On the Hardness of Dominant Strategy Mechanism Design

Shahar Dobzinski
Weizmann Institute of Science
Israel
shahar.dobzinski@weizmann.ac.il

Shiri Ron
Weizmann Institute of Science
Israel
shiriron@weizmann.ac.il

Jan Vondrák
Stanford University
USA
jvondrak@stanford.edu

ABSTRACT

We study the communication complexity of dominant strategy implementations of combinatorial auctions. We start with two domains that are generally considered "easy": multi-unit auctions with decreasing marginal values and combinatorial auctions with gross substitutes valuations. For both domains we have fast algorithms that find the welfare-maximizing allocation with communication complexity that is poly-logarithmic in the input size. This immediately implies that welfare maximization can be achieved in ex-post equilibrium with no significant communication cost, by using VCG payments. In contrast, we show that in both domains the communication complexity of any dominant strategy implementation that achieves the optimal welfare is polynomial in the input size.

We then move on to studying the approximation ratios achievable by dominant strategy mechanisms. For multi-unit auctions with decreasing marginal values, we provide a dominant-strategy communication FPTAS. For combinatorial auctions with general valuations, we show that there is no dominant strategy mechanism that achieves an approximation ratio better than $m^{1-\epsilon}$ that uses $poly(m, n)$ bits of communication, where m is the number of items and n is the number of bidders. In contrast, a *randomized* dominant strategy mechanism that achieves an $O(\sqrt{m})$ approximation with $poly(m, n)$ communication is known. This proves the first gap between computationally efficient deterministic dominant strategy mechanisms and randomized ones.

En route, we answer an open question on the communication cost of implementing dominant strategy mechanisms for more than two players, and also solve some open problems in the area of simultaneous combinatorial auctions.

CCS CONCEPTS

• **Theory of computation → Communication complexity**.

KEYWORDS

Communication Complexity, Mechanism Design

ACM Reference Format:
Shahar Dobzinski, Shiri Ron, and Jan Vondrák. 2022. On the Hardness of Dominant Strategy Mechanism Design. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3519935.3520013

STOC '22, June 20–24, 2022, Rome, Italy

ACM ISBN 978-1-4503-9264-8/22/06…$15.00
https://doi.org/10.1145/3519935.3520013

1 INTRODUCTION

In his seminal 1961 paper [31], Vickrey considers single item auctions: there is one item and n bidders, the value of each bidder i for the item is v_i. Vickrey defines the second-price auction: the highest bidder wins the item and pays the second highest bid. It is shown that in a second-price auction, bidding truthfully is a dominant strategy for each bidder. However, observe that our definition of a second-price auction was a bit careless. Bidding truthfully is indeed a dominant strategy when the second price auction is held by asking the bidders to simultaneously submit their bids, or when implemented iteratively, by conducting a (continuous) ascending auction. However, this is not always the case. Consider a "serial" implementation of a second price auction in which the bids of players $1, \ldots, i-1$ are publicly revealed before player i makes a bid. Truth-telling is no longer a dominant strategy for, e.g., player 1: if the strategy of all other players is "bid 0 unless player 1 bids 10, in which case bid 9", then player 1 is better off bidding 11 when his true value is $v_1 = 10$.

The Setting. In this paper we analyze the hardness of dominant strategy implementations in combinatorial auctions. Recall that in a combinatorial auction there is a set M of heterogeneous items ($|M| = m$) and a set N of bidders ($|N| = n$). The private information of each bidder i is his value for every subset of the items: $v_i : 2^M \to \mathbb{R}$. The standard assumptions are that the valuations are non-decreasing (for each $S \subseteq T$, $v(T) \geq v(S)$) and normalized ($v_i(\emptyset) = 0$), though we will sometimes impose additional restrictions on the valuations. Our goal is to find an allocation of the items $(S_1, \ldots, S_n)$ that maximizes the social welfare: $\Sigma_i v_i(S_i)$.

We use communication protocols to model mechanisms. Specifically, we work in the blackboard model, so all messages sent are observable by all players. The input of each player i is his valuation v_i. As usual, the communication protocol is represented by a tree that dictates which players (simultaneously) speak at each node, and the identity of the next node given the messages. The leaves of the protocol specify the outcome: the allocation and payments. We assume that each player is interested in maximizing his profit: the value of his assignment minus his payment. We assume that all mechanisms are *normalized*, i.e. that a player that wins the empty bundle pays zero.

A strategy $\mathcal{S}_i$ of player i dictates (given the valuation v_i) which messages player i sends at each node. We say that $\mathcal{S}_i$ is *dominant* for player i if for every set of possible strategies $\mathcal{S}'_{-i}$ of the other players, every valuation profile $(v_1, \ldots, v_n)$ and every strategy $\mathcal{S}'_i$

of player i it holds that:

$$v_i(f_i(\mathcal{S}_i(v_i), \mathcal{S}'_{-i}(v_{-i}))) - p_i(\mathcal{S}_i(v_i), \mathcal{S}'_{-i}(v_{-i})) \geq v_i(f_i(\mathcal{S}'_i(v_i), \mathcal{S}'_{-i}(v_{-i}))) - p_i(\mathcal{S}'_i(v_i), \mathcal{S}'_{-i}(v_{-i}))$$

where $f_i(\mathcal{S}_i(v_i), \mathcal{S}'_{-i}(v_{-i}))$ and $p_i(\mathcal{S}_i(v_i), \mathcal{S}'_{-i}(v_{-i}))$ specify the allocation and payment of player i, respectively, given that player i follows the actions specified by $\mathcal{S}_i(v_i)$ and the other players follow the actions specified in the vector $\mathcal{S}'_{-i}(v_{-i}) = (\mathcal{S}'_1(v_1), \ldots, \mathcal{S}'_{i-1}(v_{i-1}), \mathcal{S}'_{i+1}(v_{i+1}), \ldots, \mathcal{S}'_n(v_n))$. Consider a mechanism in which each player i has a dominant strategy $\mathcal{S}_i$. Let V_i be the set of possible valuations of a player and let $\mathcal{A}$ is the set of all possible allocations. Let $f : V_1 \times \cdots \times V_n \to \mathcal{A}$ be the social choice function defined by $f(v_1, \ldots, v_n)_i = f_i(\mathcal{S}_i(v_i), \mathcal{S}_{-i}(v_{-i}))$. In this case we say that the mechanism implements f in dominant strategies.

The importance of the specifics of the implementation and how they affect the solution concept are well known. The notion of ex-post equilibrium, defined by Cremer and McLean [12], attempts in a sense to get around this by ignoring the specifics of the implementation. A function $f : V_1 \times \cdots \times V_n \to \mathcal{A}$ is implementable in *ex-post equilibrium* if there are functions $p_1, \ldots, p_n : V_1 \times \cdots \times V_n \to \mathbb{R}$ such that for every player i, valuations v_i and v'_i of player i, and valuations v_{-i} of the other players:

$$v_i(f_i(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f_i(v'_i, v_{-i})) - p_i(v'_i, v_{-i})$$

Roughly speaking, in an ex-post equilibrium none of the players regrets playing according to his true value, if the other players are playing according to their true values as well. This rules out "unreasonable" strategies like in the serial second price auction described above. An alternative description would be that an ex-post incentive compatible implementation of a function f is a communication protocol that computes a function f and the associated payments, where f can be implemented in dominant strategies. However, in this protocol the players might not have dominant strategies. Clearly, every dominant-strategy implementation is also an ex-post implementation. The other direction is not true, as the serial implementation of a second price auction demonstrates.[1] Thus the communication cost of ex-post implementations is potentially much smaller than the cost of dominant-strategy implementations.

The goal of this paper is to determine whether the communication cost of implementations in dominant strategy is significantly larger than the cost of ex-post implementations. Intuitively, one might suspect that dominant-strategy mechanisms require significantly more communication than ex-post mechanisms. However, prior research can offer only mixed evidence to support this. First, the revelation principle implies, in particular, that every ex-post implementable function f is also dominant strategy implementable (the implementation is simple: each player simultaneously reveals his valuation, and the outcome is determined accordingly). However, as was already observed by Conitzer and Sandholm [11], this naive implementation method might easily result in an exponential blow-up in the communication complexity. Yet, this method works well in domains in which the private information of the players can be succinctly described, e.g., single-parameter domains.

[1] Admittedly, in some algorithmic mechanism design papers that analyze the communication complexity of incentive compatible mechanisms the distinction between the two notions is less explicit than it should be.

Our interest is in the more complicated multi-parameter domains. Almost all known *deterministic* incentive compatible mechanisms[2] [18, 20, 23, 29] are maximal-in-range mechanisms. Moreover, in each of them each bidder sends in the first round his value of all the (polynomially many) bundles he might win. Hence these mechanisms are dominant strategy.

The evidence that implementation in ex-post equilibrium does not buy much computational power comparing to dominant-strategy implementation is more than anecdotal. In [14] it is shown – perhaps counter-intuitively – that *every* two player ex-post mechanism for combinatorial auctions in a rich enough domain (in particular, one that includes all XOS valuations) can be implemented in a dominant-strategy equilibrium with only a polynomial blow-up in the communication complexity.[3]

In contrast, there is evidence supporting the idea that ex-post mechanism design is significantly less costly, communication-wise. Very recently, [30] presented a carefully crafted setting in which there is a mechanism that implements a welfare maximizer in an ex-post equilibrium with c bits, but every dominant-strategy implementation requires $exp(c)$ bits.

Our Results. We begin our explorations by considering the result of [14] discussed above, that shows that every function f for two players in a "rich enough" auction domain that can be implemented in an ex-post equilibrium can also be implemented in dominant strategies with only a polynomial blow-up in the communication. The paper [14] leaves open the question of whether this result holds also for mechanisms with more than two players. We answer this question in the negative by showing that the equivalence in implementations is unique for two player mechanisms: there is a three-player social choice function for general valuations that has an ex-post implementation that uses only c bits, but $exp(c)$ bits are required for any dominant-strategy implementation. The proof can be found in the full version.

Next, in Section 3 we consider two auction domains that are largely considered "easy" in the algorithmic mechanism design literature: multi-unit auctions with decreasing marginal values and combinatorial auctions with gross substitutes valuations (see, e.g., the surveys [8, 27]). In multi-unit auctions with decreasing marginal values, the welfare maximizing solution can be found with $poly(n, \log m)$ communication, and in combinatorial auctions with gross substitutes valuations, the welfare maximizing solution can be found with $poly(n, m)$ communication [28]. We thus get that in both settings the function that outputs the welfare maximizing allocation is implementable with low communication (since VCG payments can be computed with only a polynomial blow up in the communication).

However, these results hold only in an ex-post equilibrium. In a sharp contrast, we show that an exponential blow up is required for dominant strategy implementations (again, in the blackboard model):

[2] The only exception is [7] which is the only known example of a mechanism that is not maximal in range and achieves the state of the art results in a well-studied domain.
[3] In [14] an analogous result is proved also for domains that include all submodular valuations, under certain constraints.

Theorem:

(1) The communication complexity of every normalized mechanism that finds a welfare-maximizing allocation for two players in dominant strategies in multi unit auctions when the valuations exhibit decreasing marginal values is $\Omega(m \cdot \log m)$. In contrast, there is a mechanism with communication complexity $poly(\log m)$ that finds the welfare-maximizing allocation in an ex-post equilibrium.
(2) The communication complexity of every normalized mechanism that finds a welfare-maximizing allocation even for two players in dominant strategies in combinatorial auctions with gross substitutes valuations is $exp(m)$. In contrast, there is a mechanism with communication complexity $poly(m)$ that finds the welfare-maximizing allocation in an ex-post equilibrium.

We note that these results echo the recent result of [30] which was the first to show that the communication cost of dominant-strategy implementations of welfare maximizers might be exponential comparing to the communication cost of ex-post implementations but in an artificial domain. In contrast, our results prove an exponential blow-up of welfare maximizers in well-studied auction domains.

Perhaps in contrast to the common perception, the theorem demonstrates that these domains are not "easy" from the point of view of dominant-strategy mechanism design. This immediately raises the question of whether we can have good approximations to the social welfare by low-communication dominant-strategy mechanisms. For multi-unit auctions, we answer this question in the affirmative:

Theorem: Let $\varepsilon > 0$. There is a dominant-strategy $(1+\varepsilon)$-approximation mechanism for multi-unit auctions with valuations that exhibit decreasing marginal values that makes $poly(n, \log m, \frac{1}{\varepsilon})$ value queries.

Whether one can get good approximation ratios for combinatorial auctions with gross substitutes valuations remains an open question. The maximal-in-range mechanism of [20] achieves an approximation ratio of $O(\sqrt{m})$ in dominant strategies for the much larger class of subadditive valuations. However, we do not even know whether dominant strategy *maximal-in-range* mechanisms with polynomial communication can achieve a better approximation ratio (known impossibilities for maximal-in-range mechanisms [13, 17] hold for ex-post mechanisms but not for gross-substitutes valuations).

We then move on to analyzing the approximation ratios achievable by dominant-strategy mechanisms in the standard domain of combinatorial auctions with general (monotone) valuations. From a pure optimization point of view, there is an $O(\sqrt{m})$ approximation algorithm that is not incentive compatible and this is the best achievable with polynomial communication [25, 26]. Whether this is achievable with a deterministic ex-post incentive compatible mechanism remains a major open question, but we are able to answer this question in the negative for dominant-strategy mechanisms (Section 4):

Theorem: Fix $\varepsilon > 0$. The communication complexity of a mechanism that provides an $m^{1-\varepsilon}$ approximation for combinatorial auctions with general valuations in dominant strategies is $exp(m)$.

The best currently known mechanism (dominant-strategy or ex-post incentive compatible) is the simultaneous maximal-in-range algorithm of [23] that guarantees an approximation ratio of $O(\frac{m}{\sqrt{\log m}})$. To put the theorem in context, so far, following a long line of research, the only separation between the approximation ratios achievable by ex-post mechanisms and non incentive compatible algorithms for combinatorial auctions that use polynomial communication was achieved in [4]. This separation applies to two-player combinatorial auctions with XOS valuations, and relies on the taxation framework [14]. Recall that [14] shows the equivalence of ex-post and dominant strategy implementations for two player settings, thus the result of [4] is also the first to separate dominant-strategy mechanisms for combinatorial auctions and their non-truthful counterparts.

However, a proof for our theorem requires more players, since for two players a second-price auction on the bundle of all items provides an approximation ratio of 2.[4] Thus, new tools are required to prove a bound that is worse than 2.

The proof consists of two mains steps. First, we prove in Section 5 that:

Theorem: Fix $\varepsilon > 0$. The communication complexity of a simultaneous algorithm that provides an $m^{1-\varepsilon}$ approximation for combinatorial auctions with general valuations is $exp(m)$.

Simultaneous combinatorial auctions were introduced by [19]: in these (not necessarily incentive compatible) algorithms, all players simultaneously send a message of length $poly(n, m)$ and the allocation is determined based only on these messages. Previous work (e.g., [1, 2, 9, 10]) considered simultaneous combinatorial auctions with restricted classes of valuations, e.g., subadditive valuations.

In the second step, we leverage the hardness result to dominant-strategy mechanisms by showing that the existence of a deterministic dominant-strategies mechanism with approximation ratio c implies a simultaneous algorithm with approximation ratio (close to) c.

We note that for general valuations, there exists a *randomized* dominant strategy mechanism that achieves an approximation ratio of $O(\sqrt{m})$ [21]. The mechanism is a probability distribution over dominant-strategy mechanisms. Hence, we also obtain a separation of the approximation ratio possible by polynomial communication randomized dominant-strategy mechanisms and deterministic dominant-strategy mechanisms. An analogous separation for *ex-post* mechanisms is not known.

Open Questions and Future Directions. We conclude with some open questions. We showed that dominant-strategy mechanisms cannot exactly maximize the welfare in polynomial communication in combinatorial auctions with gross substitutes valuations. As was already mentioned, it is an open question to determine the approximation ratio achievable for this class or for other classes of

[4] The taxation framework [14] offers also a different path to proving bounds for more than 2 players by providing lower bounds on the taxation complexity, but this path was not applied successfully so far.

valuations that were extensively studied in the literature, such as subadditive, XOS, and submodular.

We do provide some evidence that good dominant-strategy mechanisms do not exist. Observe that all useful constructions of deterministic dominant strategy mechanisms that we know are based on simultaneous algorithms. In Section 5 we prove that:

Theorem: Fix $\varepsilon > 0$. The communication complexity of a simultaneous algorithm that provides an $m^{\frac{1}{16}}$ approximation for combinatorial auctions with gross substitutes valuations is $exp(m)$.

This answers an open question of [19]. Before our work, it was not even known if there is a simultaneous algorithm for combinatorial auctions with submodular valuations that achieves a constant approximation ratio.

Another exciting direction is proving impossibilities for randomized mechanisms. A recent line of work provides sub-logarithmic approximation ratios for various classes of valuations [3, 5, 15]. All these mechanisms are a probability distribution over dominant-strategy mechanisms.[5] Are randomized ex-post mechanisms more powerful than dominant-strategy mechanisms?[6]

We end by noting that our mechanisms work in the blackboard model and all messages sent are observable by all players. A more relaxed model would allow private channels between the players and the center. This assumes that the players trust the center not to leak their messages and the private communication channel itself is not leaky. We do not know how to take advantage of this relaxed model, except for the case of combinatorial auctions with k copies from each good, where the mechanism of [7] (the usual outlier) cannot be implemented in dominant strategies but can be implemented in the relaxed model. We leave studying this model to future research.

2 FORMALITIES AND BASIC OBSERVATIONS

In this section we discuss some basic properties of dominant-strategy mechanisms. These properties hold for every possible domain, not only for combinatorial auctions. Thus, in this section $\mathcal{A}$ is the set of alternatives (which are not necessarily allocations) and the valuation of each player is $v_i : \mathcal{A} \to \mathbb{R}$.

Here and subsequently, when we talk about a fixed mechanism $\mathcal{M}$ together with its dominant strategies $\mathcal{S}_1, \ldots, \mathcal{S}_n$ we will slightly abuse notation: We say that player i with valuation v_i sends a message z at vertex r instead of saying that the dominant strategy of player i with valuation v_i is to send message z in at vertex r. We also say that valuations $v_1 \ldots, v_n$ reach a leaf of a protocol, instead of saying that the strategy profile $(\mathcal{S}_1(v_1), \ldots, \mathcal{S}_n(v_n))$ leads to it.

2.1 Minimal Dominant Strategy Mechanisms

In this section, we show that all dominant strategy mechanisms can be simplified without harming their dominant strategy equilibria and without any communication burden. Since our main interest in this paper is in impossibility results, it implies that we can analyze the power of "minimal" dominant strategy mechanisms without loss of generality. Formally:

Definition 2.1. *We say that a mechanism $\mathcal{M}$ is* minimal *with respect to the strategies $(\mathcal{S}_1, \ldots, \mathcal{S}_n)$ and the valuations $V = V_1 \times \cdots \times V_n$ if it satisfies the following properties:*

1. *There are no useless messages in the protocol, i.e. if some player i can send some message in some particular vertex, we assume that it is a dominant strategy for some type v_i to send this message. It immediately implies that for every leaf in the protocol there exist valuations $(v_1, \ldots, v_n)$ such that the strategies $(\mathcal{S}_1(v_1), \ldots, \mathcal{S}_n(v_n))$ reach this leaf.*
2. *There is at least one player i that has two valuations $v_i, v_i' \in V_i$ such that the strategies $\mathcal{S}_i(v_i)$ and $\mathcal{S}_i(v_i')$ dictate sending different messages at the root of the protocol.*

Lemma 2.2. *Let $\mathcal{M}$ be mechanism and strategies $(\mathcal{S}_1, \ldots, \mathcal{S}_n)$ that realize a social choice function $f : V \to \mathcal{A}$ with payments $P_1, \ldots, P_n : V \to \mathbb{R}^n$ in dominant strategies with communication complexity of c bits. Then, there exists a* minimal *mechanism $\mathcal{M}'$ and strategies $(\mathcal{S}_1', \ldots, \mathcal{S}_n')$ that realize f with the payments schemes $P_1, \ldots, P_n$ in dominant strategies with at most c bits.*

Proof. Given a mechanism, we can assume that it has no useless messages because otherwise we can simplify the protocol by not letting player i send this message. Note that removing actions that are dominant strategy for none of the players does not make dominant strategies not dominant.

Similarly, if the second condition does not hold, then due to the fact that there are no useless messages, the root r has only one child. Then, we can delete the root and take his child to be the new root. We continue with this iterative trimming until we reach a vertex that has a player i with a "meaningful" message.

If no such vertex is found, it means that the social choice function and payment schemes are constant for all valuations, so the empty mechanism implements them (it has no root so it satisfies the second condition trivially). □

2.2 Induced Trees of Mechanisms

We now introduce the notion of induced trees and prove a simple property of them. Consider some vertex u in a minimal dominant strategy mechanism. Let $Z_{j,u}$ denote the set of possible messages that player j can send at node u (assume that $Z_{j,u} = \emptyset$ if player i does not send any message at node u). Fix some player i with $Z_{i,u} \neq \emptyset$ and some message profile for the other players $z_{-i}^u = (z_1, \ldots, z_{i-1}, z_{i+1}, \ldots, z_n)$ where each $z_j \in Z_{j,u}$. The *induced tree of player i at vertex u given z_{-i}^u* is the tree that is rooted by u and contains all subtrees that are connected to u by an edge (z_i, z_{-i}^u) for every possible $z_i \in Z_{i,u}$. I.e., we fix the messages of all other players except player i and think about each message z_i as leading to the subtree that the set of messages (z_i, z_{-i}^u) leads to. For an illustration, see Figure 1.

Lemma 2.3. *Fix some player i, vertex u, and messages of the other players z_{-i}^u in a minimal dominant strategy mechanism. Consider the induced tree of player i at vertex u given z_{-i}^u. If alternative $A \in \mathcal{A}$*

[5] Only [15] claim explicitly that the mechanism is dominant strategy and not just ex-post incentive compatible, but this is likely to be the case also for the other papers as they follow the basic structure that was introduced in [15].

[6] In contrast, many of the truthful-in-expectation mechanisms in the literature are based on solving an LP and are not dominant strategies [22, 24], though some dominant-strategy truthful-in-expectation mechanisms do provide an optimal approximation ratio [16]. Analyzing the power of dominant-strategy truthful-in-expectation mechanisms is also a fascinating avenue for future research.

(a) An illustration of the full tree protocol at vertex u.

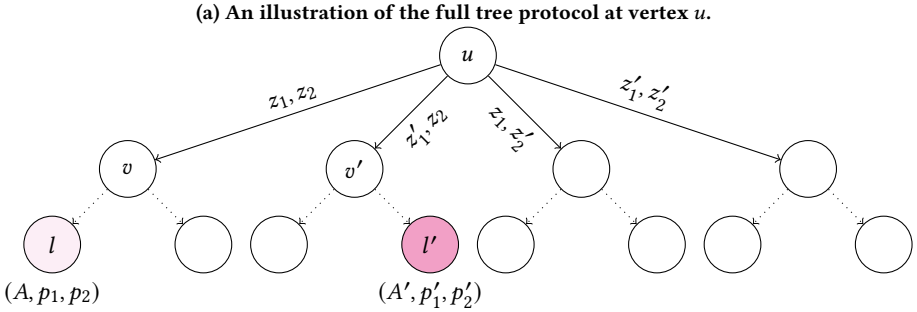

(b) The induced tree of player 1 at vertex u given the message z_2 of player 2. The tree has two subtrees: a left subtree that contains node v and its descendants and a right subtree with node v' and its descendants.

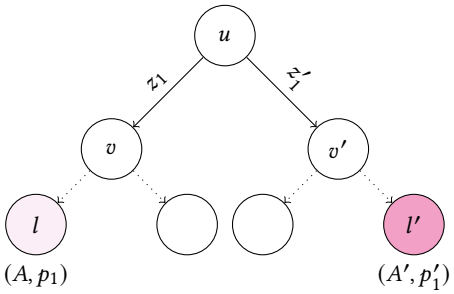

Figure 1: Illustration of the tree rooted at u of a two-player protocol and one of its induced trees. The vertex u satisfies that $Z_{1,u} = \{z_1, z_1'\}$ and $Z_{2,u} = \{z_2, z_2'\}$, i.e. each player has two possible messages. The leaf l that is labeled with (A, p_1, p_2) satisfies that the mechanism outputs alternative A, player pays p_1 and player 2 pays p_2. The same holds for the leaf l' with respect to its outcome (A', p_1', p_2'). The induced tree at Figure 1b describes how the protocol looks like from the point of view of player 1 when player 2 sends the message z_2.

appears in two different subtrees, then all the leaves in this induced tree that are labeled with A have the same payment for player i.

Proof. Let ℓ and ℓ' be two leaves labeled with (A, p_A) and with (A, p_A') that belong in different subtrees, t and t'. By the minimality of the mechanism, every leaf in the protocol has valuations such that $(\mathcal{S}_1(v_1), \ldots, \mathcal{S}_n(v_n))$ reach this leaf. Thus, there exist valuations $v, v' \in V_i, v_{-i}, v'_{-i} \in V_{-i}$ such that:

$$(\mathcal{S}_i(v), \mathcal{S}_{-i}(v_{-i})) \to \ell, \quad (\mathcal{S}_i(v'), \mathcal{S}_{-i}(v'_{-i})) \to \ell'$$

Observe the following strategy profile $\mathcal{S}''_{-i}$: For every valuation v''_{-i}, choose the actions specified by $\mathcal{S}_{-i}(v_{-i})$ until vertex u. Afterwards, at the subtree t, pick the actions that $\mathcal{S}_{-i}(v_{-i})$ specifies, and at the subtrees t' pick the actions that $\mathcal{S}_{-i}(v'_{-i})$ specifies. Since s_{-i} and s'_{-i} do not diverge until vertex u, we have that

$$(\mathcal{S}_i(v), \mathcal{S}''_{-i}(v''_{-i})) \to \ell, \quad (\mathcal{S}_i(v'), \mathcal{S}''_{-i}(v''_{-i})) \to \ell'$$

where the profit of player i with valuation v has to be larger than her profit at ℓ', since $\mathcal{S}_i(v)$ is a dominant strategy for her. Thus, $v(A) - p_A \geq v(A) - p_A'$, so $p_A' \geq p_A$. By applying the same argument for the valuation v', we get that $p_A \geq p_A' \implies p_A = p_A'$. Thus, we have that every two leaves labeled with alternative A in the induced tree of player i given z^u_{-i} have the same payment for player i, which completes the proof. □

3 HARDNESS OF EXACT WELFARE MAXIMIZATION

We now consider two domains that are generally considered "easy" in the sense that the welfare maximizing allocation can be found in time that is polylogarithmic in the representation size of the valuations. For both domains we show that – in contrast to what is perhaps a common misconception – incentive compatible mechanisms that maximize the welfare are incentive compatible only in ex-post equilibrium. For dominant strategy mechanisms, we show that the communication complexity is linear in the size of the representation of the valuations.

Let us first recall how to obtain an ex-post incentive compatible algorithm for combinatorial auctions with two players. Denote the valuations by v_1 and v_2, and for every $1 \leq x \leq m$ let $v_1'(x) = v_1(x) - v_1(x-1)$ and $v_2'(x) = v_2(x) - v_2(x-1)$ be the marginal values. The decreasing marginal values property guarantees that the welfare-maximizing allocation (o_1, o_2) is a point where v_1' and v_2' "cross" each other, i.e. where $v_1'(o_1) \geq v_2'(o_2+1)$ and $v_1'(o_1+1) \leq v_2'(o_2)$ (see also Lemma 6.1). v_1' and v_2' are monotone, so we have to find where two ordered arrays "cross" each other. Thus, a simple binary search will find the optimal allocation with $poly(\log m)$ value queries. VCG prices (player 1 pays $v_2(m) - v(o_2)$, player 2 pays $v_1(m) - v(o_1)$) guarantee incentive compatibility in an ex-post equilibrium.

For combinatorial auctions with gross-substitutes bidders the optimal allocation can be found with communication $poly(m, n)$ for valuations that can be represented by $exp(m)$ bits [28].

Despite the fact that in ex-post equilibrium the optimal welfare can be achieved efficiently, if we require dominant strategies equilibrium, we get an exponential blowup in the communication complexity in both domains.

Theorem 3.1. *Fix a normalized mechanism which implements in dominant strategies a welfare-maximizer for a multi-unit auction where the valuations have decreasing marginal utilities, and the value of a bundle can be represented with $O(\log(m))$ bits. Then, the communication complexity of the mechanism is $\Omega(m \log(m))$.*

Theorem 3.2. *Fix a normalized mechanism which implements in dominant strategies a welfare-maximizer for a combinatorial auction with gross substitutes valuations, where the value of each bundle can be represented with $poly(m)$ bits. Then, the communication complexity of the mechanism is exponential in m.*

The proof Theorem 3.1 can be found in Section 6.1, whilst the proof of Theorem 3.2 can be found in the full version of the paper. Both proofs share a similar structure.

We now give some intuition for the proof in the context of multi-unit auctions with decreasing marginal values. Consider the following scenario. We restrict ourselves to some (large) set of valuations. Suppose that Bob is decisive: for (almost) every allocation $(s, m-s)$, there exist two valuations of Bob v_b^1, v_b^2, such that for each valuation v_a of Alice that is in this set, the optimal allocation in the instances (v_a, v_b^1) and (v_a, v_b^2) is $(s, m-s)$. Furthermore, assume that the dominant strategy of Bob dictates a different message when his valuation is v_b^1 than when it is v_b^2.

Let v_a^1, v_a^2 be two valuations of Alice that are in the set. Since we are implementing a welfare maximizer, Bob must get $m - s$ items for every valuation v_a^1, v_a^2 of Alice. For simplicity, we assume for now (but not in the proof) that we are using VCG payments, so Bob's payment might be different: it can be either $v_a^1(m) - v_a^1(s)$ or $v_a^2(m) - v_a^2(s)$. Thus, if Bob sends a different message for v_b^1 than that of v_b^2 and Alice sends the *same* message for both v_a^1, v_a^2, Bob does not have a dominant strategy, since Alice can "force" him to choose one such message by guaranteeing that his payment will be higher otherwise.

To avoid this, Alice has to "commit" on her value for s items. That is, if v_a^1 and v_a^2 have a different value for s items, then the message that the dominant strategy of Alice dictates cannot be the same for both of them. In fact, we show that this implies, roughly speaking, that Alice's first message must be so informative that we can fully reconstruct Alice's valuation from her first message. Thus, her first message is very big, and the proof is complete. The main challenge of the proof is to construct a big enough set of valuations that satisfies all those properties.

To complement this hardness result, we show that for multi-unit auctions with decreasing marginal values, arbitrarily good approximations are possible in dominant strategies (a "communication FPTAS"):

Theorem 3.3. *For every $\varepsilon > 0$, there is a dominant strategy algorithm for multi-unit auctions with decreasing marginal values that makes $poly(\frac{1}{\varepsilon}, n)$ value queries and provides an allocation with social welfare at least $(1-\varepsilon) \cdot OPT$, where OPT is the value of the optimal social welfare.*

In contrast, the only known upper bound on the approximation ratio of efficient dominant strategy mechanisms for combinatorial auctions with gross substitutes valuations is $O(\sqrt{m})$ [20]. Determining the approximation ratio possible for this class remains an open problem.

4 INAPPROXIMABILITY OF MECHANISMS FOR GENERAL VALUATIONS

In this section we prove that no deterministic dominant strategy mechanism with polynomial communication for general valuations achieves an approximation ratio better than $m^{1-\varepsilon}$. In contrast, there is a *randomized* dominant strategy mechanism that achieves an approximation ratio of $O(\sqrt{m})$ [21]. Note that an approximation ratio of $O(\sqrt{m})$ is the best possible with polynomial communication even when ignoring incentives [28]. We refer the reader to the full version for the exact statement.

The proof is composed of two main steps which we now describe.

Step I: A Lower Bound on Simultaneous Algorithms (Section 5). In general, our approach is to show that dominant-strategy mechanisms for combinatorial auctions with general valuations are as powerful as simultaneous (non-incentive compatible) algorithms. Recall that perhaps the "easiest" way to obtain a dominant strategy mechanism is by designing an ex-post mechanism and making it "simultaneous". Indeed, almost all deterministic dominant-strategy mechanisms in the literature are simultaneous. Thus, the first step is done in Subsection 5.1: a proof that no simultaneous algorithm can achieve an approximation ratio better than $m^{1-\varepsilon}$ with polynomial communication.

Step II: Efficient Dominant Strategy Mechanisms Imply Efficient Simultaneous Mechanisms. Note that not all dominant strategy mechanisms are simultaneous. Consider the following example of a combinatorial auction with two players with additive valuations v_A, v_B. All values are integers between 1 and $\binom{m}{\frac{m}{2}}$. Split the items arbitrarily to two equal sets A and B. Alice can win only items from A, and Bob wins only items from B. We associate each possible value of Alice $v_A(\{b\})$ for some item $b \in B$ with a distinct pair of items in B, and similarly we associate Bob's value $v_B(\{a\})$ for some item $a \in A$ with a distinct pair of items in A. According to the social choice function, Alice wins her more valuable item among the pair that $v_B(\{a\})$ points to and Bob wins his more valuable item among the pair that $v_A(\{b\})$ points to.

A protocol with $O(\log m)$ bits where they simultaneously send $v_A(\{b\})$ and $v_B(\{a\})$ in the first round and then each reports the preferred item among the possible two items is clearly truthful in dominant strategies. However, it is not hard to show that any simultaneous mechanism for this auction requires $\Omega(m \cdot \log m)$ bits. Thus, this instance exhibits a separation between dominant strategy and simultaneous implementations.

On the other hand, we will show that if a mechanism provides an approximation ratio better than $m^{1-\epsilon}$ to the welfare for general valuations, it can used to construct a simultaneous algorithm with comparable approximation ratio. We relegate the proof of this step to the full version of the paper.

5 SIMULTANEOUS ALGORITHMS FOR COMBINATORIAL AUCTIONS

In this section we consider simultaneous combinatorial auctions. The hardness results that we obtain will be used to prove impossibility result for dominant strategy mechanisms for combinatorial auctions with general valuations.

The setup is as follows: as usual, there is a set of items M, $|M| = m$, and n bidders with valuation functions $v_1, \ldots, v_n : 2^M \to \mathbb{R}_+$. Each of them simultaneously sends a message s_i to a central authority; the messages all together are bounded by bit-length L. The algorithm, given the messages, produces an allocation $\mathcal{A}(s_1, \ldots, s_n) = (A_1, A_2, \ldots, A_n)$. The goal is the maximize the social welfare $\sum_{i=1}^{n} v_i(A_i)$. We impose no computational constraints on the bidders or the central authority.

Theorem 5.1. *For m items and $n = \Omega(m^{2-\epsilon})$ bidders with general monotone (binary) functions as valuations, there is no simultaneous mechanism with messages of size at most $\frac{2^{m^{\frac{\epsilon^2}{2}}}}{n}$ which achieves an approximation ratio better than $m^{1-\epsilon}$, for any fixed $\epsilon > 0$.*

Theorem 5.2. *For m items and $n = \Omega(m^{\frac{3}{32}})$ bidders with matroid rank functions as valuations, there is no simultaneous mechanism with messages of length $\frac{2^{m^{\frac{1}{32}}}}{n}$ which achieves an approximation ratio better than $m^{\frac{1}{16}}$.*

The first theorem is used to prove a lower bound for dominant strategy mechanisms. The second one solves an open problem of [19] that asks whether there is a simultaneous algorithm that provides a constant approximation for submodular valuations. Therefore, Theorem 5.2 answers this question negatively, even for matroid rank functions (which are also gross substitutes valuations). We note that the result almost settles completely the approximation ratio achievable in this setting, as a simultaneous $\tilde{O}(m^{\frac{1}{3}})$-approximation algorithm for all subadditive valuations exists [19].

5.1 Proof of Theorem 5.1: An Impossibility for General Valuations

The Hard Distribution. We prove our impossibility for randomized mechanisms by applying Yao's principle. Thus, we now describe a distribution over instances and analyze the performance of deterministic mechanisms on it.

Fix $\epsilon > 0$. Let the number of bidders be $n = m^{2-\epsilon} - m$, divided into $\ell = m^{1-\epsilon} - 1$ groups $G_1, \ldots, G_\ell$ of m bidders each. Let $(A_1, A_2, \ldots, A_\ell, B)$ be a random partitioning of the m items, such that for all j, $|A_j| = |B| = m^\epsilon$ (note that $m^\epsilon(\ell + 1) = m$). For each group G_j, the set of relevant items is $A_j \cup B$. Let $\mathcal{A}_j$ be a family of $t = 2^{\Theta(\epsilon^2 m^\epsilon)}$ subsets of $A_j \cup B$ of size m^ϵ, such that one of the sets is always A_j and the other sets are chosen uniformly at random. By standard concentration bounds, with high probability, these sets overlap pseudo-randomly in the sense that the intersection of any two sets in $\mathcal{A}_j$ has size $(\frac{1}{2} \pm \epsilon)m^\epsilon$. In the following, we will only use a weaker statement which is that for any two sets $A \in \mathcal{A}_j, A' \in \mathcal{A}_{j'}$ such that $A \neq A_j, A' \neq A_{j'}$, we have $A \cap A' \neq \emptyset$ w.h.p. For any two such sets A, A', we have $A \subseteq B \cup A_j$ and $A' \subseteq B \cup A_{j'}$, and the probability that they are disjoint is at most $e^{-\Omega(m^\epsilon)}$, since for every $b \in B$, the probability that $b \in A \cap A'$ is $1/4$ and these events are negatively correlated. The number of such pairs of sets is $2^{\Theta(\epsilon^2 m^\epsilon)}$; i.e. by the union bound, all pairs of sets $A \in \mathcal{A}_j \setminus \{A_j\}, A' \in \mathcal{A}_{j'} \setminus \{A_{j'}\}$ intersect with probability $1 - e^{-\Omega(m^\epsilon)}$.

For each bidder i in group G_j, the valuation is supported on the set of items $A_j \cup B$. For each bidder $i \in G_j$, we choose a random sub-family $\mathcal{B}_i \subseteq \mathcal{A}_j$ such that each set in $\mathcal{A}_j$ appears in $\mathcal{B}_i$ independently with probability $\frac{1}{m}$. More specifically, we do this in such a way that for each set $A \in \mathcal{A}_j$, we choose independently a random bidder $i \in G_j$ for whom $A \in \mathcal{B}_i$; for the other bidders $i' \neq i$, $A \notin \mathcal{B}_{i'}$.

We define the valuation of bidder i as:

$$v_i(S) = \begin{cases} 1 & S \supseteq B \text{ for some } B \in \mathcal{B}_i, \\ 0 & \text{otherwise.} \end{cases}$$

I.e., a bidder i is satisfied if she gets the items of some set in $\mathcal{B}_i$. We call each subset in $\mathcal{B}_i$ a set that bidder i is *interested in*. In particular, if $A_j \in \mathcal{B}_i$, one way to satisfy a bidder in group G_j is to allocate the set A_j. However, this set is valuable only for those bidders $i \in G_j$ such that $A_j \in \mathcal{B}_i$. We call such bidders *special* in group G_j. Note also that only a small number of non-special bidders can be satisfied overall, since these bidders want random sets which intersect with each other with high probability. This leads to the following lemma.

Lemma 5.3. *With probability $1 - e^{-\Omega(m^\epsilon)}$, the welfare of an allocation is at most 1 plus the number of special bidders who receive the respective set A_j.*

Proof. Any player who is not special can get value 1 only if she gets a set in $\mathcal{B}_i$, which does not include the special set A_i. As we argued above, all the sets in $\mathcal{B}_i \setminus \{A_i\}$, for different values of i, intersect pairwise with probability $1 - e^{-\Omega(m^\epsilon)}$. Hence, at most one bidder can be satisfied this way. Any additional value comes from special bidders who receive the respective set A_j. □

Lemma 5.4. *The expected optimal welfare for this instance is $OPT \geq m^{1-\epsilon} - 1$.*

Proof. Each group G_j contains exactly 1 bidder who wants the special set A_j. Hence, a solution which allocates A_j to the special bidder in group G_j, achieves value exactly $\ell = m^{1-\epsilon} - 1$. □

We now analyze the expected welfare achieved by any mechanism on the random instance described above. By Yao's principle, we assume that that the mechanism is deterministic. A good mechanism should ensure that many of the sets A_j go to some special bidder in group G_j. But how can it determine who the special bidders are? For that, it would intuitively need to know the value of A_j for each bidder, but the bidders do not know which of their sets is special and there are too many sets to encode in a message. Our goal is to prove that this indeed implies an impossibility result in the simultaneous model.

We prove that the messages $(s_i : i \in G_j)$ sent by the bidders in group G_j typically do not give us much information about who the special bidder is. Suppose that the messages $(s_i : i \in G_j)$ altogether have bit-length bounded by L. These messages are chosen depending on the random valuations $(v_i : i \in G_j)$, so each choice

of messages appears with a certain probability. We distinguish between "frequent" and "rare" message sets.

Definition 5.5. *We call a message set* $(s_i : i \in G_j)$ frequent *if it appears with probability at least* $\frac{1}{4^L}$*; otherwise it is* rare.

Observe that since the total number of messages is at most 2^L, all rare messages together appear with probability less than $\frac{1}{2^L}$. Next, we prove that a frequent set of messages cannot give us much information about the distribution of high-value sets. Recall that without any conditioning, for a particular bidder $i \in G_j$, each set in $\mathcal{A}_j$ is chosen to be in $\mathcal{B}_i$ with probability $\frac{1}{|G_j|} = \frac{1}{m}$. The key lemma is the following.

Lemma 5.6. *Let* $\bar{s} = (s_i : i \in G_j)$ *be a frequent set of messages. Then for every bidder* $i \in G_j$*, there are fewer than* $L \cdot |G_j|$ *sets* $A \in \mathcal{A}_j$ *such that conditioned on bidders in* G_j *sending* $\bar{s}$*,* $\Pr[A \in \mathcal{B}_i \mid \bar{s}] > \frac{7}{|G_j|}$.

Proof. Suppose that $\bar{s}$ is a frequent set of messages and there is a family of $L \cdot |G_j|$ sets $A \in \mathcal{A}_j$ with $\Pr[A \in \mathcal{B}_i \mid \bar{s}] > \frac{7}{|G_j|}$; denote it by $\mathcal{S} \subset \mathcal{A}_j$.

Consider the choices whether $A \in \mathcal{B}_i$ for $A \in \mathcal{S}$. Without any conditioning, each A is chosen to be in $\mathcal{B}_i$ independently with probability $\frac{1}{|G_j|}$. In expectation, the number of sets in $\mathcal{S} \cap \mathcal{B}_i$ is $\frac{|\mathcal{S}|}{|G_j|} = L$. Hence, by the Chernoff bound,

$$\Pr[|\mathcal{S} \cap \mathcal{B}_i| > (1+\delta)L] \le \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^L < \frac{1}{2^{\delta L}}$$

for $\delta \ge 5$. Consider now the conditioning on $\bar{s}$. Since $\Pr[\bar{s}] \ge \frac{1}{4^L}$, this conditioning cannot increase the probability of any event by more than a factor of 4^L. Therefore,

$$\Pr[|\mathcal{S} \cap \mathcal{B}_i| > (1+\delta)L \mid \bar{s}] < \frac{4^L}{2^{\delta L}}.$$

For $\delta = 5$, we get $\Pr[|\mathcal{S} \cap \mathcal{B}_i| > 6L \mid \bar{s}] < \frac{1}{8^L}$ and the tail probability decays exponentially beyond that. This implies that:

$$\mathbf{E}[|\mathcal{S} \cap \mathcal{B}_i| \mid \bar{s}] < 7L$$

We assumed that $\Pr[A \in \mathcal{B}_i \mid \bar{s}] > \frac{7}{|G_j|}$ for every $A \in \mathcal{S}$, and $|\mathcal{S}| = L \cdot |G_j|$, a contradiction. □

We are now ready to conclude the proof of Theorem 5.1 by showing that for any simultaneous mechanism with messages of total size at most $L = 2^{m^{\epsilon^2/2}}$, executed on a random instance as described above, the expected welfare is $O(1)$, while the optimum is $OPT = \Omega(m^{1-\epsilon})$.

We already showed above that $OPT = \Omega(m^{1-\epsilon})$. Let us bound the expected welfare achieved by bidders in group G_j, assuming that their messages together are bounded by L bits. The contribution of cases where $\bar{s} = (s_i : i \in G_j)$ is a rare message set is small, because they happen with total probability less than $\frac{1}{2^L}$; hence the expected contribution from these cases is negligible (less than $\frac{m}{2^L}$).

In the case of a frequent message set $\bar{s}$, consider the partitioning of the items $B \cup A_j$ among the bidders in group G_j. This partitioning is determined by $\bar{s}$. Lemma 5.6 says that for each bidder $i \in G_j$, fewer than $L \cdot |G_j|$ sets $A \in \mathcal{A}_j$ have the property that $\Pr[A \in \mathcal{B}_i \mid s_i] > \frac{7}{|G_j|}$. Recall that $|G_j| = m$. Hence, among all the sets in $\mathcal{A}_j$, at most $L \cdot |G_j|^2 = Lm^2$ sets are "biased" in the sense that the value is 1 for some bidder with conditional probability more than $\frac{7}{m}$.

Considering group G_j in isolation, the special set A_j is uniformly random among all sets in $\mathcal{A}_j$, and this is true even conditioned on the valuations in group G_j, and hence also conditioned on the message set $\bar{s}$. (Recall that given the set of items $B \cup A_j$ relevant to group G_j, there is no way to distinguish the subset A_j, which is equally likely to be any of the sets in $\mathcal{A}_j$). Furthermore, unless the special set A_j is one of the at most Lm^2 biased sets discussed above, however the items in A_j are allocated, each bidder is the special bidder for it with conditional probability at most $\frac{7}{m}$. If A_j is split among multiple bidders, none of them receives all of A_j. If A_j goes to a particular bidder, then this bidder is special with conditional probability at most $\frac{7}{m}$. Hence, conditioned on a message set $\bar{s}$, we satisfy a special bidder with conditional probability at most $\frac{7}{m}$.

Finally, in case the special set A_j is one of the biased sets, we can assume that we derive value of 1 from it; however this happens with probability at most $\frac{Lm}{|\mathcal{A}_j|} = O(m \cdot 2^{-m^{\frac{\epsilon^2}{2}}})$. The contribution of these cases is negligible.

We have $\ell = m^{1-\epsilon} - 1$ groups of bidders. There are also the items in B, which can contribute value at most 1 in total, with high probability. Hence, the total expected welfare is at most $1 + \frac{7\ell}{m} = O(1)$.

5.2 Proof of Theorem 5.2: Simultaneous Algorithms for Matroid Rank Functions

Here we combine the ideas of Section 5.1 with a construction of matroids by Balcan and Harvey, which we recap here.

Theorem 5.7 ([6]). *For any* $k \ge 8$ *with* $k = 2^{o(\tilde{m}^{\frac{1}{3}})}$*, there exists a family of sets* $\mathcal{A} \subseteq 2^{[\tilde{m}]}$ *and a family of matroids* $\{\mathcal{M}_\mathcal{B} : \mathcal{B} \subseteq \mathcal{A}\}$ *with the following properties:*

- $|\mathcal{A}| = k$ *and* $|A| = \tilde{m}^{\frac{1}{3}}$ *for every* $A \in \mathcal{A}$.
- *For every* $\mathcal{B} \subseteq \mathcal{A}$ *and every* $A \in \mathcal{A}$*, we have:*

$$rank_{\mathcal{M}_\mathcal{B}}(A) = |A|, \quad \text{if } A \in \mathcal{B}.$$
$$rank_{\mathcal{M}_\mathcal{B}}(A) = 8 \log k, \quad \text{if } A \in \mathcal{A} \setminus \mathcal{B},$$

For an instance of combinatorial auctions with m items, we will use this construction with $\tilde{m} = m^{\frac{3}{4}}$ and $k = 2^{m^{\frac{1}{16}}}$; hence $\text{rank}_{\mathcal{M}_\mathcal{B}}(A)$ is either $m^{\frac{1}{4}}$ or $8 \cdot m^{\frac{1}{16}}$, depending on the choice of $\mathcal{B}$.[7]

The Hard Distribution. We prove our impossibility for randomized mechanism by applying Yao's principle. Thus, we now describe a distribution over instances and analyze the performance of deterministic mechanisms on it. We define instances as follows. Let the number of bidders be $n = m^{\frac{1}{8}}(m^{\frac{3}{4}} - m^{\frac{1}{2}} + 1)$, divided into $\ell = m^{\frac{3}{4}} - m^{\frac{1}{2}} + 1$ groups $G_1, \ldots, G_\ell$ of $m^{\frac{1}{8}}$ bidders each. Let $(A_1, A_2, \ldots, A_\ell, B)$ be a random partitioning of the m items, such that $|A_j| = m^{\frac{1}{4}}$ and $|B| = m^{\frac{3}{4}} - m^{\frac{1}{4}}$. (Note that $m^{\frac{1}{4}} \cdot \ell + m^{\frac{3}{4}} - m^{\frac{1}{4}} = m$.)

[7] Note that compared to Balcan-Harvey, we switch the meaning of $\mathcal{B}$ and $\mathcal{A} \setminus \mathcal{B}$; we find it more natural to use $\mathcal{B}$ to denote bases of the matroid. However, the reader should keep in mind that there are also other bases in $\mathcal{M}_\mathcal{B}$.

For each bidder i in group G_j, the valuation is supported on the set of items $A_j \cup B$; it is a matroid rank function of a Balcan-Harvey matroid on $\tilde{m} = m^{\frac{3}{4}}$ elements, with parameter $k = 2^{m^{\frac{1}{16}}}$, defined by set families $\mathcal{B}_i \subseteq \mathcal{A}_j$ and embedded in $A_j \cup B$ so that a random one of the sets in $\mathcal{A}_j$ is mapped onto A_j, and the remaining elements are mapped randomly onto B. (Note that we use a j subscript for $\mathcal{A}_j$, because this family is shared among all the bidders in G_j.) The sub-family $\mathcal{B}_i \subseteq \mathcal{A}_j$ of high-value sets for bidder i is chosen randomly in the following way: For each set $A \in \mathcal{A}_j$, we choose independently and uniformly at random one bidder i in group G_j such that $A \in \mathcal{B}_i$. For all the other bidders $i' \in G_j$, we don't include A in $\mathcal{B}_{i'}$. Note that in expectation we have $\mathbf{E}[|\mathcal{B}_i|] = \frac{|\mathcal{A}_j|}{|G_j|} = m^{-\frac{1}{8}} \cdot 2^{m^{\frac{1}{16}}}$, and $|\mathcal{B}_i|$ is tightly concentrated. Exactly one bidder in group G_j has a high value for the set mapped to A_j, and we call this bidder the *special bidder* in G_j.

Lemma 5.8. *The optimal welfare for this instance is $OPT = m$.*

Proof. In each group G_j, we allocate the special set A_j to the special bidder, who receives value $|A_j| = m^{\frac{1}{4}}$. The items in B can be allocated arbitrarily to some non-special bidders (since $|B| = m^{\frac{3}{4}} - m^{\frac{1}{4}}$ and the number of non-special bidders is $\Omega(m^{\frac{7}{8}})$), who get value 1 each. Hence, each item contributes exactly 1 and $OPT = m$. □

We analyze the expected welfare achieved by any mechanism on the random instance described above. We make the following simple claim.

Lemma 5.9. *If at most m_j of the items in A_j are allocated to the special bidder in group G_j, then the welfare of the allocation is at most $O(m^{\frac{15}{16}}) + \sum_j m_j$.*

Proof. The items in B contribute at most $|B| = m^{\frac{3}{4}} - m^{\frac{1}{4}}$ altogether. Any player who is not special can get value at most $O(m^{\frac{1}{16}})$ from the items in A_j, hence all these players together can get at most $m^{\frac{3}{4}} + O(n \cdot m^{\frac{1}{16}}) = O(m^{\frac{15}{16}})$. Finally, the special players can get at most m_j each from the items in A_j; hence $\sum_j m_j$. □

From here, the proof is similar to the proof of Theorem 5.1. We complete the proof by showing that for any simultaneous mechanism with messages of total size at most $L = 2^{m^{\frac{1}{32}}}$, executed on a random instance as described above, the expected welfare is $O(m^{\frac{15}{16}})$, while the optimum is $OPT = m$.

We already showed above that $OPT = m$. Let us bound the expected welfare achieved by bidders in group G_j, assuming that their messages together are bounded by L bits. The contribution of cases where $\bar{s} = (s_i : i \in G_j)$ is a rare message set is small, because they happen with total probability less than $\frac{1}{2^L}$; hence the expected contribution from these cases is negligible (less than $\frac{m}{2^L}$).

In the case of a frequent message set $\bar{s}$, consider the partitioning of the items $B \cup A_j$ among the bidders in group G_j. This partitioning is determined by $\bar{s}$. Lemma 5.6 says that for each bidder $i \in G_j$, fewer than $L \cdot |G_j|$ sets $A \in \mathcal{A}_j$ have the property that $\Pr[A \in \mathcal{B}_i \mid s_i] > \frac{7}{|G_j|}$. Here, we have $|G_j| = m^{\frac{1}{8}}$. Hence, among all the sets in $\mathcal{A}_j$, at most $L|G_j|^2 = Lm^{\frac{1}{4}}$ sets are "biased" in the sense that the value is high for some bidder with conditional probability more than $\frac{7}{m^{1/8}}$.

The special set A_j is uniformly random among all sets in $\mathcal{A}_j$, and this is true even conditioned on the valuations in group G_j, and hence also conditioned on the message set $\bar{s}$. (Recall that given the set of items $B \cup A_j$ relevant to group G_j, there is no way to distinguish the subset A_j, which is equally likely to be any of the sets in $\mathcal{A}_j$). Furthermore, unless the special set A_j is one of the at most $L \cdot m^{1/4}$ biased sets discussed above, however the items in A_j are split, each bidder is the special bidder for it with conditional probability at most $\frac{7}{m^{1/8}}$. Suppose that bidder i receives k_i items from A_j in this allocation. Then the expected value that the bidders derive from A_j is at most

$$\sum_{i \in G_j} \frac{7}{m^{1/8}} \cdot k_i + \sum_{i \in G_j} \left(1 - \frac{7}{m^{1/8}}\right) O(m^{\frac{1}{16}}) < \frac{7 \cdot |A_j|}{m^{1/8}} + O(m^{1/16}|G_j|) = O(m^{3/16})$$

because a bidder who is special gets value 1 for each item received from A_j, $|A_j| = m^{\frac{1}{4}}$, and a bidder who is not special receives value at most $O(m^{\frac{1}{16}})$ from A_j. Finally, in case the special set A_j is one of the biased sets, we can assume that we derive full value $|A_j| = m^{\frac{1}{4}}$ from it; however this happens with probability at most $L \cdot \frac{m^{\frac{1}{4}}}{|\mathcal{A}_j|} = O(m^{\frac{1}{4}} \cdot 2^{-m^{\frac{1}{32}}})$. The contribution of these cases is negligible.

We have $\ell \leq m^{\frac{3}{4}}$ groups of bidders. There are also the items in B, $|B| \leq m^{\frac{3}{4}}$, which can contribute at most $|B|$ in total. Hence, the total expected welfare is at most $|B| + O(\ell \cdot m^{\frac{3}{16}}) = O(m^{\frac{15}{16}})$.

6 MULTI-UNIT AUCTIONS WITH DECREASING MARGINAL VALUATIONS

Consider a social choice function f that always outputs an allocation that maximizes the welfare. This social choice function can be implemented in dominant strategies by the VCG mechanism. The next theorem shows that even if we restrict ourselves to a subset of the valuations such that each valuation can be represented by $O(m \cdot \log m)$ bits, any dominant-strategy normalized implementation of f requires $\Omega(m \cdot \log m)$ bits, even when there are only two players. In contrast, recall that an ex-post implementation of this set with VCG payments requires only $poly(\log m)$ bits.

We also show an exponential blow up also in the implementation of dominant-strategy welfare maximizers for combinatorial auctions with gross substitute valuations (Theorem 3.2). The two hardness proofs share a very similar structure.

6.1 Hardness Result For Multi-Unit Auctions - Proof of Theorem 3.1

Consider a multi-unit auction of $m \geq 5$ items and two players (Alice and Bob). The valuations that we consider belong to three families: "semi-decisive" valuations V^D, non-decisive valuations V^{ND} and another set of valuations V^P that we will use to show that payments can be used as sketches of valuations.

Every semi-decisive and non-decisive valuation will have a "weight" which is a scalar $\gamma \in \{1, \ldots, m^5\}$ that captures its

magnitude. We now define the set $V^{D,\gamma}$ of semi-decisive valuations with the scalar γ. Every $v \in V^{D,\gamma}$ has two parameters: a special bundle $x^* \in \{2,\ldots,m-2\}$ and $d_m \in \{\frac{1}{2},1\}$ such that:

$$v(x) = \begin{cases} 0 & x = 0, \\ \gamma \cdot 3m^8 & x = 1, \\ \gamma \cdot (m^2 - m + 1) + v(x-1) & x \in \{1,\ldots,x^*\}, \\ \gamma + v(x-1) & x \in \{x^*+1,\ldots,m-1\}, \\ d_m + v(m-1) & x = m. \end{cases}$$

To define the set of non-decisive valuations, we define for every number of items $x \in \{2,\ldots,m\}$ its set of all its possible marginal utilities:

$$\forall x \in \{2,\ldots,m-1\},$$
$$D_x = \{m^2 - mx, m^2 - mx + 1, \ldots, m^2 - m(x-1)\}$$
$$D_m = \{\frac{1}{2}, 1\}$$

For every weight $\gamma \in \{1,\ldots,m^5\}$, every valuation in the set $V^{ND,\gamma}$ is parameterized by a vector $(d_2,\ldots,d_m) \in D_2 \times \cdots \times D_m$ such that:

$$v(x) = \begin{cases} 0 & x = 0, \\ \gamma \cdot 3m^8 & x = 1, \\ \gamma \cdot d_x + v(x-1) & x \in \{1,\ldots,m-1\}, \\ d_m + v(m-1) & x = m. \end{cases}$$

Throughout the proof, we use the notations $V^{ND} = \bigcup_{\gamma=1}^{m^5} V^{ND,\gamma}$ and $V^D = \bigcup_{\gamma=1}^{m^5} V^{D,\gamma}$.

We are now going to define another set of valuations V^P with the purpose of guaranteeing that different valuations in $V^D \cup V^{ND}$ induce different payments. We use this fact later on to sketch valuations. Every $v \in V^P$ has a valuation $v' \in V^{ND} \cup V^D$, a sign $sn \in \{0,1\}$ and a special bundle $t^* \in \{1,\ldots,m\}$ such that:

$$v(x) = \begin{cases} 0 & x = 0, \\ m^{15} + v(x-1) & x < t^*, \\ v'(m-x+1) - v'(m-x) + \frac{(-1)^{sn}}{8m^2} + v(x-1) & x = t^*, \\ v(x-1) & x > t^*. \end{cases}$$

It is easy to see that all the valuations in all three families are normalized, monotone and have decreasing marginal utilities. Also, the value of each bundle can be represented with $O(\log m)$ bits. We begin with a simple observation regarding the properties of welfare maximizing allocations:

Lemma 6.1. *Let $v_A, v_B : [m] \to \mathbb{R}_+$ be multi-unit valuations with decreasing marginal values. Suppose that $s \in \{1,\ldots,m-1\}$ is a number of items such that:*

(1) $v_B(m-s) - v_B(m-s-1) > v_A(s+1) - v_A(s)$.
(2) $v_A(s) - v_A(s-1) > v_B(m-s+1) - v_B(m-s)$.

Then, $(s, m-s)$ is the unique *welfare maximizing allocation. If $v_B(m) - v_B(m-1) > v_A(1) - v_A(0)$, then the unique welfare maximizing allocation is $(0,m)$. Equivalently, $v_A(m) - v_A(m-1) > v_B(1) - v_B(0)$ implies that the only welfare maximizing allocation is $(m,0)$.*

Its proof is relegated to the full version. It is easy to see that the following two propositions together imply Theorem 3.1:

Proposition 6.2. *Let $\mathcal{M}$ be a normalized mechanism with c bits which implements in dominant strategies a welfare maximizer for a multi-unit auction where the valuations have decreasing marginal utilities and the value of a bundle can be represented with $O(\log m)$ bits. Then, there exists $\gamma \in \{1, m^5\}$ such that every element of $V^{ND,\gamma}$ can be represented with at most $c + O(\log(m))$ bits.*

Proposition 6.3. *For every $\gamma \in \{1,\ldots,m^5\}$, The representation size of a valuation in $V^{ND,\gamma}$ is $\Omega(m\log(m))$.*

Proof of Proposition 6.3. By definition, for every $\gamma \in \{1,\ldots,m^5\}$, $|V^{ND,\gamma}| = 2\cdot(m+1)^{m-2}$. Thus, by the pigeonhole principle, the representation size of an element in $V^{ND,\gamma}$ is $\Omega(m\log(m))$ bits. □

6.2 Proof of Proposition 6.2

Fix a dominant strategy normalized two-player mechanism $\mathcal{M}, \mathcal{S}_A, \mathcal{S}_B$ that implements a welfare-maximizer f[8] with payment schemes P_A, P_B for a multi-unit auction where the valuations have decreasing marginal utilities and the value of a bundle can be represented with $O(\log m)$ bits. Observe that $\mathcal{M}$ is in particular dominant strategies when the domain of each player is $V^D \cup V^{ND} \cup V^P$. Denote with c the communication complexity of the mechanism $\mathcal{M}$.

Observe that $\mathcal{M}$ is incentive compatible, so by the taxation principle every valuation v_A of Alice is associated with a menu of prices to Bob, such that for every valuation v_B of Bob the action profile $(\mathcal{S}_A(v_A), \mathcal{S}_B(v_B))$ reaches a leaf that is labeled with a profit-maximizing bundle given this menu. The same can be said of Bob's valuation and the menu presented to Alice.

The proof idea is as follows. We begin by showing that the payments in the menu associated with a valuation are closely related to its values (Subsection 6.2.1). In Section 6.2.2, we show that there exists a set of valuations of Bob such that he sends the price of some bundle (e.g., the price of 1 item), or otherwise Alice's strategy $\mathcal{S}_A$ is not dominant. Consider now two valuations v_B, v_B' from this set that differ only in the price of 1 item. Assume towards a contradiction that Alice has two valuations v_A, v_A' with the same message such that the optimal solution in every one of the four possible instance is $(s, m-s)$ but $P_B(m-s, v_A) \neq P_B(m-s, v_A')$. In this case, the worry is that Alice can determine Bob's payment to be either $P_B(m-s, v_A)$ or $P_B(m-s, v_A')$ without changing Bob's allocation, based only on the price of v_B, v_B' for one item. Thus, Bob will not have a dominant strategy in this case unless Alice commits on the price she displays for $m-s$ items (Subsection 6.2.3). However, if this happens for too many bundles, we can reconstruct Alice's valuation from her first message (Subsection 6.2.4).

6.2.1 Payments Are Good Sketches. We now prove that the payments in the menu that each player presents to the other player are tightly related to the valuation.

[8] There is more than one welfare-maximizer due to tie breaking.

Lemma 6.4. *Let $v_A \in V^{ND} \cup V^D$ and let $x \in \{1, \ldots, m-1\}$ be a number of items. Then:*

$$P_B(x, v_A) \in \left[v_A(m) - v_A(m-x) \pm \frac{1}{8m}\right]$$

where $P_B(x, v_A)$ is the price of x items presented to Bob when Alice has the valuation v_A. Similarly, every valuation of Bob $v_B \in V^{ND} \cup V^D$ and every $x \in \{1, \ldots, m-1\}$ satisfy that:

$$P_A(x, v_B) \in \left[v_B(m) - v_B(m-x) \pm \frac{1}{8m}\right]$$

We relegate the proof of Lemma 6.4 to the full version.

Corollary 6.5. *Fix $v_A \in V^{ND} \cup V^D$ and a number of items $x \in \{1, \ldots, m-1\}$. Given every $P_B(m-x, v_A)$ and $v_A(m)$, the exact value of $v_A(x)$ can be deduced.*

Proof. By Lemma 6.4, we have that $v_A(x) \in [v_A(m) - P_B(m-x, v_A) \pm \frac{1}{8m}]$. Thus, given $P_B(m-x, v_A)$ and $v_A(m)$, we can construct an interval of size $\frac{1+1}{8m} \leq \frac{1}{4}$ such that $v_A(x)$ belongs in it. Recall that $x \leq m$ so by definition $v_A(x)$ is an integer and an interval of size at most $\frac{1}{4}$ has only one integer in it, so we can immediately identify it. □

6.2.2 Bob Reveals Information That Does Not Affect the Allocation. From now on, we focus on the following subsets of valuation sets of Alice and Bob:

$$V_A = V_B = \left\{\bigcup_{\gamma=1}^{m^5} V^{ND,\gamma}\right\} \bigcup \left\{\bigcup_{\gamma=1}^{m^5} V^{D,\gamma}\right\}$$

Observe that the mechanism $\mathcal{M}$ together with the strategies $\mathcal{S}_A, \mathcal{S}_B$ is also a dominant strategy implementation of f, P_A, P_B with respect to $V_A \times V_B$, since they have decreasing marginal values and the value of a bundle can be described with $O(\log m)$ bits. By Lemma 2.2, given the valuations $V_A \times V_B$ there exists a *minimal* dominant strategy mechanism $\mathcal{M}'$ with strategies $(\mathcal{S}'_A, \mathcal{S}'_B)$ that realize the welfare-maximizer f with payment schemes P_A, P_B with $c' \leq c$ bits.

We remind that throughout the proof we slightly abuse notation: we say that a player with valuation v sends a message z at vertex r instead of saying that the dominant strategy of the player is to send message z given the valuation v. We also use the notations V^{γ}, $V^{\leq\gamma}$ or $V^{\geq\gamma}$ to denote all the valuations in V_A or V_B with weight γ, or the valuations with a weight which is smaller or larger than γ.

Observe that there exists a player, without loss of generality Alice, that sends different messages for different valuations in V_A at the root vertex of the protocol, which we denote with r. The reason for that is that $\mathcal{M}'$ is minimal and there exist $(v_A, v_B), (v'_A, v'_B) \in V_A \times V_B$ such that the optimal allocation for them differs. We will show that since she sends non-trivial message in the first round, she has a dominant strategy in $\mathcal{M}'$ only if Bob discloses very specific information that, in certain situations, does not affect the allocation. Formally:

Claim 6.6. *One of the two conditions below necessarily holds:*

1. *For every $v^1_B, v^2_B \in V_B^{D,\gamma=m^5}$ such that $P_A(1, v^1_B) \neq P_A(1, v^2_B)$, Bob sends different messages at vertex r.*
2. *For every $v^1_B, v^2_B \in V_B^{D,\gamma=1}$ such that $P_A(m-1, v^1_B) \neq P_A(m-1, v^2_B)$, Bob sends different messages at vertex r.*

For the proof of Claim 6.6, we prove the following lemma, which is the main working horse of this subsection:

Lemma 6.7. *Let $v^1_A, v^2_A \in V_A$ be two valuations of Alice, and let v^1_B, v^2_B be two valuations of Bob such that:*

1. *The unique optimal solution for the instances (v^1_A, v^1_B) and (v^2_A, v^2_B) is $(x, m-x)$.*
2. *$P_A(x, v^1_B) \neq P_A(x, v^2_B)$.*
3. *Alice sends different messages at the root vertex r for v^1_A and v^2_A.*

Then, *Bob sends different messages at the root vertex r for the valuations v^1_B and v^2_B.*

Proof. Denote with z^1_A and z^2_A the messages that Alice sends for v^1_A, v^2_A. Assume towards a contradiction that Bob sends the same message z_B for the valuations v^1_B, v^2_B at the root vertex r. Let t_1, t_2 be the subtrees that the message profiles (z^1_A, z_B) and (z^2_A, z_B) lead to. Denote with l_1, l_2 the leaves that (v^1_A, v^1_B) and (v^2_A, v^2_B) reach (respectively). For an illustration, see Figure 2.

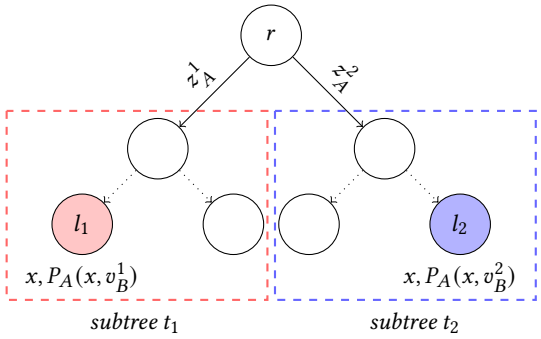

Figure 2: An illustration for the proof of Lemma 6.7. It describes two subtrees t_1, t_2 in the tree that the message z_B of Bob induces for Alice at the root vertex r. The leaves l_1, l_2 are the leaves that (v^1_A, v^1_B) and (v^2_A, v^2_B) reach, so by assumption they are labeled with the allocation x items for Alice with a price of $P_A(x, v^1_B)$ and $P_A(x, v^2_B)$, respectively.

Note that the leaf l_1 is labeled with the allocation $(x, m-x)$ and with the payment $P_A(x, v^1_B)$ for Alice, and similarly the leaf l_2 is labeled with the allocation $(x, m-x)$ and with the payment $P_A(x, v^2_B)$ for Alice. Observe that l_1, l_2 appear in different subtrees t_1, t_2, so by Lemma 2.2, they are labeled with the same payment for Alice. However, $P_A(x, v^1_B) \neq P_A(x, v^2_B)$ by assumption, so we reach a contradiction.

□

The following two lemmas are immediate corollaries of Lemma 6.7:

Lemma 6.8. *Assume that there exist two valuations $v^1_A, v^2_A \in V_A^{\leq m^2}$ that Alice sends different messages for at the root vertex r. Let $v^1_B, v^2_B \in V_B^{D,\gamma=m^5}$ be two semi-decisive valuations of Bob such that $P_A(1, v^1_B) \neq P_A(1, v^2_B)$. Then, Bob sends different messages at the root vertex r for the valuations v^1_B and v^2_B.*

Proof. We begin by showing that for every $v_A \in V_A^{\leq m^2}$ and for every $v_B \in V_B^{D,\gamma=m^5}$, the unique optimal allocation is $(1, m-1)$. By Lemma 6.1, it suffices to prove the inequalities $v_B(m-1) - v_B(m-2) > v_A(2) - v_A(1)$ and $v_A(1) - v_A(0) > v_B(m) - v_B(m-1)$, which hold by definition:

$$v_B(m-1) - v_B(m-2) \geq m^5 > m^2 \cdot (m^2 - m + 1) \geq v_A(2) - v_A(1)$$
$$\implies v_B(m-1) - v_B(m-2) > v_A(2) - v_A(1)$$
$$v_A(1) - v_A(0) \geq 1 \cdot 3m^8 > 1 \geq v_B(m) - v_B(m-1)$$
$$\implies v_A(1) - v_A(0) > v_B(m) - v_B(m-1)$$

Thus, the unique optimal allocation for the instances $(v_A^1, v_B^1), (v_A^2, v_B^2)$ is $(1, m-1)$. Recall that by assumption Alice sends different messages for v_A^1, v_A^2 and that $P_A(1, v_B^1) \neq P_A(1, v_B^2)$, so by Lemma 6.7, Bob sends different messages at the root vertex for v_B^1 and v_B^2, as needed. □

Lemma 6.9. *Assume that there exist two valuations $v_A^1, v_A^2 \in V_A^{\geq m^2}$ that Alice sends different messages for at the root vertex r. Let $v_B^1, v_B^2 \in V_B^{D,\gamma=1}$ be two semi-decisive valuations of Bob such that $P_A(m-1, v_B^1) \neq P_A(m-1, v_B^2)$. Then, Bob sends different messages at the root vertex r for the valuations v_B^1 and v_B^2.*

The proof of Lemma 6.9 is analogous to the proof of Lemma 6.8, and can be found in the full version of the paper. We can now prove Claim 6.6:

Proof of Claim 6.6. Recall that we have assumed (without loss of generality) that there exist two valuations of Alice that she sends different messages for at the root vertex r. It implies that the mechanism $\mathcal{M}'$ satisfies at least one of the following conditions: either Alice sends different messages for two valuations in $V_A^{\leq m^2}$ or she sends different messages for two valuations in $V_A^{\geq m^2}$ (otherwise, she sends the same message for all valuations in V_A, since $V_A^{\leq m^2}$ and $V_A^{\geq m^2}$ are intersecting and $V_A = V_A^{\leq m^2} \cup V_A^{\geq m^2}$).

If she sends different messages for two valuations in $V_A^{\leq m^2}$ at the root vertex r, by Lemma 6.8, we get that for every $v_B^1, v_B^2 \in V_B^{D,\gamma=m^5}$ such that $P_A(1, v_B^1) \neq P_A(1, v_B^2)$, Bob sends different messages at vertex r. Similarly, if she sends different messages for two valuations in $V_A^{\geq m^2}$, then by applying Lemma 6.9 we have that for every $v_B^1, v_B^2 \in V_B^{D,\gamma=1}$ with $P_A(m-1, v_B^1) \neq P_A(m-1, v_B^2)$, Bob sends different messages at vertex r. □

6.2.3 Alice Commits to Bob's Payment. We now use the information revealed by Bob about the semi-decisive valuations in $V_B^{\gamma=1}$ or in $V_B^{\gamma=m^5}$ to show that there exists "large" set of valuations such Alice has to commit to Bob's payment for every possible allocation in the first round of the mechanism. In Section 6.2.4, we will show how to use the payment to reconstruct these valuations.

Observe that we now use the fact that $\mathcal{M}'$ is dominant strategies for Bob. For the statement of the claim, we define $v_{m-s}^{\gamma} \in V^{D,\gamma}$ as the semi-decisive valuation parameterized with weight γ, the special bundle $x^* = m - s$ and $d_m = \frac{1}{2}$.

Claim 6.10. *The following holds for either $\gamma = 1$ or for $\gamma = m^5$. Let $v_A \in V_A^{ND,\gamma}$ be a valuation, and let z_A be the message that Alice sends for it at the root of the protocol. Fix a number of items $s \in \{2, \ldots, m-2\}$ and let z_B be the message that Bob sends at the root if his valuation is the decisive valuation v_{m-s}^{γ} defined above. Denote with t the subtree that the message profile (z_A, z_B) leads to.* Then:

1. *There exists a leaf at subtree t labeled with the allocation $(s, m-s)$.*
2. *Every leaf at subtree t that is labeled with the allocation $(s, m-s)$ satisfies that it is labeled with the payment $P_B(m-s, v_A)$ for Bob.*

Proof. We show that condition 1 of Claim 6.6 implies that Claim 6.10 holds for $\gamma = m^5$. The proof that condition 2 of Claim 6.6 implies that Claim 6.10 holds for $\gamma = 1$ is analogous. Claim 6.10 follows since by Claim 6.6 at least one of the conditions specified in the statement of Claim 6.6 holds.

Assume that condition 1 holds. Let $v_A \in V_A^{ND,\gamma=m^5}$ be a valuation, and let $s \in \{2, \ldots, m-2\}$ be a number of items. Define $v_B, v_B' \in V_B^{D,\gamma=m^5}$ as follows. $v_B = v_{m-s}^{m^5}$ and

$$v_B'(x) = \begin{cases} 0 & x = 0, \\ m^5 \cdot m^8 & x = 1, \\ m^5(m^2 - m + 1) + v'(x-1) & x \in \{2, \ldots, m-s\}, \\ m^5 + v'(x-1) & x \in \{m-s+1, \ldots, m-1\}, \\ v'(m-1) + 1 & x = m. \end{cases}$$

In words, v_B and v_B' are the two decisive valuations with weight $\gamma = m^5$ and special bundle $x^* = m - s$. Note that the only difference between v_B, v_B' is the marginal value of the m'th item.

We begin by explaining why the unique welfare maximizing allocation for the instance (v_A, v_B) is $(s, m-s)$. By Lemma 6.1, it suffices to prove that:

$$v_B(m-s) - v_B(m-s-1) = m^5 \cdot (m^2 - m + 1) >$$
$$m^5 \cdot (m^2 - m) \geq v_A(2) - v_A(1) \geq v_A(s+1) - v_A(s)$$
$$v_A(s) - v_A(s-1) \geq v_A(m-1) - v_A(m-2) \geq m^5 \cdot m >$$
$$m^5 \geq v_B(m-s+1) - v_B(m-s)$$

Thus, the leaf l that $(v_A, v_B = v_{m-s}^{m^5})$ reaches is labeled with the allocation $(s, m-s)$. By definition, this leaf belongs in the subtree t, so we have part 1 of the claim. For the proof of the second part, recall that by Lemma 6.4 we have that:

$$P_A(1, v_B) \leq v_B(m) - v_B(m-1) + \frac{1}{8m},$$
$$P_A(1, v_B') \geq v_B'(m) - v_B'(m-1) - \frac{1}{8m}$$

Therefore:

$$P_A(1, v_B) \leq v_B(m) - v_B(m-1) + \frac{1}{8m} <$$
$$v_B'(m) - v_B'(m-1) - \frac{1}{8m} \leq P_A(1, v_B')$$
$$\implies P_A(1, v_B) < P_A(1, v_B')$$

where the strict inequality holds because $v_B'(m) - v_B(m) = \frac{1}{2}$ and $v_B'(m-1) = v_B(m-1)$. Therefore, by condition 1 of Claim 6.6 we

have that Bob sends different message z'_B for v'_B than the message z_B he sends for v_B at vertex r.

Denote with t' the subtree that the messages (z_A, z'_B) lead to, and denote the leaf in t' that (v_A, v'_B) reaches with l'. Since v_B and v'_B are equal for all the coordinates in $\{1, \ldots, m-1\}$, we have that the unique welfare-maximizing allocation for (v_A, v'_B) is also $(s, m-s)$, so l' is labeled with it. For an illustration, see Figure 3.

Since the mechanism $\mathcal{M}'$ realizes the welfare-maximizer f with the payment schemes P_A, P_B, we have that the leaf l that is labeled with the allocation $(s, m-s)$ is labeled with the payment $P_B(m-s, v_A)$ for Bob. By Lemma 2.3 all the leaves in t and in t' that are labeled with the allocation $(s, m-s)$ have the same price for Bob. By combining these two facts, we get that all the leaves in the subtree t labeled with the allocation $(s, m-s)$ are labeled with the payment $P_B(m-s, v_A)$ for Bob, which completes the proof. □

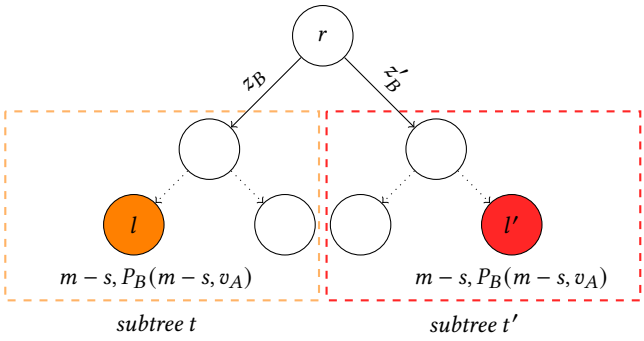

Figure 3: An illustration for the proof of Claim 6.10. It describes the subtrees t, t' in the tree that the message z_A of Alice induces for Bob at the root vertex r. The leaves l, l' are the leaves that (v_A, v_B) and (v_A, v'_B) reach, so as we prove they are labeled with the allocation $(s, m-s)$.

6.2.4 Reconstructing Alice's Valuation. We can now complete the proof of Proposition 6.2. Let $\gamma \in \{1, m^5\}$ be the scalar that Claim 6.10 holds for. We will show how to represent every valuation in $V^{ND,\gamma}$ with at most $c' + O(\log(m)) \leq c + O(\log(m))$ bits (we remind that c, c' stand for the communication complexity of the mechanisms $\mathcal{M}, \mathcal{M}'$).

The representation of a valuation v is composed of the values $v(1), v(m-1), v(m)$ and the message z_A Alice sends at the root vertex r given the valuation v. For every number of items $x \in [m]$, we show how to compute $v(x)$ without any additional communication.

$v(1), v(m-1)$ and $v(m)$ are specified in the sketch. Let $s \in \{2, \ldots, m-2\}$. Let z_B be the message that Bob sends at the root vertex r when his valuation is the decisive valuation $v_B = v^{\gamma}_{m-s}$. Let ℓ be an arbitrary leaf in the subtree that (z_A, z_B) leads to that is labeled with the allocation $(s, m-s)$. By Claim 6.10, such a leaf exists and it is labeled with the payment $P_B(m-s, v_A)$ for Bob. Recall that $v(m)$ is included in the representation, so by Corollary 6.5 we can extract $v(s)$.

6.3 An FPTAS for Multi-Unit Auctions with Decreasing Marginal Values- Proof of Theorem 3.3

In Section 6.1 we showed that no mechanism finds the welfare maximizing allocation in dominant strategies and $poly(\log m)$ communication. In this section we show that this result is tight.

The mechanism is an adaptation of the maximal in range 2-approximation algorithm for general multi unit auctions of [18]. A maximal in range algorithm (see [17],[18]) is an algorithm that finds the welfare maximizing solution in some pre-defined set of allocations. VCG payments are used to guarantee incentive compatibility.

Our maximal-in-range algorithm will split the items into $t = \frac{m}{q}$ bundles of size $q = \lfloor \frac{\varepsilon \cdot m}{n^2} \rfloor$, and (possibly) one additional bundle of size $l = m - t \cdot q$. The maximal-in-range algorithm will optimally distribute these items among the bidders. We implement the algorithm by asking each bidder i with valuation v_i to send, simultaneously with the others, his values for all possible combinations of the bundles: $\{v_i(z \cdot q)\}_{z \leq t}$ and $\{v_i(z \cdot q + l)\}_{z \leq t}$.

It is clear that the number of value queries that the algorithm makes is $poly(n, \frac{1}{\varepsilon})$. In fact, the running time of the algorithm is also polynomial, the proof is essentially identical to that of [18]. The dominant strategy of each bidder is to send the true values, since this is a simultaneous maximal-in-range algorithm. It remains to prove the claimed approximation ratio.

Lemma 6.11. *The social welfare of the allocation that the algorithm outputs is at least* $(1-\varepsilon) \cdot OPT$.

Proof. We will show that there is an allocation in the range with social welfare at least $(1-\varepsilon) \cdot OPT$. Since the algorithm is maximal-in-range, it must output a solution with at least that welfare.

Fix some optimal allocation of the items $(o_1, \ldots, o_n)$. Without loss of generality assume that all items are allocated: $\sum_i o_i = m$. Thus, there must be some bidder, without loss of generality bidder 1, such that $o_1 \geq m/n$.

For each $i > 1$, obtain o'_i by rounding up o_i to the nearest multiple of q. Let $o'_1 = m - \sum_{i>1} o'_i$. Note that this allocation is indeed in the range (each bidder $i > 1$ gets a multiple of q, bidder 1 gets the remaining bundles of size q and the single bundle of size l).

We now analyze the social welfare of the allocation $(o'_1, \ldots, o'_n)$. By the monotonicity of the valuations, for each bidder $i' > 1$ it holds that $v_i(o'_i) \geq v_i(o_i)$. As for bidder 1, it holds that: $o_1 - o'_1 = m - \sum_{i>1} o_i - m + \sum_{i>1} o'_i \leq n \cdot q = n \cdot \lfloor \frac{\varepsilon \cdot m}{n^2} \rfloor \leq \frac{\varepsilon \cdot m}{n}$. Recall that $o_1 \geq \frac{m}{n}$ and that v_1 exhibits decreasing marginal utilities, so by taking away at most ϵ fraction of the items of player 1, his utility decreases by at most $\epsilon \cdot v_1(o_1)$. Thus, $v_1(o'_1) \geq (1-\varepsilon) \cdot v_1(o_1)$ and we have that $\sum_i v_i(o'_i) \geq (1-\varepsilon) \cdot \sum_i v_i(o_i)$, as needed. □

ACKNOWLEDGMENTS

Work supported by BSF grant 2016192, ISF grant 2185/19, and a BSF-NSF grant (BSF number: 2021655, NSF number: 2127781).

REFERENCES

[1] Noga Alon, Noam Nisan, Ran Raz, and Omri Weinstein. 2015. Welfare Maximization with Limited Interaction. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. 1499–1512. https://doi.org/10.1109/FOCS.2015.95

[2] Sepehr Assadi. 2020. Combinatorial auctions do need modest interaction. *ACM Transactions on Economics and Computation (TEAC)* 8, 1 (2020), 1–23.

[3] Sepehr Assadi, Thomas Kesselheim, and Sahil Singla. 2021. Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 653–661.

[4] Sepehr Assadi, Hrishikesh Khandeparkar, Raghuvansh R. Saxena, and S. Matthew Weinberg. 2020. Separating the Communication Complexity of Truthful and Non-Truthful Combinatorial Auctions. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (Chicago, IL, USA) *(STOC 2020)*. Association for Computing Machinery, New York, NY, USA, 1073–1085. https://doi.org/10.1145/3357713.3384267

[5] Sepehr Assadi and Sahil Singla. 2019. Improved truthful mechanisms for combinatorial auctions with submodular bidders. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 233–248.

[6] Maria-Florina Balcan and Nicholas J.A. Harvey. 2011. Learning Submodular Functions. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (San Jose, California, USA) *(STOC '11)*. Association for Computing Machinery, New York, NY, USA, 793–802. https://doi.org/10.1145/1993636.1993741

[7] Yair Bartal, Rica Gonen, and Noam Nisan. 2003. Incentive Compatible Multi Unit Combinatorial Auctions. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge* (Univerity of Indiana, Indiana) *(TARK '03)*. Association for Computing Machinery, New York, NY, USA, 72–87. https://doi.org/10.1145/846241.846250

[8] Liad Blumrosen and Noam Nisan. 2007. *Combinatorial Auctions*. Cambridge University Press, 267–300. https://doi.org/10.1017/CBO9780511800481.013

[9] Mark Braverman, Jieming Mao, and S Matthew Weinberg. 2018. On simultaneous two-player combinatorial auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2256–2273.

[10] Mark Braverman and Rotem Oshman. 2017. A Rounds vs. Communication Tradeoff for Multi-Party Set Disjointness. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. 144–155. https://doi.org/10.1109/FOCS.2017.22

[11] Vincent Conitzer and Tuomas Sandholm. 2004. Expressive Negotiation over Donations to Charities *(EC '04)*. Association for Computing Machinery, New York, NY, USA, 51–60. https://doi.org/10.1145/988772.988781

[12] Jacques Cremer and Richard P McLean. 1985. Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent. *Econometrica* 53, 2 (March 1985), 345–61.

[13] Amit Daniely, Michael Schapira, and Gal Shahaf. 2015. Inapproximability of Truthful Mechanisms via Generalizations of the VC Dimension. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing* (Portland, Oregon, USA) *(STOC '15)*. Association for Computing Machinery, New York, NY, USA, 401–408. https://doi.org/10.1145/2746539.2746597

[14] Shahar Dobzinski. 2016. Computational efficiency requires simple taxation. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 209–218.

[15] Shahar Dobzinski. 2021. Breaking the Logarithmic Barrier for Truthful Combinatorial Auctions with Submodular Bidders. *SIAM J. Comput.* 50, 3 (2021). https://doi.org/10.1137/16M1088594 arXiv:https://doi.org/10.1137/16M1088594

[16] Shahar Dobzinski and Shaddin Dughmi. 2013. On the power of randomization in algorithmic mechanism design. *SIAM J. Comput.* 42, 6 (2013), 2287–2304.

[17] Shahar Dobzinski and Noam Nisan. 2007. Limitations of VCG-Based Mechanisms. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing* (San Diego, California, USA) *(STOC '07)*. Association for Computing Machinery, New York, NY, USA, 338–344. https://doi.org/10.1145/1250790.1250842

[18] Shahar Dobzinski and Noam Nisan. 2010. Mechanisms for multi-unit auctions. *Journal of Artificial Intelligence Research* 37 (2010), 85–98.

[19] Shahar Dobzinski, Noam Nisan, and Sigal Oren. 2014. Economic Efficiency Requires Interaction. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing* (New York, New York) *(STOC '14)*. Association for Computing Machinery, New York, NY, USA, 233–242. https://doi.org/10.1145/2591796.2591815

[20] Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2005. Approximation Algorithms for Combinatorial Auctions with Complement-Free Bidders. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing* (Baltimore, MD, USA) *(STOC '05)*. Association for Computing Machinery, New York, NY, USA, 610–618. https://doi.org/10.1145/1060590.1060681

[21] Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2006. Truthful Randomized Mechanisms for Combinatorial Auctions. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing* (Seattle, WA, USA) *(STOC '06)*. Association for Computing Machinery, New York, NY, USA, 644–652. https://doi.org/10.1145/1132516.1132607

[22] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. 2011. From Convex Optimization to Randomized Mechanisms: Toward Optimal Combinatorial Auctions. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (San Jose, California, USA) *(STOC '11)*. Association for Computing Machinery, New York, NY, USA, 149–158. https://doi.org/10.1145/1993636.1993657

[23] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. 2004. Bundling Equilibrium in Combinatorial Auctions. *Games and Economic Behavior* 47 (2004), 104–123.

[24] Ron Lavi and Chaitanya Swamy. 2005. Truthful and near-optimal mechanism design via linear programming. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. 595–604. https://doi.org/10.1109/SFCS.2005.76

[25] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. 2002. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *JACM 49(5)*. 577–602.

[26] Noam Nisan. 2002. The Communication Complexity of Approximate Set Packing and Covering". In *International Colloquium on Automata, Languages, and Programming*. Springer Berlin Heidelberg, 868–875.

[27] Noam Nisan. 2015. Algorithmic mechanism design: Through the lens of multiunit auctions. In *Handbook of Game Theory with Economic Applications*. Vol. 4. Elsevier, 477–515.

[28] Noam Nisan and Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129, 1 (2006), 192–224. https://doi.org/10.1016/j.jet.2004.10.007

[29] Christos Papadimitriou, Michael Schapira, and Yaron Singer. 2008. On the Hardness of Being Truthful. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. 250–259. https://doi.org/10.1109/FOCS.2008.54

[30] Aviad Rubinstein, Raghuvansh R Saxena, Clayton Thomas, S Matthew Weinberg, and Junyao Zhao. 2021. Exponential communication separations between notions of selfishness. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 947–960.

[31] W. Vickrey. 1961. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance* (1961), 8–37.