at least $c \cdot \mathbb{B}(TV_I)$.²

(Let us illustrate Definition 2 by virtue of an explicit example. Suppose we run an auction with twenty goods and ten players, where the last five players in fact belong to two secret coalitions, namely the subsets of players $\{6, 7, 8\}$ and $\{9, 10\}$. The first coalition decides to further its interests by having its members bid valuations respectively denoted as BID_6 , BID_7 and BID_8 . The players of the second coalition instead bid

² $TV_I \sqcup BID_{-I}$ denotes the profile X for which $X_I = TV_I$ and $X_{-I} = BID_{-I}$.

valuations denoted by BID_9 and BID_{10} . As for the other players, since they are independent and M is DST, they will bid their true valuations, denoted TV_1, \dots, TV_5 . Definition 2 then guarantees that, setting $I = \{1, \dots, 5\}$, the collected revenue is at least $c \cdot \mathbb{B}(TV_{\{1, \dots, 5\}})$: that is, the same revenue we would be “content with” had our five collusive players vanished from the auction.)

1.4 Our Contributions

- We introduce a natural and very aggressive benchmark, $\text{MSW}_{-\star}$, to evaluate the revenue performance of mechanisms in truly combinatorial auctions. (In essence, $\text{MSW}_{-\star}$ generalizes to truly combinatorial auctions the second-highest-bid revenue benchmark of single-good auctions).
- We exhibit a DST mechanism \mathbb{M} that in such auctions yields (asymptotically) optimal performance relative to the $\text{MSW}_{-\star}$ benchmark. We prove both \mathbb{M} ' performance and the impossibility of any DST mechanism, whether deterministic or stochastic, to perform essentially better than \mathbb{M} .
- We show that $\text{MSW}_{-\star}$ is in fact a member of a special class of benchmarks, the *bidder-monotone* benchmarks, membership in which immediately transforms any revenue guarantee (relative to such a benchmark) into a stronger guarantee that such revenue is impervious to collusion.

We present this group of tightly complementary results in the hope to lay the theoretical groundwork for resilient mechanism design, and in particular truly combinatorial auctions that are both profitable and robust.

2 Prior Work

Several notions of “collusion-resiliency” appear in the literature. Focusing on auctions, the subject matter of this paper, we note the notion of a *group strategyproof* mechanism. Essentially, such a mechanism discourages collusion in that any gain for a collusive player is accompanied by a loss for another collusive player. Notable examples of group strategyproof mechanisms are those of [18, 21, 11]. Such a notion of collusion resiliency, however, is only meaningful when collusive players cannot make side payments to one another. No such restrictions occur in our model.

A much stronger notion of collusion resiliency is that of a *c-truthful* mechanism [13]. Essentially, such a mechanism guarantees that a group of at most c collusive players cannot “collectively gain more than they could by bidding individually.” This notion, however, is so strong that it has very limited applicability. The authors prove that the only *c-truthful* mechanisms must work in a specific manner: for each subset S of the goods and for each player i these mechanisms must fix a price $p_{S,i}$ and offer S to i for that price. We note that without any special knowledge about the players, one cannot pick these prices so as to offer any revenue guarantee. The authors also put forward a weaker variant of their notion — *c-truthful with high probability*— for which they can approximate maximum revenue, but only for a very special type of auction: unlimited supply of a single good.

Other revenue mechanisms, sharing a similar algorithmic approach, have been developed for various restricted combinatorial auctions [5, 17, 20]. Let us sketch a specific incarnation of this approach for auctions of multiple goods in the unlimited supply model:

Let L (for “lowerbound”) and U (for “upperbound”) be such that, for any subset S of distinct goods and any player i , either $TV_i(S) = 0$ or $L \leq TV_i(S) \leq U$. Accordingly, the mechanism (1) randomly selects a power of 2 between L and U , without loss of generality 2^k , and then (2) offers any subset of the goods for fixed price 2^k to any player who wants it.³

³Such mechanisms typically achieve a fraction of roughly $\frac{1}{\log(U/L)}$ of the appropriate benchmark, essentially because, of the $\log(U/L)$ possible prices the mechanism might choose, one of them is the “right guess” (while the other possible prices typically perform poorly). Such performance, however, depends on prior knowledge of U and L . To avoid this requirement, one may consider the techniques seen in [2] or [5], which achieve (in our notation) a $O(\frac{1}{\log(U/L) \log^2 \log(U/L)})$ fraction of the appropriate benchmark, without knowing U, L in advance.

Notice that, even without explicitly discussing collusion, the above approach is collusion-resilient: there is little for collusive players to do when every bundle of goods is offered at the same take-it-or-leave-it price. Notice too that this style of mechanism is not directly applicable to **truly** combinatorial auctions. The key issue is: given prices for various sets of items, how (in a dominant-strategy truthful manner) can we choose which of the eligible bidders gets which of the possibly overlapping subsets? The response of [5, 17, 20] is (in essence) to work in the unlimited supply model, assuming that the auctioneer may freely duplicate the goods as necessary so that every eligible bidder receives (a copy of) every good he can pay for.

A natural alternative to picking a high benchmark, and achieving a log fraction of it, is to pick a lower benchmark and find mechanisms that achieve it to within constant factor. Both [10] and [14] take this approach, considering auctions of a single good in unlimited supply relative to the following benchmark: *the maximum revenue that can be generated by fixing a price p lower than the second highest player’s value for a copy of the good, and then offering a copy to any player willing to pay p for it.* Goldberg et al. ([14]) achieve a $\frac{1}{7600}$ fraction of this benchmark; this is improved to a $\frac{1}{4}$ approximation by Fiat et al. [10]. The mechanisms both papers propose are inspired by a classic machine learning technique: split the data into halves via random sampling, and use the first half to learn about the second half. Applied to auctions, the halves are different bidders, and the auction mechanism computes, based on one half, the optimal pricing structure to charge the other half.

This machine learning framework is adapted to the combinatorial auctions setting in [4], but again restricted to the unlimited supply case. Balcan et al. take the further step of taking computational considerations into account, using the same random sampling framework to show that if there exists a computationally efficient algorithm A that β approximates any of a class of a certain class of revenue benchmarks (omniscient optimization over a bounded-size class of mechanisms) then there exists a computationally efficient DST mechanism which $\beta(1 + \epsilon)$ approximates the same benchmark.

In short, the prior literature relates either to restricted models of collusion, or to auctions of a restricted type. Prior techniques, however, have been relevant to us. In particular, as we shall discuss later on, to obtain our positive result (Theorem 1) we shall refine the above-mentioned idea of exponentially-distributed prices and then put it to our own use. As for our impossibility results (Theorems 2 and 3), perhaps the closest analogs may be found in [14]. We note, however, that the impossibility results of Goldberg et al. hold only for restricted classes of truthful mechanisms (their version of Theorem 2 holds only for “ex-post truthful” mechanisms, and their version of Theorem 3 holds only for “symmetric” mechanisms.)

3 Our Benchmark

Our definition of resilient mechanisms, Definition 2, sets the bar much higher than Definition 1: to resiliently achieve a benchmark \mathbb{B} we must “defeat a longer sequence of universal quantifiers.” From the next lemma, however, we see that proving results about Definition 2 may be no harder than proving results about Definition 1, provided we work with benchmarks satisfying a certain natural property:

Definition 3. *A benchmark \mathbb{B} is player monotone if $\mathbb{B}(S) \geq \mathbb{B}(T)$ for all valuation sub-profiles T of S .⁴*

The attractiveness of adopting a player monotone benchmark is shown by the following trivial lemma, stating that, relative to a player-monotone benchmark, “Definition 1 implies Definition 2.”

Lemma 1. *If a DST mechanism M achieves a fraction c of a player-monotone benchmark \mathbb{B} , then it resiliently achieves the same fraction c of \mathbb{B} .*

Proof. For any true-valuation profile TV , independent player subset I , and bid sub-profile BID_{-I} , the (expected) revenue of $M(TV_I \sqcup BID_{-I})$ is by hypothesis at least $c \cdot \mathbb{B}(TV_I \sqcup BID_{-I})$, and thus by \mathbb{B} ’s player monotonicity at least $c \cdot \mathbb{B}(TV_I)$. *Q.E.D.*

⁴Notice that the *player-monotonicity* of a benchmark \mathbb{B} does not imply the *revenue monotonicity* (as defined by [24]) of a mechanism guaranteeing a given fraction of \mathbb{B} .

In other words, adopting a player-monotone benchmark \mathbb{B} provides a plausible way forward to proving that a mechanism M resiliently achieves a fraction c of \mathbb{B} in our very general collusive model: namely, proving that M is DST and achieves a fraction c of \mathbb{B} for all possible bid profiles. Accordingly, in light of the leverage that Lemma 1 affords us, we are resolved to seek a player-monotone benchmark. The obvious temptation is MSW : after all, the maximum social welfare upperbounds the revenue achievable by any DST auction mechanism. The problem, however, is that no DST mechanism can achieve a positive fraction of MSW .⁵

Our benchmark consists of the maximum social welfare, after disregarding the “star player,” that is the player with the highest valuation for some subset of the goods. Let us start by defining the latter concept.

Definition 4. *Given a valuation subprofile V for a set of goods G , we say that player i is the star player (denoted by “ \star ”) if there exists $S \subset G$ such that, for any player j and any $T \subset G$, $TV_i(S) \geq TV_j(T)$.*

Definition 5. *We define the benchmark $MSW_{-\star}$ as follows: for any valuation profile V ,*

$$MSW_{-\star}(V) = MSW(V_{-\star}).$$

That is, $MSW_{-\star}$ is computed by first removing the valuation of the star player, and then computing the maximum social welfare of the remaining valuations. In other words: if $V = (V_1, \dots, V_{\star-1}, V_{\star}, V_{\star+1}, \dots, V_n)$, then $MSW_{-\star}(V) = MSW(V_1, \dots, V_{\star-1}, V_{\star+1}, \dots, V_n)$.

Notice that $MSW_{-\star}$ indeed is player-monotone: the introduction of new players will only grow the set of “non-star players” and thus the MSW of the “non-star players” may only grow. We point out two more arguments that make $MSW_{-\star}$ uniquely natural. (See Section 5 – after we have stated our main results – for a comparison between $MSW_{-\star}$ and other potential benchmarks.)

1. *Natural Generalization.* In auctions with a single good g for sale, $MSW_{-\star}$ coincides with the second-highest valuation, the revenue benchmark achieved by the second-price mechanism. Indeed, in such simple auctions, a player’s valuation is a single number (the value that the player has for g), and the star player is the one having the highest valuation for g . Therefore, after the star player is removed, the maximum social welfare of the remaining players is just the highest of the remaining valuations, and thus the second-highest of the original valuations. This analogy enables us to translate many reasons in favor of the second-highest valuation in single-good auctions to reasons in favor of $MSW_{-\star}$ in truly combinatorial auctions. In particular, $MSW_{-\star}$ is “reasonably high.”
2. *The “DST unreachability” of the star player.* Because the maximum social welfare function can be evaluated also on single valuations, the following is an alternative definition of the star player: $\star = \arg \max_i MSW(V_i)$.⁶ Unfortunately, no DST mechanism M can, on input a bid profile V , guarantee revenue that is a positive fraction of $MSW(V_{\star})$, the social welfare of the star player.⁷ It is thus quite natural to disregard from consideration the star player when setting revenue goals for DST mechanisms. In a sense, by choosing our benchmark to be $MSW_{-\star}$ we aim at “everything but the impossible.”

4 Our Theorem Statements

We prove three theorems. Stating the first two theorems requires defining three quantities, μ , c_μ and H_μ , where both c_μ and H_μ are roughly equal to $\log \mu$. Namely,

⁵That is, for any DST mechanism M , any truly combinatorial auction with n players and m goods, and positive constant $g_{n,m}$, there exists a valuation profile V such that the revenue of $M(V)$ is less than $g_{m,n} \cdot MSW(V)$. See [14] for a simple proof.

⁶Recall that our combinatorial auctions are not restricted to free disposal. Accordingly, $MSW(V_i)$ need not coincide with $V_i(G)$. Rather, $MSW(V_i) = V_i(S_i)$, where S_i is i ’s favorite subset of the goods: that is, $V_i(S_i) \geq V_i(T)$ for all $T \subset G$.

⁷This is for reasons analogous to the case of single-good auctions. Indeed, assume that the star player has an astronomical valuation for some subset of the goods, way out of scale with anyone else’s valuation of any other goods. Then, because a DST mechanism M cannot charge the star player based on his bid, and because it “cannot charge anyone else more than their bids,” M cannot produce revenue that is any fixed fraction of the star player’s social welfare.

- We consistently denote by μ the minimum of n and m , the number of players and the number of goods.
- For all μ , c_μ is the constant > 2 solving the equation $\frac{e^{x-2}}{x} = \mu$.
- For all μ , the μ th harmonic number, H_μ , is $\sum_{i=1}^\mu 1/i$.

Note that c_μ is uniquely defined. Indeed, for $\mu \geq 1$ the function $f_{n,m}(x) = e^{x-2} - x\mu$ is negative at $x = 2$, goes to infinity with x , and has positive second derivative everywhere.

Theorem 1. (Resilient-Revenue Theorem) *There exists a probabilistic DST mechanism \mathbb{M} that resiliently achieves a fraction $\frac{1}{c_\mu}$ of the revenue benchmark $\text{MSW}_{-\star}$.*

Theorem 2. (Harmonic Revenue Bound) *There is no DST mechanism, probabilistic or not, achieving a fraction $\frac{1}{H_{\mu-1}}$ of the revenue benchmark $\text{MSW}_{-\star}$.*

Taken together, Theorems 1 and 2 show that the revenue performance of mechanism \mathbb{M} is actually optimal within a small constant factor. (In addition, when μ goes to infinity $\frac{c_\mu}{H_{\mu-1}}$ goes to 1. That is, \mathbb{M} 's performance is asymptotically optimal.)

Mechanism \mathbb{M} works for *truly* combinatorial auctions, and its performance should not be confused with that of the mechanisms in the cited literature, which only apply to *restricted* combinatorial auctions.

Since in auctions of a single good, the second-price mechanism resiliently achieves revenue equal to 100% of $\text{MSW}_{-\star}$, perhaps one might have hoped for a constant-factor revenue guarantee in truly combinatorial auctions. From this viewpoint, proving that a logarithmic fraction of $\text{MSW}_{-\star}$ is the best one can guarantee for truly combinatorial auctions may be interpreted as a sad truth. On the other hand, however, we must recall that up to now the best revenue guarantee in such auctions was 0. In addition, our logarithmic guarantee holds in a most adversarial collusive model. So our results could and perhaps should be interpreted in a more optimistic manner.

Finally we prove that mechanism \mathbb{M} is necessarily probabilistic, as any *deterministic* DST mechanism perform exponentially worse, at least in some auctions.

Theorem 3. *There is no deterministic DST mechanism achieving a fraction $\frac{1}{\mu-1}$ of the revenue benchmark $\text{MSW}_{-\star}$.*

Again, we stress that our results do not rely on any restrictions.

Our Results Relative To Other Benchmarks. One may wonder what would happen to our results, had we chosen a different benchmark. For example, although we had valid reasons for dismissing the star player, one might be tempted to consider the following three benchmarks that remove many more players (1) “MSW minus the first k star players”; (2) “randomly partition the players into two halves, and take the smallest of the two MSWs”; and (3) “second-tier MSW: dismiss all the players that win goods in the MSW allocation, and take the maximum social welfare of the remaining players.” (The first alternative benchmark is inspired by our own, the second one by the mechanism of [4]; and the third is another generalization of the revenue of the second-price mechanism.) In another approach, one may consider bid-removing benchmarks such as “MSW after removing the highest bid for each subset of the goods.” All of these approaches lower the benchmark below $\text{MSW}_{-\star}$, but as it turns out, none enables a better than logarithmic performance: Our Theorems 1 and 2 remain unchanged relative to these other benchmarks (up to $o(1)$ constants). In yet another approach, one might consider benchmarks that remove items. For instance, the Louvre might be happy benchmarking the value of its collection as the value of all its artwork except the Mona Lisa. Unfortunately no positive fraction of the “MSW minus the star item” is achievable by DST mechanisms, for much the same reason that no positive fraction of MSW itself is DST-achievable. In sum, none of the above choices is advantageous from a “worst-case” perspective. In choosing $\text{MSW}_{-\star}$ we made a honest effort to find a high and reasonable benchmark.

5 Proof of Theorem 1

We use the balance of the space to prove Theorem 2 in full, and thus present just the mechanism of Theorem 1 (confining its analysis, as well as the proof of Theorem 3, to the appendix).

We obtain our mechanism \mathbb{M} by modifying the VCG so as convert some of its efficiency to revenue. We do so in three stages. First, we run the VCG on the profile of bids provided by the players and obtain an allocation A' and a profile of prices P' . Second, we raise all prices in P' by an amount ρ . Finally, we decide the final allocation and prices as follows. If player i wins a set of goods S in A' , and if the raised price $P'_i + \rho$ is less than i 's bid for S , then we finally allocate S to i for a price of $P'_i + \rho$. Else, S will go unallocated, and i pays 0.

But: How should we choose ρ to guarantee more revenue? In light of Theorem 3, we choose ρ *probabilistically*; and in light of our revenue goals we are tempted to choose ρ as a fraction α of $\text{MSW}_{-\star}(BID)$, where the scaling factor α is probabilistically chosen between 0 and 1. However, to ensure that \mathbb{M} is DST, we cannot have player i 's price depend on his bid, and $\text{MSW}_{-\star}(BID)$ may indeed depend on BID_i . We thus raise each price P'_i by $\alpha \text{MSW}(BID_{-i})$ instead. Our analysis will support that this small change does not alter the ability to achieve the chosen benchmark.

As we are shooting for a logarithmic fraction of $\text{MSW}_{-\star}$, we choose the scaling factor α according to an exponential distribution. We differ from the previous literature by adopting a *continuous* (as opposed to discrete) exponential distribution; this makes our analysis both tighter and cleaner. Our specific selection of constants is solely justified by our desire to optimize our revenue guarantee.

Mechanism \mathbb{M}

On input BID , a profile of n bids for a set of m goods, compute an outcome (A, P) as follows:

1. Pick a *scaling factor* $\alpha \in [0, 1]$ as follows. Flip a coin whose probability of Heads is $\frac{1}{c_\mu - 1}$. If Heads, choose $\alpha = 0$. If Tails, draw r uniformly from $[-(c_\mu - 2), 0]$ and choose $\alpha = e^r$.
2. Compute the *provisional allocation* A' and the profile of *provisional prices* P' , respectively the allocation and the prices of the VCG mechanism for the bid profile BID , and then the set of *provisional winners* W' consisting of all players that obtain a non-empty subset of goods in A' .
3. For each $i \in W'$ compute i 's *offer price* $P'_i + \alpha \text{MSW}(BID_{-i})$. If i 's bid $BID_i(A'_i)$ exceeds i 's offer price, set $A_i = A'_i$ and $P_i = P'_i + \alpha \text{MSW}(BID_{-i})$; otherwise set $A_i = \emptyset$ and $P_i = 0$.

Notice that although each price P_i is personalized, it is obtained via the same choice of scaling factor α . Were we in a Bayesian setting, where different players have different distributions for their valuations, then we would optimally choose a separate scaling factor α_i for each player i .

6 Proof of Theorem 2

Since there are infinitely many DST mechanisms M , one cannot construct an *ad hoc* “unprofitable” bid profile for each M . In principle, however, one could prove the Harmonic Revenue Bound by exhibiting a single and uniform algorithm that, on inputs n , m , and M , outputs the desired $BID^{n,m,M}$, and argue —again in a uniform way— that the expected revenue of $M(BID^{n,m,M})$ is at most $\text{MSW}_{-\star}(BID^{n,m,M}) / (H_\mu - 1)$. Having tried this approach for a while, we do not recommend it. We thus try a different and non-constructive approach, based on the probabilistic method.⁸

First, in a uniform way, we specify for any n and m a distribution $\mathcal{BID}^{n,m}$ over the bid profiles for combinatorial auctions with n players and m goods. Then, as promised, we prove that, for all n , m and

⁸When, as in the case of Theorem 1, we need to prove that there exists a mechanism \mathbb{M} enjoying some useful properties, explicitly finding \mathbb{M} is very desirable; but for proving that “no good mechanism exists” constructiveness is not necessary.

DST mechanisms M (probabilistic or deterministic), the ratio of the expected revenue of $M(\mathcal{BID}^{n,m})$ and the expected value of $\text{MSW}_{-\star}(\mathcal{BID}^{n,m})$ is at most $\frac{1}{H_{\mu-1}}$. Because this could not happen if M 's revenue were greater than $\frac{\text{MSW}_{-\star}(\mathcal{BID})}{H_{\mu-1}}$ for all bid profiles BID in the support of $\mathcal{BID}^{n,m}$, the existence of an unprofitable bid profile is established.

6.1 Implementing the Plan

We define the distribution $\mathcal{BID}^{n,m}$ in two steps. We start by defining a distribution, h_S^k , over single-minded bids of a single player.

Definition 6. (Bounded-Harmonic Distributions) *For any subset of goods S and positive integer k , we denote by h_S^k the distribution assigning, for each integer $j \in [1, k]$, probability $\frac{1}{k}$ to the single-minded valuation where set S (and only set S) is valued for $\frac{1}{j}$. We denote this valuation as $(S, \frac{1}{j})$.*

We now prove a very general property of DST mechanisms when a player bids according to h_S^k .

Lemma 2. (Harmonic-Pricing) *For all probabilistic DST mechanisms M , all players i , all valuation sub-profiles BID_{-i} , all positive integers k , and all subsets of goods S ,*

$$E_{BID_{-i} \leftarrow h_S^k} [E[M_p(BID_{-i} \sqcup BID_i)_i]] \leq \frac{1}{k}.$$

Proof. For each $j \leq k$ define α_j as the expected price paid by player i relative to the bid profile $BID_{-i} \sqcup (S, \frac{1}{j})$ under mechanism M , and define $\alpha_{k+1} = 0$. Thus our lemma can be restated as follows:

$$\frac{1}{k} \sum_{j=1}^k \alpha_j \leq \frac{1}{k} \quad \text{or, equivalently,} \quad \sum_{j=1}^k \alpha_j \leq 1.$$

Now suppose for the sake of contradiction that $\sum_{j=1}^k \alpha_j > 1$. Then

$$1 < \sum_{j=1}^k \alpha_j = \sum_{j=1}^k j(\alpha_j - \alpha_{j+1}). \quad (1)$$

For each $j \leq k$ define now β_j to be the probability that player i is allocated a set containing S by mechanism M when the bid profile is $BID_{-i} \sqcup (S, \frac{1}{j})$, and define $\beta_{k+1} = 0$. Then we have

$$\sum_{j=1}^k (\beta_j - \beta_{j+1}) = \beta_1 \leq 1. \quad (2)$$

Comparing the summations of inequalities 1 and 2 term by term, we conclude that there must exist a j such that the corresponding term from the first sum exceeds the term from the second sum, namely $j(\alpha_j - \alpha_{j+1}) > (\beta_j - \beta_{j+1})$. Dividing by j and rearranging terms yields

$$\frac{1}{j} \beta_j - \alpha_j < \frac{1}{j} \beta_{j+1} - \alpha_{j+1}. \quad (3)$$

Assume now that i 's true valuation is $(S, \frac{1}{j})$. Thus, since M is DST, when all other players bid according to the sub-profile BID_{-i} , i 's utility is at least as large when he bids $(S, \frac{1}{j})$ as when he bids $(S, \frac{1}{j+1})$. That is,

$$\frac{1}{j} \beta_j - \alpha_j \geq \frac{1}{j} \beta_{j+1} - \alpha_{j+1}. \quad (4)$$

Inequality 4 contradicts Inequality 3. Thus $\sum \alpha_j \leq 1$, as desired. Q.E.D.

To construct $\mathcal{BID}^{n,m}$ we “force $n = m = \mu$ and have μ separate auctions,” each with a single good and a single player. Indeed, we let each player i ’s bid to be distributed according to $h_{\{i\}}^\mu$: that is, i bids only for the subset $\{i\}$, and the amount of his bid is the inverse of an integer uniformly and independently chosen in $\{1, \dots, \mu\}$.

Definition 7. *In a combinatorial auction with n players and m goods, denoting the set of goods by $\{1, \dots, m\}$ and letting $\mu = \min\{n, m\}$, we define the distribution $\mathcal{BID}^{n,m}$ as follows. For each player i : if $i \in \{1, \dots, \mu\}$, then $\mathcal{BID}_i^{n,m} = h_{\{i\}}^\mu$; else $\mathcal{BID}_i^{n,m}$ is the null valuation.*

Finally, let us restate and prove the Harmonic Revenue Bound.

Theorem 2’. *For any $n, m > 1$, and any DST mechanism M , there exists a valuation profile BID for a truly combinatorial auction with n players and m goods such that, letting $\mu = \min\{n, m\}$ we have*

$$E \left[\sum_i M_p(BID)_i \right] \leq \frac{\text{MSW}_{-\star}(BID)}{H_\mu - 1}.$$

Proof. Let $\mathcal{BID}^{n,m}$ be the distribution of Definition 7 and fix arbitrarily a DST mechanism M . Invoking μ times (i.e., for each player $\leq \mu$) Lemma 2 with $k = \mu$ we have

$$E \left[\sum_i M_p(\mathcal{BID}^{n,m})_i \right] \leq 1. \tag{5}$$

That is, M ’s expected revenue (over $\mathcal{BID}^{n,m}$ and M ’s random choices, if any) is ≤ 1 . At the same time,

$$E [\text{MSW}_{-\star}(\mathcal{BID}^{n,m})] \leq H_\mu - 1. \tag{6}$$

In fact,

- (a) the expected value of MSW over $\mathcal{BID}^{n,m}$ is just $\sum_{j=1}^\mu 1/j = H_\mu$;
- (b) there are at least two players by hypothesis; and
- (c) the star player —whoever he may be— values his item for at most 1.

Inequalities 5 and 6 thus imply that the ratio between M ’s expected revenue and the expected $\text{MSW}_{-\star}$ is at most $\frac{1}{H_\mu - 1}$. In turn, this implies the existence of a bid profile BID as per our thesis. *Q.E.D.*

Acknowledgements

We are grateful to Daron Acemoglu and Sergei Izmalkov for several suggestions.

References

- [1] Ausubel, L.M. and Milgrom, P. The Lovely but Lonely Vickrey Auction. *Combinatorial Auctions*, pp. 17-40, 2006.
- [2] Awerbuch, B. Azar, Y. and Meyerson, A. Reducing Truth-telling Online Mechanisms to Online Optimization. *Symposium on Theory of Computing* 2003.
- [3] Babaioff, M., Lavi, R., and Pavlov, E. Single-Value Combinatorial Auctions and Implementation in Undominated Strategies. *Symposium on Discrete Algorithms*, pp. 1054-1063, 2006.
- [4] Balcan, M.-F., Blum, A., Hartline, H., and Mansour, Y. Mechanism Design via Machine Learning. *Foundations of Computer Science*, 2005.

- [5] Balcan, M.-F., Blum, A., and Mansour, Y. Single Price Mechanisms for Revenue Maximization in Unlimited Supply Combinatorial Auctions. *CMU-CS-07-111*, 2007.
- [6] Chung, K.-S., and Ely, J.C. Ex-Post Incentive Compatible Mechanism Design. mimeo. 2003.
- [7] Clarke, E.H. Multipart Pricing of Public Goods. *Public Choice* 11:17-33, 1971.
- [8] 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.
- [9] Cramton, P. and Schwartz, J. Collusion in Auctions. *Annales d'Economie et de Statistique*, 15/16:217-230, 1989.
- [10] A. Fiat, A. Goldberg, J. Hartline, A. Karlin. Competitive Generalized Auctions. *Symposium on Theory of Computing* 2002.
- [11] Feigenbaum, J., Papadimitriou, C., and Shenker, S. Sharing the Cost of Multicast Transmissions. *Symposium on Theory of Computing*, pp. 218-226, 2000.
- [12] Friedman, M. Comment on 'Collusion in the Auction Market for Treasury Bills'. *J. of Political Economy* 9:757-785, 1996.
- [13] Goldberg, A. and Hartline, J. Collusion-Resistant Mechanisms for Single-Parameter Agents. *Symposium on Discrete Algorithms*, 2005.
- [14] Goldberg, A., Hartline, J., Karlin, A., Saks, M., and Wright, A. Competitive Auctions. *Games and Economic Behavior*, 2006.
- [15] 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.
- [16] Groves, T. Incentives in Teams. *Econometrica*, 41:617-631, 1973.
- [17] 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.
- [18] Jain, K. and Vazirani, V. Applications of Approximation Algorithms to Cooperative Games. *Symposium on Theory of Computing* 2001.
- [19] 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.
- [20] Likhodedov, A. and Sandholm, T. Approximating Revenue-Maximizing Combinatorial Auctions. *AAAI Conference on Artificial Intelligence*, pp. 267-274, 2005.
- [21] Moulin, H. and Shenker, S. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, 18:511-533, 2001.
- [22] Myerson, R. Optimal Auction Design. *Mathematics of Operations Research*, 6: 58-73, 1981.
- [23] Porter, R. Detecting Collusion. *Review of Industrial Organization*, 26:2:147-167, 2005.
- [24] Rastegari, B., Condon, A., and Leyton-Brown, K. Revenue Monotonicity in Combinatorial Auctions. *AAAI Conference on Artificial Intelligence*, 2007.
- [25] Vickrey, W. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8-37, 1961.

Appendix

A Preliminaries

We consistently denote the number of players by n and the number of goods by m .

An allocation is a sequence $A = A_0, A_1, \dots, A_n$, where A_i is the subset of goods allocated to player i , and A_0 the set of unallocated goods. The set of *winners* in an allocation A , denoted by Win_A , is the subsets of all players i such that A_i is non-empty. An outcome is a pair $\Omega = (A, P)$, where A is an allocation, and P a profile of prices (non-negative numbers). The utility function u_i of player i maps i 's true valuation and an outcome $\Omega = (A, P)$ to i 's utility as follows: $u_i(TV_i, \Omega) = TV_i(A_i) - P_i$.

A valuation of the goods is a function mapping each of the 2^m subsets of the goods to a real number, such that the empty subset is mapped to 0. If V_S and V_T are two valuation sub-profiles such that the subsets of the players S and T are disjoint, then by $V_S \sqcup V_T$ be denote the sub-profile mapping each player $i \in S \cup T$ to $(V_S)_i$ if $i \in S$, and to $(V_T)_i$ otherwise. A bid is a valuation of the goods.

A mechanism M is a (possibly probabilistic) function mapping a profile of bids BID to an outcome (A, P) satisfying the *opt-out condition*: $P_i = 0$ whenever BID_i is the null valuation. We view each mechanism M as two separate functions: an *allocation function* M_a and a *price function* M_p . That is, for all bid profiles BID : $M(BID) = (M_a(BID), M_p(BID))$. The *expected revenue* of mechanism M on bid profile BID is $E[\sum_{i \in N} M_p(BID)_i]$. We say that M is DST if for all players i and bid sub-profile BID_{-i} : $E[u_i(TV_i, M(TV_i \sqcup BID'_{-i}))] \geq E[u_i(TV_i, M(BID'))]$.

The social welfare, best-allocation, and maximum social welfare functions — SW , BA , and MSW — are so defined. For each valuation sub-profile V_C and allocation A :

- $SW(V_C, A) = \sum_{i \in C} V_i(A_i)$,
- $BA(V_C) = \operatorname{argmax}_{A \in \mathbb{A}(G)} SW(V_C, A)$, where $\mathbb{A}(G)$ denotes the set of all possible allocations of G , and
- $MSW(V_C) = SW(V_C, \operatorname{Best}A(V_C))$.

By convention, (1) *argmax*'s ties are broken lexicographically, and (2) $\operatorname{Best}A(V_C)_i \neq X$ for any subset of the goods X such that $V_i(X) = 0$.

A valuation v of a finite set of goods G is *single-minded* if there exists a single subset of goods S and $x \in \mathbb{R}^+$ such that $v(T) = x$ whenever $S \subset T$ and 0 otherwise. We compactly represent such a single-minded valuation v by the pair (S, x) .

Let us explicitly highlight two properties of DST mechanisms which we are going to use extensively. (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: \forall (probabilistic or not) DST mechanisms M , players i , and bid profile BID , we have:

$$0 \leq E[M_p(BID)_i] \leq E[BID_i(M_a(BID)_i)].$$

DST-2: \forall deterministic DST mechanisms M , players i , and bid profiles BID and BID' such that

$$BID_{-i} = BID'_{-i}, \text{ we have: } M_a(BID)_i = M_a(BID')_i \text{ implies } M_p(BID)_i = M_p(BID')_i.$$

B Proof of Theorem 1

Let us restate Theorem 1 using the notation of Appendix A.

Theorem 1': Mechanism M is DST and, \forall bid profiles BID in a truly combinatorial auction with n players and m goods,

$$E \left[\sum_{i \in N} M_p(BID)_i \right] \geq \frac{MSW_{-\star}(BID)}{c_\mu}. \tag{7}$$

Proof. The fact that \mathbb{M} is DST is due to the fact that it has been obtained from the VCG mechanism via modifications that are well known to preserve the VCG's dominant-strategy truthfulness. (See Appendix D for details.) Therefore, let us focus on \mathbb{M} 's revenue.

Recall that, for each player i , A'_i is the (possibly empty) set player i provisionally wins, and P'_i is the provisional price $VCG_p(BID)_i$. We divide our proof into two cases: in the first case the star player bids large enough that we derive the revenue bound solely based on the revenue \mathbb{M} extracts from the star player. In the second case we must sum up the revenue that \mathbb{M} extracts from each set-winning player.

Case 1: $BID_\star(A'_\star) > P'_\star + \text{MSW}_{-\star}(BID)$.

Note that the right-hand side of the inequality of this case is always ≥ 0 , thus $BID_\star(A'_\star) > 0$ always. This implies that $A'_\star \neq \emptyset$; namely that \star is a provisional winner. As such, when mechanism \mathbb{M} “makes \star the offer” $P'_\star + \alpha \cdot \text{MSW}_{-\star}(BID)$ where $\alpha \leq 1$, the offer price will always be at most player \star 's bid for A'_\star , and hence player \star will always pay his offer price. Thus the expected revenue from player \star is just the expected offer price, namely

$$\begin{aligned} E[\mathbb{M}_p(BID)_\star] &= \frac{1}{c_\mu - 1} P'_\star + \left(1 - \frac{1}{c_\mu - 1}\right) \int_{-(c_\mu - 2)}^0 \frac{1}{c_\mu - 2} (P'_\star + e^r \text{MSW}_{-\star}(BID)) dr \\ &= \left(\frac{1}{c_\mu - 1} + \left(1 - \frac{1}{c_\mu - 1}\right)\right) P'_\star + \left(1 - \frac{1}{c_\mu - 1}\right) \frac{1}{c_\mu - 2} \text{MSW}_{-\star}(BID) \int_{-(c_\mu - 2)}^0 e^r dr \\ &= P'_\star + \frac{1}{c_\mu - 1} \text{MSW}_{-\star}(BID) \int_{-(c_\mu - 2)}^0 e^r dr \\ &= P'_\star + \text{MSW}_{-\star}(BID) \frac{1 - e^{-(c_\mu - 2)}}{c_\mu - 1} \\ &\geq \text{MSW}_{-\star}(BID) \frac{1 - \mu e^{-(c_\mu - 2)}}{c_\mu - 1} = \text{MSW}_{-\star}(BID) \frac{1 - \frac{1}{c_\mu}}{c_\mu - 1} = \frac{\text{MSW}_{-\star}(BID)}{c_\mu}, \end{aligned}$$

where the inequality follows because $P'_\star \geq 0$ and $\mu \geq 1$, and the second to last equality is by the definition of c_μ , namely that $c_\mu \mu e^{-(c_\mu - 2)} = 1$. Thus we have the desired result in this case.

Case 2: $BID_\star(A'_\star) \leq P'_\star + \text{MSW}(BID_{-\star})$.

Consider a provisional winner i , and consider his offer price $P'_i + \alpha \cdot \text{MSW}(BID_{-i})$. We note that when $\alpha = 0$ the offer price for player i is just P'_i , which is less than or equal to $BID_i(A'_i)$ since the VCG mechanism never charges players more than their bid; thus when $\alpha = 0$ player i will “accept the offer” and pay P'_i . Since player i will pay the offer price whenever it is less than $BID_i(A'_i)$, we have that i will pay whenever $\alpha < \frac{BID_i(A'_i) - P'_i}{\text{MSW}(BID_{-i})}$. Recall that, by the definition of \mathbb{M} , when $\alpha \neq 0$ we have $\alpha = e^r$. Thus this condition becomes $r < \log_e \frac{BID_i(A'_i) - P'_i}{\text{MSW}(BID_{-i})}$. We note that r is also bounded to be at most 0, but the other condition takes precedence since $\log_e \frac{BID_i(A'_i) - P'_i}{\text{MSW}(BID_{-i})} \leq 0$, as we show by an analysis of two cases: when $i = \star$ the claim is equivalent to the condition of Case 2, $BID_\star(A'_\star) \leq P'_\star + \text{MSW}(BID_{-\star})$; otherwise, when $i \neq \star$ we have $BID_i(A'_i) - P'_i \leq BID_i(A'_i) \leq \text{MSW}(BID_\star) \leq \text{MSW}(BID_{-i})$, where $\text{MSW}(BID_\star)$ denotes the highest bid of the star player (which is higher than any other bid, including $BID_i(A'_i)$ by assumption), which implies the ≤ 0 bound we wanted to prove. Thus the expected price paid by player i is exactly expressed as the following integral:

$$\frac{P'_i}{c_\mu - 1} + \left(1 - \frac{1}{c_\mu - 1}\right) \int_{-(c_\mu - 2)}^{\max\left\{- (c_\mu - 2), \log_e \frac{BID_i(A'_i) - P'_i}{\text{MSW}(BID_{-i})}\right\}} \frac{1}{c_\mu - 2} (P'_i + e^r \text{MSW}(BID_{-i})) dr.$$

We lower-bound this expression using the following two observations: first, since $P'_i \geq 0$ we may remove this term from inside the integral; second, since the integrand is always positive, if we decrease the upper

limit of the integral to $\log_e \frac{BID_i(A'_i) - P'_i}{MSW(BID_{-i})}$ the integral can only decrease (where we use the standard convention that an integral with upper limit less than its lower limit is evaluated with limits reversed and negated). Thus we have

$$\begin{aligned} E [\mathbb{M}_p(BID)_i] &\geq \frac{P'_i}{c_\mu - 1} + \left(1 - \frac{1}{c_\mu - 1}\right) \int_{-(c_\mu - 2)}^{\log_e \frac{BID_i(A'_i) - P'_i}{MSW(BID_{-i})}} \frac{1}{c_\mu - 2} (e^r MSW(BID_{-i})) dr \\ &= \frac{P'_i}{c_\mu - 1} + MSW(BID_{-i}) \frac{1}{c_\mu - 1} \left(e^{\log_e \frac{BID_i(A'_i) - P'_i}{MSW(BID_{-i})}} - e^{-(c_\mu - 2)} \right) \\ &= \frac{1}{c_\mu - 1} \left(BID_i(A'_i) - e^{-(c_\mu - 2)} MSW(BID_{-i}) \right) \end{aligned}$$

Summing up this inequality over all provisional winners i , we get

$$E \left[\sum_{i \in W'} \mathbb{M}_p(BID)_i \right] \geq \frac{1}{c_\mu - 1} \left(\sum_{i \in W'} BID_i(A'_i) - e^{-(c_\mu - 2)} \sum_{i \in W'} MSW(BID_{-i}) \right).$$

Now notice that $\sum_{i \in W'} BID_i(A'_i) = MSW(BID)$. Further since $|W'| \leq \mu$ and $MSW(BID_{-i}) \leq MSW(BID)$ we have $\sum_{i \in W'} MSW(BID_{-i}) \leq \mu \cdot MSW(BID)$. Thus we have

$$\begin{aligned} E \left[\sum_{i \in W'} \mathbb{M}_p(BID)_i \right] &\geq MSW(BID) \frac{1 - \mu e^{-(c_\mu - 2)}}{c_\mu - 1} \\ &= MSW(BID) \frac{1 - \frac{1}{c_\mu}}{c_\mu - 1} = \frac{MSW(BID)}{c_\mu} \geq \frac{MSW_{-*}(BID)}{c_\mu}, \end{aligned}$$

where we invoke the definition of c_μ to derive the first equality. Thus we have the desired conclusion. *Q.E.D.*

Remarks.

- Notice that our mechanism \mathbb{M} requires that its underlying DST mechanism be reasonably efficient. Indeed, in the analysis of Case 2, we rely on the fact that the VCG algorithm is 100% efficient: namely, we rely $\sum_{i \in W'} BID_i(A'_i) = MSW(BID)$. If another DST mechanism is used, one should make sure that, for its provisional allocation A' , $\sum_i BID_i(A'_i)$ is a sufficient fraction of $MSW(BID)$.
- Notice that, when lower-bounding the revenue generated by \mathbb{M} , the profile of prices returned by the underlying DST mechanism are essentially ignored. However, were we to “simplify” the definition of \mathbb{M} by replacing the provisional prices with zeros, the resulting mechanism would not be DST.

C Proof of Theorem 3

Let us restate and then prove Theorem 4. Namely,

Theorem 4: *For any $n, m > 1$, and for any deterministic auction mechanism M , there exists a valuation profile BID for truly combinatorial auctions with n players and m goods such, letting $\mu = \min\{n, m\}$, the revenue of $M(BID)$ is at most*

$$\frac{MSW_{-*}(BID)}{\mu - 1}.$$

Proof. We construct the desired bid profile within three steps.

Step 1. Define the single-minded valuation profile BID^0 as follows: BID_i^0 equals $(\{i\}, 1)$ if player $i \leq \mu$, and the null valuation otherwise. It is thus clear that $\text{MSW}_{-\star}(BID^0) = \mu - 1$, so that $\frac{\text{MSW}_{-\star}(BID^0)}{\mu - 1} = 1$. We distinguish two cases: namely, (1) $M_p(BID^0)_i > 0$ for no i and (2) $M_p(BID^0)_i > 0$ for some i . In the first case, the revenue of M on bid profile BID^0 is 0, and thus BID^0 is the profile required by the theorem. Otherwise, we proceed to Step 2.

Step 2. Let j be a player such that $M_p(BID^0)_j > 0$, and define for each integer $\alpha \geq 2$ the valuation profile $BID^\alpha = BID_{-j}^0 \sqcup (\{j\}, \mu^\alpha)$. It is thus evident that, for all $\alpha \geq 2$, $\text{MSW}_{-\star}(BID^\alpha) = \mu - 1$ and thus $\frac{\text{MSW}_{-\star}(BID^\alpha)}{\mu - 1} = 1$.

Let us now analyze the price side. Notice three facts: by construction, $BID_{-j}^0 = BID_{-j}^\alpha$; by Property DST-1, j is allocated the subset of goods $\{j\}$ under bid profile BID^0 ; and, for all $\alpha \geq 2$, j 's bid value for $\{j\}$ is higher in BID^α than in BID^0 . Thus, because M is deterministic, Property DST-2 implies that, for all $\alpha \geq 2$, j continues to win the set $\{j\}$ in BID^α and to pay the same price he pays in BID^0 , which is at most 1 —because of Property DST-1 and because $BID_j^0(\{j\}) = 1$.

We now distinguish two cases: (a) there is some integer $\bar{\alpha} \geq 2$ such that $M_p(BID^{\bar{\alpha}})_i = 0$ for all $i \neq j$, or (b) for each integer $\alpha \geq 2$ there is a player k_α , $k_\alpha \neq j$, such that $M_p(BID^\alpha)_{k_\alpha} > 0$. In the first case, the total revenue for M under bid profile $BID^{\bar{\alpha}}$ is at most 1, and thus $BID^{\bar{\alpha}}$ is the profile required by the theorem. Otherwise, we proceed to Step 3.

Step 3. By the opt-out condition, for each integer $\alpha > 2$, we have $k_\alpha \in \{1, \dots, \mu\} \setminus \{j\}$. Thus, the pigeonhole principle implies the existence of $\beta, \gamma \in \{2, \dots, \mu + 1\}$ such that $k_\beta = k_\gamma$. Without loss of generality, let $\beta < \gamma$. Define now

$$k = k_\beta (= k_\gamma) \quad \text{and} \quad BID' = BID_{-k}^\gamma \sqcup (\{k\}, \mu^{\gamma+1}).$$

Since the star player in BID' is k , we have $\text{MSW}_{-\star}(BID') = \mu^\gamma + \mu - 2 \geq \mu^\gamma$. Further, because γ and β are integers, $\gamma > \beta$, and $\beta \geq 2$, we have $\mu^\gamma > (\mu^\beta + \mu)(\mu - 1)$ and thus

$$\mu^\beta + \mu < \frac{\text{MSW}_{-\star}(BID')}{\mu - 1}. \quad (8)$$

Let us now analyze the price situation. We consider the following two mutually exclusive cases.

Case 1: $M_p(BID')_j \leq \mu^\beta$.

The definition of BID' and Property DST-1 clearly imply that $\sum_{i \in -\{j, k\}} M_p(BID')_i \leq \mu - 2$. As for player k , we note that under bid profile BID^γ , he wins his set $\{k\}$, without paying more than 1, his bid for $\{k\}$. Further, the bid profile BID' is identical to BID^γ except for k 's bid for $\{k\}$, which is higher in BID' than in BID^γ . Thus, k will continue to win $\{k\}$ in BID' , without paying more than 1. Thus the revenue in this case is at most the sum of μ^β (from player j), 1 (from player k), and $\mu - 2$ (from all other players), totalling less than $\mu^\beta + \mu$.

Thus Inequality 8 implies that the valuation profile BID' satisfies our thesis.

Case 2: $M_p(BID')_j > \mu^\beta$.

Define

$$BID'' = BID_{-k}^\beta \sqcup (\{k\}, \mu^{\gamma+1}).$$

It is clear that $\text{MSW}_{-\star}(BID'') = \mu^\beta + \mu - 2 \geq \mu^\beta$ and thus, since $\beta \geq 2$, we have

$$\mu < \frac{\text{MSW}_{-\star}(BID'')}{\mu - 1}. \quad (9)$$

Turning our attention to prices, as for BID' , it is clear that $\sum_{i \in -\{j, k\}} M_p(BID'')_i \leq \mu - 2$.

Let us now analyze the price paid by player j . Notice that BID' and BID'' differ only in the bid of player j , and that j bids higher for $\{j\}$ in BID' than in BID'' . Thus, if j won $\{j\}$ in BID'' , then he would

win it too in BID' for the same price. However, by the assumption of this case, j 's price is greater than μ^β , namely, greater than his valuation for $\{j\}$ under BID'' , which implies that he cannot win it under bid profile BID'' .

Finally, let us analyze the price of player k . Notice that BID'' is identical to BID^β except for k 's bid for $\{k\}$, which is higher in BID^β than in BID'' . Since k wins his set $\{k\}$ under BID^β paying at most 1, he continues to win $\{k\}$ in BID'' , for at most 1.

Thus the revenue in this case is at most the sum of 0 (from player j), 1 (from player k), and $\mu - 2$ (from all other players), totalling less than μ .

Thus Inequality 9 implies that the valuation profile BID'' satisfies our thesis.

And thus, in all cases, we have exhibited a valuation profile that satisfies the theorem.

Q.E.D.

D DST-Preserving Modifications of DST Mechanisms

Definition 8. A function $f : N \times \mathbb{V}(G)^N \rightarrow \mathbb{R}^+$ is called stable if, for each $i \in N$, there exists a function g_i such that, for each $BID \in \mathbb{V}(G)^N$, $f(i, BID) = g_i(BID_{-i})$ —i.e., f is “independent” of BID_i .

Definition 9. Given a deterministic mechanism \mathcal{M} and a stable function f define \mathcal{M}^{+f} to be the mechanism defined as follows: on input $BID \in \mathbb{V}(G)^N$,

1. Compute the provisional allocation $A' = \mathcal{M}_a(BID)$, the profile of provisional prices $P' = \mathcal{M}_p(BID)$, and the set of provisional winners $W' = Wins(A')$.
2. For each $i \in W'$, if $BID_i(A'_i) \geq P'_i + f(i, BID)$ then let $P_i = P'_i + f(i, BID)$ and $A_i = A'_i$; otherwise let $P_i = 0$ and $A_i = \emptyset$.

Lemma 3. If \mathcal{M} is DST, so is \mathcal{M}^{+f} for any stable f .

Proof. Given an auction context $\mathcal{C} = (N, G, TV)$, a bid profile $BID \in \mathbb{V}(G)^N$, and a player $i \in N$ we have

$$\begin{aligned}
u_i(TV_i, \mathcal{M}^{+f}(TV_i \sqcup BID_{-i})) &\stackrel{(1)}{=} TV_i(\mathcal{M}_a^{+f}(TV_i \sqcup BID_{-i})_i) - \mathcal{M}_p^{+f}(TV_i \sqcup BID_{-i})_i \\
&\stackrel{(2)}{=} \max\{0, TV_i(\mathcal{M}_a(TV_i \sqcup BID_{-i})_i) - \mathcal{M}_p(TV_i \sqcup BID_{-i})_i - f(i, BID_{-i})\} \\
&\stackrel{(3)}{=} \max\{0, u_i(TV_i, \mathcal{M}(TV_i \sqcup BID_{-i})) - f(i, BID_{-i})\} \\
&\stackrel{(4)}{\geq} \max\{0, u_i(TV_i, \mathcal{M}(BID)) - f(i, BID_{-i})\} \\
&\stackrel{(5)}{\geq} u_i(TV_i, \mathcal{M}^{+f}(BID)),
\end{aligned}$$

where: Equality (1) holds by the definition of the utility function u ; equality (2) holds by the definition of \mathcal{M}^{+f} , and can be easily checked in each of the two cases $\mathcal{M}_a^{+f}(TV_i \sqcup BID_{-i})_i = \mathcal{M}_a(TV_i \sqcup BID_{-i})_i$ and $\mathcal{M}_a^{+f}(TV_i \sqcup BID_{-i})_i = \emptyset$; equality (3) holds by the definition of the utility function u ; inequality (4) holds because \mathcal{M} is dominant-strategy truthful; and inequality (5) holds because $\mathcal{M}^{+f}(BID)$ either yields i utility 0 or utility $u_i(TV_i, \mathcal{M}(BID)) - f(i, BID_{-i})$, and thus at most the maximum of these two quantities. *Q.E.D.*

We now define a special class of probabilistic mechanisms.

Definition 10. Let \mathcal{D} a distribution over a set of deterministic $n \times m$ auction mechanisms. Then we denote by $\mathbb{S}_{\mathcal{D}}$ the probabilistic $n \times m$ auction mechanism that, on input a profile of bids BID , first selects a mechanism \mathcal{M} according to \mathcal{D} and then returns the outcome $\mathcal{M}(BID)$.

We refer to such a mechanism $\mathbb{S}_{\mathcal{D}}$ as above as a $(n \times m)$ sampler.

Lemma 4. *If $\mathbb{S}_{\mathcal{D}}$ is a sampler and all mechanisms in \mathcal{D} 's support are DST, then $\mathbb{S}_{\mathcal{D}}$ is DST.*

Proof. Let our $\mathbb{S}_{\mathcal{D}}$ actually be an $n \times m$ sampler, and let $\mathcal{C} = (N, G, TV)$ be an $n \times m$ auction context, $BID \in \mathbb{V}(G)^n$, and $i \in N$. Then we have

$$\begin{aligned}
& E[u_i(TV_i, \mathbb{S}_{\mathcal{D}}(TV_i \sqcup BID_{-i}))] \stackrel{(1)}{=} \\
& \underset{\mathcal{M} \leftarrow \mathcal{D}}{E} [u_i(TV_i, \mathcal{M}(TV_i \sqcup BID_{-i}))] \stackrel{(2)}{\geq} \\
& \underset{\mathcal{M} \leftarrow \mathcal{D}}{E} [u_i(TV_i, \mathcal{M}(BID))] \stackrel{(3)}{=} \\
& E[u_i(TV_i, \mathbb{S}_{\mathcal{D}}(BID))],
\end{aligned}$$

where equality (1) holds because the definition of $\mathbb{S}_{\mathcal{D}}$; inequality (2) because \mathcal{M} is DST; and equality (3) because of the definition of $\mathbb{S}_{\mathcal{D}}$. This chain of inequality proves that $\mathbb{S}_{\mathcal{D}}$ is DST.

Q.E.D.