Communication Separations for Truthful Auctions: Breaking the Two-Player Barrier

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

Clayton Thomas
Microsoft Research
Cambridge, USA
clathomas@microsoft.com

S. Matthew Weinberg
Princeton University
Princeton, USA
smweinberg@princeton.edu

Qianfan Zhang
Princeton University
Princeton, USA
qianfan@princeton.edu

2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS) | 979-8-3315-1674-1/24/$31.00 ©2024 IEEE | DOI: 10.1109/FOCS61266.2024.00032

***Abstract*—We study the communication complexity of truthful combinatorial auctions, and in particular the case where valuations are either subadditive or single-minded, which we denote with SubAdd∪SingleM. We show that for three bidders with valuations in SubAdd∪SingleM, any deterministic truthful mechanism that achieves at least a 0.366-approximation requires $\exp(m)$ communication. In contrast, a natural extension of [Fei09] yields a non-truthful $\text{poly}(m)$-communication protocol that achieves a $\frac{1}{2}$-approximation, demonstrating a gap between the power of truthful mechanisms and non-truthful protocols for this problem.**

Our approach follows the taxation complexity framework laid out in [Dob16b], but applies this framework in a setting not encompassed by the techniques used in past work. In particular, the only successful prior application of this framework uses a reduction to simultaneous protocols which only applies for *two* bidders [AKSW20], whereas our three-player lower bounds are stronger than what can possibly arise from a two-player construction (since a trivial truthful auction guarantees a $\frac{1}{2}$-approximation for two players).

I. Introduction

1) Background: Combinatorial Auctions: Combinatorial Auctions are a paradigmatic problem at the intersection of Economics and Computation; see e.g. [LOS02, LLN06, NS06, MSV08, DNS10, DSS15, AKS21] and many others. Here, an auctioneer has m items to allocate among n bidders, where each bidder i has a combinatorial valuation v_i over subsets of items (that is, $v_i : 2^{[m]} \to \mathbb{R}_{\geq 0}$), and seeks to do so in a way that maximizes the *welfare*. That is, the auctioneer seeks to partition the items into n subsets, $A_1, \ldots, A_n$, so as to maximize $\sum_i v_i(A_i)$. The challenge is that only Bidder i knows her valuation $v_i(\cdot)$, and therefore: (a) the auctioneer must communicate with the bidders to learn sufficient information about $v_i(\cdot)$ to find a high-welfare allocation, and (b) the bidders may strategically manipulate the protocol, and therefore the auctioneer must further design a *truthful* communication protocol.[1]

Subject only to constraint (a), the problem is fairly well understood for several canonical valuation classes. For example, $\text{poly}(m)$-communication (not necessarily truthful) protocols are known to achieve: an asymptotically tight $\Theta(1/\sqrt{m})$-approximation for arbitrary monotone valuations [LS05, NS06], a tight $1/2$-approximation for subadditive valuations [Fei09], a tight $(1-1/e)$-approximation is known for XOS valuations [Fei09, DNS10], and the tight constant for submodular valuations is known to lie in $[1 - 1/e + 1/10^5, 1 - 1/2e]$ [DV13, FV10, Von07]. Thus, for all of these problems, the optimal approximation guarantee of poly-communication non-truthful protocols are well-understood.

Subject only to constraint (b), the problem can be solved optimally by the classical Vickrey-Clark-Groves (VCG) mechanism [Vic61, Cla71, Gro73]; however, this solution requires exponential communication for any of the above-mentioned valuation classes.[2] Subject to *both* constraints (a) and (b), the problem is still quite poorly understood despite significant effort over the past two decades. For example, the state-of-the-art deterministic truthful mechanisms with $\text{poly}(m)$-communication achieve approximation guarantees of $\Omega(\ln(m)/m)$ for arbitrary monotone valuations, and just $\Omega(\sqrt{\ln(m)/m})$ for each of subadditive/XOS/submodular valuations [QW24]. Thus, these canonical settings each have a $\widetilde{\Theta}(\sqrt{m})$ gap in approximation guarantees for state-of-the-art truthful vs. non-truthful protocols. Despite these massive gaps, corresponding impossibility results remain starkly rare; indeed, for each of the above examples, no lower bounds for truthful mechanisms are known beyond those that also hold for non-truthful protocols. Thus, for each of these canonical settings, *it remains unknown whether there should be any gap at all*!

This question has received significant attention since [NS06] introduced the study of communication complexity for combinatorial auctions, and it is generally conjectured that the more significant missing piece is stronger impossibility results. Developing such impossibility results necessarily requires leveraging truthfulness, which imposes significant technical barriers. For example, a natural first attempt would be to characterize all truthful auctions in a given setting, perhaps a là Roberts' Theorem [Rob79], and then prove impossibility results for communication-efficient auctions using this classification. [LMN03, DN15] push the classification approach

S. Matthew Weinberg and Qianfan Zhang are supported by NSF CCF-1955205.

[1] See Section II for a formal definition – informally, a protocol is truthful if every player is incentivized to follow the protocol.

[2] The VCG mechanism reduces non-truthful protocols to truthful ones, but only for protocols that *exactly* optimize welfare. Thus, the VCG mechanism makes the problem easy in any setting where describing a valuation function only takes polynomially-many bits. Richer classes, like the four mentioned, require exponentially-many bits to describe, and require high communication for exact welfare maximization, thus motivating approximations.

2575-8454/24/$31.00 ©2024 IEEE
DOI 10.1109/FOCS61266.2024.00032

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

roughly to its limit, but bizarre deterministic truthful mechanisms exist and it appears that any classification attempt may be intractable (see the discussion in [Dob11], for example).

Interest in the question revitalized when [Dob16b] proposed an alternative framework termed *Taxation Complexity*. This framework does not attempt to classify truthful mechanisms, but merely identifies one key complexity measure that bounds the communication complexity of a truthful auction: namely, the logarithm of the maximum number of menus presented to a player, also called the taxation complexity. The Taxation Complexity framework further implies the following corollary: communication lower bounds on *two-player simultaneous (non-truthful) protocols* imply communication lower bounds on *two-player interactive truthful mechanisms* (whereas simultaneous lower bounds certainly do not generally imply interactive lower bounds for non-truthful protocols [Yao79]). Even this corollary proved challenging to leverage, but [AKSW20] later established an impossibility of $3/4 - 1/240$ for two XOS bidders, which exhibits the first separation of truthful and non-truthful protocols, since there is a non-truthful protocol that gives a $3/4$-approximation for two XOS bidders [Fei09, DNS10].

[AKSW20] remains to-date the only known separation between what is achievable by $\text{poly}(m)$-communication truthful mechanisms and $\text{poly}(m)$-communication (non-truthful) protocols. Yet, this separation holds only for two bidders. Indeed, while a $\text{poly}(m)$-communication non-truthful protocol can guarantee a $3/4$-approximation for two XOS bidders, the best possible guarantee for even three XOS bidders is $18/27 < 3/4 - 1/240$.[3] That is, while we now know that two-bidder $\text{poly}(m)$-communication protocols achieve strictly better guarantees than two-bidder $\text{poly}(m)$-communication truthful mechanisms, it is still plausible that truthful mechanisms are just as powerful as non-truthful protocols for $n \geq 3$ XOS bidders. Moreover, the approach of [AKSW20] heavily leverages the two-bidder corollary of [Dob16b]: their result follows from a communication lower bound on simultaneous protocols, which does not imply anything about interactive truthful mechanisms for $n \geq 3$ bidders.

2) Our Contributions: In our main result, we provide the first separation of $\text{poly}(m)$-communication truthful vs. non-truthful protocols for more than 2 bidders. Specifically, we consider the class of bidders that are either Single-Minded or Subadditive (and term the class $\mathsf{SubAdd} \cup \mathsf{SingleM}$). We show that for this class, any deterministic truthful mechanism beating a $(\sqrt{3}-1)/2 \approx 0.366$ approximation requires $\exp(m)$ communication, whereas a simple extension of [Fei09] achieves a $1/2$-approximation with a non-truthful protocol.[4]

3) Brief Overview of Approach: We leverage the Taxation Complexity framework of [Dob16b], which hinges around the concept of a *menu*. The menu for Bidder i of an auction $\mathcal{A}$ is a function $M = \mathsf{Menu}_i^{\mathcal{A}}(v_{-i})$, parameterized by v_{-i}, that takes as input v_i and outputs the set of items (and price paid) that Bidder i gets on input (v_i, v_{-i}). That is, a menu fixes the valuations of bidders other than i, and stores the impact that Bidder i's valuation v_i has on i's own allocation and prices.

Informally speaking, one can think of the results of [Dob16b] as showing that, in terms of the communication-efficiency of truthful auctions (in sufficiently rich settings), the auction might as well be implemented by fully learning a bidder's menu, and then allocating to that bidder according to their menu and their valuation. Indeed, in these settings, [Dob16b] proves that for any truthful auction, the communication complexity of the auction is at least (some polynomial function of) the communication complexity of learning one bidder's menu.[5] To prove our main result, we establish that in order for a truthful mechanism to beat a 0.366-approximation for $\mathsf{SubAdd} \cup \mathsf{SingleM}$, the Communication Complexity of learning a bidder's menu must be exponential in the number of items m. The Communication Complexity of all truthful mechanisms achieving the same approximation ratio then follows via results from [Dob16b].

We defer technical details, but now briefly give intuition for the role of both Subadditive bidders and the Single-Minded bidder, and how the 0.366-approximation factor arises. A construction of [EFN+19] establishes that exponential communication is necessary in order for a (not necessarily truthful) protocol to beat a $1/2$-approximation for two subadditive bidders. So intuitively, truthful welfare-maximizing auctions face the following challenge. Suppose there is one Single-Minded and two Subadditive bidders. Imagine that the maximum possible welfare between the two Subadditive bidders is c. Then if, based on the two Subadditive bidders, the price of set S for the Single-Minded bidder is set to $p(S)$, it could be that:

- Perhaps the Single-Minded bidder has interest set S, has value barely exceeding $p(S)$ (and therefore chooses to purchase set S), and yet the maximum welfare achievable between the two Subadditive bidders for $\overline{S}$ is ≈ 0. Therefore, we could have welfare as bad as $\approx p(S)$ when the optimum is c.
- Perhaps the Single-Minded bidder has interest set S, has value barely below $p(S)$ (and therefore chooses not to purchase anything), and yet the maximum welfare achievable between the two Subadditive bidders for $\overline{S}$ is also c. Moreover, even when allocating all items to the two Subadditive bidders, we don't expect to achieve welfare greater than $c/2$ without $\exp(m)$-communication, based on the results of [EFN+19]. Therefore, we could have welfare as bad as $c/2$ when the optimum is $c+p(S)$.

Therefore, the best ratio we can hope to achieve is $\min\{\frac{p(S)}{c}, \frac{c/2}{c+p(S)}\}$, which (it turns out) is at most $\frac{\sqrt{3}-1}{2} \approx 0.366$ no matter how we set $p(S)$.

[3] Moreover, [BMW18] design a simultaneous protocol for two XOS bidders achieving an approximation guarantee of $23/32 > 18/27$, so no lower bound via two-player simultaneous protocols can possibly achieve a separation for more than 2 players.

[4] We also include novel results on related classes of valuations with two bidders; see Section I-A.

[5] [Dob16b] provides other ways to bound the communication complexity of truthful auctions, as we discuss below.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

Of course, this is just intuition—we need an actual construction where the above arguments hold even after the bidders have engaged in polynomial communication. For example, if we know the Single-Minded Bidder is interested in a particular set S, then it is trivial to determine whether the maximum welfare achievable between the two Subadditive bidders for $\overline{S}$ is ≈ 0 or $\approx c$. Nevertheless, the high-level intuition of our proof closely follows the outline above. In particular, the role of the two Subadditive Bidders in the is to: (a) have uncertainty regarding whether $\overline{S}$ can generate non-trivial welfare between them, and (b) require $\exp(m)$ communication to beat welfare $c/2$ allocating any set of items to them. The role of the Single-Minded Bidder is for a clean and direct analysis of what the price of a single set S might mean for the resulting allocation, thus allowing us to reason about the menu presented to the Single-Minded Bidder.

A. Roadmap and Conclusions

We give preliminaries in Section II. In Section III, we give a "warmup" to our main result, namely, we give a false proof attempting to implement the intuitive hardness outline discussed above, but leaving a gap in one step of the proof. After explaining this gap and the main ideas used to fix it, in Section IV, we prove our main result: a welfare approximation lower bound of 0.366 for poly-communication truthful auctions for three bidders with valuations in $\mathsf{SubAdd} \cup \mathsf{SingleM}$. We then observe that this gives a separation between truthful auctions and non-truthful protocols, by showing that non-truthful protocols can achieve a higher welfare approximation of $\frac{1}{2}$.[6]

In Section V, we study auctions for two bidders in the class of valuations $\mathsf{XOS} \cup \mathsf{SingleM}$. Observe that to obtain a separation for two-bidder mechanisms, it is necessary to consider a "smaller" class than $\mathsf{SubAdd} \cup \mathsf{SingleM}$.[7] We show that for two bidders in the class of valuations $\mathsf{XOS} \cup \mathsf{SingleM}$, the communication complexity of every truthful mechanism that has an approximation better than $(\sqrt{5}-1)/2 \approx 0.618$ is $\exp(m)$. We show this bound by two different proofs based on two different approaches from [Dob16b]. Finally, in Section V-C we also discuss difficulties for generalizing it to a stronger lower bound for more than two $\mathsf{XOS} \cup \mathsf{SingleM}$ bidders. Holisticly, our results in Section V illustrate additional ways to achieve impossibility results, and in fact served as stepping stones en-route to our main result.

[6]Additionally, we note that our results get a separation for any $n = O(\log n)$ bidders (since our upper bound of a $1/2$-approximation extends to any $n = O(\log n)$ bidders, and not only 3; see Lemma II.3). This is in contrast to [AKSW20]: while they prove upper and lower bounds yielding a separation for 2 bidders, there is no known separation with 3 or more bidders in their setting (since, while their lower bounds trivially also holds for more bidders, their upper bound degrades as the number of bidders increases).

[7]The reason for it is that for $\mathsf{SubAdd} \cup \mathsf{SingleM}$, there is provably no gap between the power of communication efficient truthful mechanisms and non-truthful protocols for two bidders, as the second-price auction on the grand bundle is truthful, communication-efficient, and $1/2$-approximates the social welfare, which is optimal even for non-truthful protocols [EFN+19]. Thus, to show a gap for two bidders, it is necessary to consider a "smaller" class.

We conclude by restating that our main result is the first separation between the approximation guarantees of $\mathrm{poly}(m)$-communication truthful and non-truthful combinatorial auctions beyond two bidders for any class of valuations. Still, the major open problem of whether $\mathrm{poly}(m)$-communication deterministic truthful mechanisms can achieve an approximation guarantee better than $\widetilde{\Omega}(\sqrt{m})$ remains open. An obvious direction for future work is to make further progress, extending our result either by removing the need for a single-minded bidder, or by achieving super-constant lower bounds for more bidders. It is also important to investigate whether our techniques open doors for similar results with other canonical valuation classes (e.g., submodular valuations, arbitrary monotone valuations), or the related setting of multi-unit auctions.

B. Related Work

1) Communication Complexity of Deterministic Truthful Combinatorial Auctions: The most related work to ours concerns the Communication Complexity of Truthful Combinatorial Auctions. Here, the best approximation ratios are guaranteed by "VCG-Based" (also called "Maximal-in-Range") mechanisms [HKMT04, DNS10, DN11, QW24], which are a factor of $\widetilde{\Theta}(\sqrt{m})$ worse than those of the best non-truthful protocols [LOS02, Fei09, DNS10, FV10] for the canonical settings of Submodular, XOS, Subadditive, and Monotone. It is plausibly conjectured that the aforementioned VCG-based mechanisms are optimal among deterministic truthful mechanisms, which would imply that strong separations between truthful and non-truthful protocols exist.

The lone prior separation is [AKSW20]. This separation relies on a result of [Dob16b] which reduces the problem to showing impossibilities for two-bidder simultaneous protocols; this reduction does not extend beyond two bidders.[8] In contribution to this line of works, we provide the first separation for > 2 bidders and directly via the Taxation Complexity framework [Dob16b] in a manner that is not restricted to a particular number of bidders. Finally, we also note that the particular notion of truthfulness we consider is the standard "Ex-Post Nash," meaning that it is always a Nash equilibrium for Bidders to follow the protocol, no matter their valuations. A stronger notion of truthfulness asks that it is a "Dominant Strategy" to follow the protocol, and [RST+21, DRV22] establish separations between the achievable guarantees of $\mathrm{poly}(m)$-communication Dominant Strategy Truthful mechanisms and non-truthful protocols.[9]

2) Communication Complexity of Randomized Truthful Combinatorial Auctions: There is also a significant line of work developing randomized truthful combinatorial auctions. Here, the state-of-the-art approximation guarantees are much

[8][DNO14, DRV22] also prove lower bounds for simultaneous combinatorial auctions. However, their results hold only for a large number of bidders.

[9]The distinction between Ex-Post Nash and Dominant-Strategy Truthful is that the former need only incentivize each bidder to behave truthfully when the other bidders are truthful *according to some, but possibly arbitrary, valuation*, whereas the latter faces the much stiffer task of incentivizing each bidder to be truthful *even when other bidders are behaving in a bizarre manner inconsistent with any valuation function.*

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

better, $\text{poly}(1/\log\log m)$ for Submodular, XOS, and Subadditive [AKS21] (building upon [AS19, Dob16a, Dob07, KV12]), and $\Theta(\sqrt{m})$ for arbitrary monotone valuations [DNS12]. The latter is asymptotically tight, whereas it remains a major open problem whether the former can be improved to a constant.

3) Computational Complexity of Truthful Combinatorial Auctions: A significant body of work also considers the computational complexity of truthful Combinatorial Auctions. This line of research indeed concludes strong separations between the constant-factor approximations achievable by poly-time algorithms [LLN06, MSV08, Von08] and poly-time truthful mechanisms for Submodular valuations [Dob11, DV12b, DV12a, DV16, DV11]. These results establish an $\widetilde{\Omega}(\sqrt{m})$ lower bound on the approximation guarantees of randomized poly-time truthful mechanisms. The simple posted-price mechanisms of the prior paragraph 'break' these bounds because they use demand queries, and [CTW20] identifies a relaxed notion of "truthfulness" under which these mechanisms achieve their guarantees in poly-time. This counterintuitive interaction between computation and incentives further motivates the communication model for combinatorial auctions.

II. Preliminaries

1) Combinatorial auctions: Recall that in a combinatorial auction, there are m heterogeneous items and n bidders. We denote items by $j \in [m] = \{1, 2, \ldots, m\}$, and bidders by $i \in [n]$. The set of allocations of items to the n bidders, i.e., sets $(A_1, \ldots, A_n)$ with $A_i \subseteq [m]$ and $A_i \cap A_j = \emptyset$ for all $i \neq j$, is denoted by Σ. Each bidder i holds a private valuation function $v_i : 2^{[m]} \to \mathbb{R}$ which is drawn from some set $\mathcal{V}_i$. The sets $\mathcal{V}_i$ differ depending on the auction problem considered; we discuss different canonical cases below. We assume that all valuations are monotone non-decreasing ($v_i(S) \leq v_i(T)$ for $S \subseteq T$) and normalized ($v_i(\emptyset) = 0$). The goal is to find an allocation of items $(A_1, \ldots, A_n) \in \Sigma$ that (approximately) maximizes the social welfare $\sum_{i \in [n]} v_i(A_i)$.

An *auction rule*, equivalently known as a direct-revelation mechanism, is a function $\mathcal{A} : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \Sigma \times \mathbb{R}^n$, that maps a valuation profile $(v_1, \ldots, v_n)$ to an outcome $\big((A_1, \ldots, A_n), p_1, \ldots, p_n\big)$, where S_i is the allocation of player i and p_i is her payment. The auction rule can also be equivalently be defined as a tuple $(f, p_1, \ldots, p_n)$, where the allocation is determined via the function $f : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \Sigma$ and the payments are determined via the functions $p_i : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \mathbb{R}$ for each player $i \in [n]$. When no confusion can arise, we do not distinguish between these two representations of the auction rule. We denote by $f_i(v_1, \ldots, v_n)$ the bundle that bidder i receives in $f(v_1, \ldots, v_n)$. When bidder i receives bundle S_i and is charged payment p_i, bidder i's utility is $v_i(S_i) - p_i$. We assume that bidders are rational and strategic, and thus aim to maximize their utility.

We say that an auction rule is *truthful* (for the relevant classes of valuations $\mathcal{V}_1, \ldots, \mathcal{V}_n$) if it satisfies the following: for every bidder i, for every two valuations $v_i, v_i' \in \mathcal{V}_i$ for bidder i, and valuations $v_{-i} \in \mathcal{V}_{-i}$ for other players,

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

Intuitively, this means truthfully-reporting is always i's best strategy.[10] An auction rule gives an *α-approximation to the optimal welfare* (for the valuation classes $\mathcal{V}_1, \ldots, \mathcal{V}_n$) if for every valuation profile $(v_1, \ldots, v_n) \in \mathcal{V}_1 \times \cdots \times \mathcal{V}_n$:

$$\sum_{i=1}^{n} v_i(f_i(v_1, \ldots, v_n)) \geq \alpha \cdot \max_{\substack{(S_1, \ldots, S_n) \\ \in \Sigma}} \sum_{i=1}^{n} v_i(S_i).$$

2) Communication complexity and protocols: We study the communication complexity of implementing auction rules (which, we recall, calculate both the allocation and the payments charged). Technically, we study the n-player deterministic number-in-hand blackboard communication model, where the input known to each player i is her valuation function $v_i \in \mathcal{V}_i$ and all the messages sent are visible to all the bidders. The communication cost of a protocol is the maximum number of bits that are written to the blackboard in the worst case. The communication complexity of some problem, denoted with $\text{cc}(\cdot)$, is the minimum communication cost of a protocol that computes it. For clarity, we phrase our results (e.g.,) as "for any truthful auction rule $\mathcal{A}$ for valuations $\mathcal{V}$ achieving an α-approximation to the optimal welfare, the communication complexity of $\mathcal{A}$ is at least C". It is well known and easy to see that the communication complexity of a protocol is at least the log of the number of transcripts; see e.g. [KN97]. Our lower bounds will be based on this fact.

In our main results, we discuss three-player protocols for three-bidder auctions. We refer to these three bidders as Alice, Bob, and Charlie, and we denote $i \in \{A, B, C\}$ for shorthand (so that, for example, these bidders' valuation functions are v_A, v_B, v_C).

3) Valuation classes: As mentioned above, different combinatorial auction problems are defined by different classes of bidder valuations; typically, one studies valuations with some canonical property.[11] Our main result concerns the following two classes.

Definition II.1 (Single-minded valuations)**.** A valuation function $v : 2^{[m]} \to \mathbb{R}$ is *single-minded* if there exists a weight

[10] Note that, following most of the algorithmic mechanism design literature, we consider a mechanism truthful so long as the underlying auction rule is truthful, i.e., our definition is independent of the communication protocol used to implement the mechanism. In technical jargon, this means we we study implementations of truthful auction rules in ex-post equilibria, i.e., *ex-post* incentive-compatible auction rules, as opposed to the stronger notion of *dominant-strategy* incentive compatibility (where no bidder would regret acting according the true value even if other bidders act in a way which is inconsistent with any valuation function, see e.g. [RST+21, DRV22]). Note that working with a weaker notion of truthfulness makes our lower bounds only stronger.

[11] Formally, specifying a class of valuation functions for use in a communication protocol requires specifying a number of items, and a range / precision of possible numerical values. Formally, one can say that a class of valuations uses *precision* k if for all v in the class and all bundles S, we have $v(S) \in \{0\} \cup \{x/2^k \mid x \in [2^{2k}]\}$. Following most of the algorithmic mechanism design literature, we leave k implicit in all our communication complexity bounds. Formally, this means that we always hide factors of $\text{poly}(k)$ in the communication costs of the protocols we construct, and that all of our lower bounds / impossibility theorems hold for some $k = \text{poly}(m, n)$. For more discussion of the issue of representation of numbers, see [Dob16b, Appendix A.4.1] and our Section D.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

$w \in \mathbb{R}_{\geq 0}$ and a set $T \subseteq [m]$ such that $v(S) = w \cdot \mathrm{I}[S \supseteq T]$ for all $S \in 2^{[m]}$, where $\mathrm{I}[\cdot]$ denotes the indicator function. We denote the family of all single-minded valuation functions over m items by $\mathsf{SingleM} = \mathsf{SingleM}_m$.

Definition II.2 (Subadditive valuations)**.** A valuation function $v : 2^{[m]} \to \mathbb{R}$ is *subadditive* if, for all bundles $S, T \subseteq [m]$, we have $v(S \cup T) \leq v(S) + v(T)$. We denote the family of all subadditive valuation functions over m items by $\mathsf{SubAdd} = \mathsf{SubAdd}_m$.

SingleM is in some sense the most basic class of valuations with economic complementarities—the bidder gets positive utility w if and only if she gets her desired bundle T. On the other hand, SubAdd is often considered the most general canonical class of valuations *without* complementarities—when two bundles are combined, the bidder's utility can never increase beyond the sum of her value on the individual bundles.

4) Black-box welfare approximation reduction with $O(\log m)$ single minded bidders: In our paper, we focus on showing a separation between the power of communication-efficient truthful auctions and their non-truthful counterparts for the class $\mathsf{SubAdd} \cup \mathsf{SingleM}$. However, the known non-truthful protocols are defined only for subadditive bidders. Fortunately, we can easily extend the existent protocols to take care of single-minded bidders without loss in the approximation guarantee, via the following black-box reduction:

Lemma II.3. *Let $\mathcal{P}$ be a protocol that achieves an approximation α for n bidders with valuations in the class $\mathcal{C}$ with $\mathrm{poly}(m)$ bits. Then, there exists a protocol $\mathcal{P}'$ that achieves an approximation α for n bidders with valuations in the class $\mathcal{C} \cup \mathsf{SingleM}$ with $\mathrm{poly}(m, 2^n)$ bits.*

Briefly, the protocol $\mathcal{P}'$ proceeds by asking each bidder whether they are single-minded, and running protocol $\mathcal{P}$ separately for each of the $\leq 2^n$ sub-problem assuming that a given subset of single-minded bidders are allocated. We defer the full proof of Lemma II.3 to Section A.

5) Menus: Following the seminal work of [Dob16b], our approach to studying the communication complexity of truthful auctions centers around the notion of bidders' *menus*. Given a truthful mechanism $\mathcal{A}$, bidder i's *menu* given the valuations v_{-i} of the other players specifies a price p_S for every subset of items $S \subseteq [m]$. Thus, if player i wins S, then she pays p_S; the fact that such a p_S is well-defined follows immediately from truthfulness, via Proposition II.5 below. We use the following notation for menus:

Definition II.4 (Menus)**.** For any truthful auction rule $\mathcal{A} = (f, p_1, \ldots, p_n)$, any player i, and any valuation profile $v_{-i} \in \mathcal{V}_{-i}$, define the *menu* of player i as the function $\mathsf{Menu}_i^{\mathcal{A}}(v_{-i}) : 2^{[m]} \to \mathbb{R}$ such that, for all bundles $S \subseteq [m]$ such that there exists a v_i with $f_i(v_i, v_{-i}) = S$, we have $\mathsf{Menu}_i^{\mathcal{A}}(v_{-i})(S) = p_i(v_i, v_{-i})$. We often write $\mathsf{Menu}(v_{-i})$ where the mechanism and the player presented with the menu are clear from context.

Note that the menu is well-defined only if for every player i, and for every valuation profile $v_{-i} \in \mathcal{V}_{-i}$ of the other players, the price a bidder is charged can depend only on the bundle S that player i receives, and cannot vary with valuation v_i (so long as player i still receives bundle S). Fortunately, this fact follows immediately from truthfulness, goes back to at least [Ham79], and (following [Gue81]) is known as the taxation principle:

Proposition II.5 (Taxation Principle)**.** *Consider a truthful auction rule $\mathcal{A} = (f, p_1, \ldots, p_n) : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \Sigma \times \mathbb{R}^n$. For every bidder $i \in [n]$, every valuation profile $v_{-i} \in \mathcal{V}_{-i}$, and every bundle $S \subseteq [m]$, if $f_i(v_i, v_{-i}) = f_i(v_i', v_{-i}) = S$ for two valuations $v_i, v_i' \in \mathcal{V}_i$, then $p_i(v_i, v_{-i}) = p_i(v_i', v_{-i})$.*

Note also that truthfulness directly implies that for all (v_i, v_{-i}), the bundle of items i receives is (one of) the bundles S that maximizes $v_i(S) - M(S)$, where $M = \mathsf{Menu}_i(v_{-i})$.

A property of menus that will be useful is that we can assume that they are be non-decreasing and normalized without loss of generality, just like valuation functions.

Proposition II.6 (Menu monotonicity [Dob16b])**.** *Let $\mathcal{A} = (f, p_1, \ldots, p_n) : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \Sigma \times \mathbb{R}^n$ be any deterministic truthful auction rule. There exists some other mechanism $\mathcal{A}' = (f, p_1', \ldots, p_n') : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \Sigma \times \mathbb{R}^n$ with the same allocation function f, such that for every $i \in [n], v_{-i} \in \mathcal{V}_{-i}$, the menu $\mathsf{Menu}_i^{\mathcal{A}'}(v_{-i})$ for bidder i is non-decreasing, i.e., $\mathsf{Menu}_i^{\mathcal{A}'}(v_{-i})(S) \leq \mathsf{Menu}_i^{\mathcal{A}'}(v_{-i})(T)$ for $S \subseteq T$, and normalized, i.e., $\mathsf{Menu}_i^{\mathcal{A}'}(v_{-i})(\emptyset) = 0$.*

6) Lower Bounds through Menus: [Dob16b] gives several techniques for lower bounding the communication complexity of truthful auctions; all of these techniques hinge around the concept of the menu. To get some intuition as to how this works, observe that one way to implement a truthful combinatorial auction is to select a player i, run an $(n-1)$-player communication protocol among the bidders other than i to determine i's menu $M = \mathsf{Menu}_i(v_{-i})$, then query i to determine i's highest-utility bundle according to M.[12] Informally, [Dob16b] shows that something quite special happens for truthful auctions (in sufficiently rich domains), namely, truthful auctions can do *essentially no better* in terms of communication-efficiency than this approach of explicitly learning bidders' menus.

The above intuition is actually formalized multiple different ways in [Dob16b], with a large emphasis on the Taxation Complexity, which counts (the log of) the number of distinct menus. We recall this technique (and the reduction to simultaneous protocols exploited in [AKSW20]) when we need them in Section V. To prove our main result, we use the following result from [Dob16b], which gives an even stronger lower bound technique than Taxation Complexity:

Theorem II.7 (Follows from [Dob16b])**.** *Consider any deterministic, truthful auction rule $\mathcal{A}$ for m items and $n \geq 2$*

[12]Note that this intuitive sketch ignores potentially-complex issues of tie-breaking; these considerations are handled in full detail in [Dob16b].

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

bidders with valuation functions in $\mathcal{V}$. Suppose $\mathcal{V} \supseteq$ SubAdd. Let $\mathrm{cc}(\mathcal{A})$ be the communication complexity of $\mathcal{A}$. Then, for any player $i \in [n]$, there exists an $(n-1)$-player protocol $\mathcal{P}$ with $\mathrm{poly}(\mathrm{cc}(\mathcal{A}), m, n)$ bits of communication that computes (the index of) the menu $\mathsf{Menu}_i^{\mathcal{A}}(v_{-i})$.

Theorem II.7 is a direct corollary of [Dob16b]'s Theorem 3.1, Theorem 2.3, and Proposition F.1. See the proof in Section A.

III. Building Up To Our Main Result

The main result of this paper is a lower bound on the communication complexity of getting good welfare in a three-bidder truthful auctions. In particular, we will show that for the class of bidder valuations SubAdd $\cup$ SingleM, any truthful auction for three bidders which gets $\frac{\sqrt{3}-1}{2} \approx 0.366$ approximation requires $\exp(m)$ communication.

In this section, we give exposition into the proof of this result via a simplified, but ultimately faulty, version of the proof which conveys much of the ideas (and how our real proof overcomes the faults of this expository version).

A. Failed Proof Attempt

To give exposition into our proof, we first give a simpler "failed attempt". In this "failed attempt", there will be a gap in the proof (as we explain when the gap appears); we will also gloss over techniques from previous works and claim without proof that a certain constructions exist. In our actual proof, we will correct the gap, and also provide full details and proofs for all constructions.

1) (Failed) proof outline: Let $\mathcal{A}$ be a truthful auctions for SubAdd $\cup$ SingleM with three-bidder and m items. By Theorem II.7, the communication complexity of $\mathcal{A}$ is at most (a polynomial function of) the communication complexity of determining a bidder's menu. Thus, our goal is to show that if $\mathcal{A}$ gets a good approximation to the optimal welfare, then the communication complexity of determining a bidder's menu is $\exp(m)$.

Our approach is inspired by the classical rectangle argument in communication complexity. This argument proves lower bounds on communication complexity by constructing a large set of inputs (the "fooling set"), considering any pair of inputs in the fooling set, and showing that every protocol computing a function f must use a different transcript for those two inputs. However, in order to make our argument work, it turns out we need to extend this approach to consider 4-*tuples* of inputs rather than pairs.

After constructing the valuations used in our hard instance, we lower bound the communication complexity of Bob and Charlie calculating Alice's menu. Our (failed) proof outline for this proceeds in two steps.

- First, we show that, if a communication-efficient and truthful auction rule induces the same menu on a 4-tuple of "sufficiently different" inputs, then the auction cannot get a good approximation to the optimal welfare. If such a 4-tuple uses the same menu, we call this a "bad" 4-tuple; for a precise definition, see below. This step uses direct arguments about our hard instance, and about the welfare obtainable under different menus and different possible valuations of Alice.
- Second, we consider any protocol for calculating Alice's menu which avoids all bad 4-tuples, and show that all such protocols must have high communication. This step is a direct, simple lemma in communication complexity.

Combining these two steps finishes the (failed) proof.

2) Description of the Class of Valuations: We begin by defining the class of subadditive valuation functions that our simplified construction uses. This construction is based on exponentially many overlapping 4-cell partitions of all of the items; say these partitions are denoted $\mathcal{G}_{i,1} \cup \mathcal{G}_{i,2} \cup \mathcal{G}_{i,3} \cup \mathcal{G}_{i,4} = [m]$ for many different values of i, say $i \in [K]$ for some large K. The key property of the construction is that, for each way to select one cell of each partition, there is a valuation in this class who, for each partition i simultaneously, wants *only* elements of the selected cell of partition i. Intuitively, [EFN+19] construct valuations like this in order to hide an allocation with good welfare by reducing it to an equality problem—high welfare can be achieved among two bidders with valuations in this class if and only if there is some partition on which the bidders want different items.

Formally, for some family of partitions $\left\{\{\mathcal{G}_{i,1}, \mathcal{G}_{i,2}, \mathcal{G}_{i,3}, \mathcal{G}_{i,4}\}\right\}_{i,j}$, we define a class of valuations $\mathcal{V}$ such that there exists a constant ℓ such that for every pair of $i \in [K]$ and $j \in [4]$, there exists a valuation $v \in \mathcal{V}$ such that $v(\mathcal{G}_{i,j}) = \ell - 1$ and $v([m] \setminus \mathcal{G}_{i,j}) = 1$. We use strings $\mathbf{b} \in [4]^K$ to denote ways to pick a cell of each partition, where $\mathbf{b}[i]$ denotes the ith character in the string and $\mathbf{b}[i] = j \in [4]$ denotes selecting cell $\mathcal{G}_{i,j}$ from partition i.

Formally, the class of valuation functions we consider is given by the following:

Proposition III.1. *Fix any constant ℓ. For all sufficiently large m, and for some $K = \Omega(\exp(m))$, there exists:*

1) *a family of 4-cell partitions $\{\mathcal{G}_{i,j} \mid i \in [K], j \in [4]\}$ of $[m]$ (i.e., for all i, we have $\bigcup_{j \in [4]} \mathcal{G}_{i,j} = [m]$ and $\mathcal{G}_{i,j} \cap \mathcal{G}_{i,j'}$ for all $j \neq j'$), and*
2) *a family of subadditive valuation functions $v_{\mathbf{b}} : 2^{[m]} \to \mathbb{R}_{\geq 0}$ indexed by strings $\mathbf{b} \in [4]^K$, such that for all $\mathbf{b} \in [4]^K$ and all $i \in [K]$:*
   - *a) $v_{\mathbf{b}}([m]) = \ell$,*
   - *b) $v_{\mathbf{b}}(\mathcal{G}_{i,\mathbf{b}[i]}) = \ell - 1$, and*
   - *c) $v_{\mathbf{b}}([m] \setminus \mathcal{G}_{i,\mathbf{b}[i]}) = 1$.*

   We name this class $\mathcal{V}_{\text{4-cell}}$.

For an illustration of the construction, see Figure 1. To prove this proposition, one can use the main construction of [EFN+19], along with an argument using the probabilistic method. We will later specify and formally prove an even stronger construction; for now, we assert this proposition without proof.

[EFN+19] use a construction analogous to the valuations in Proposition III.1 to lower bound the communication complexity of maximizing welfare with two subadditive bidders.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

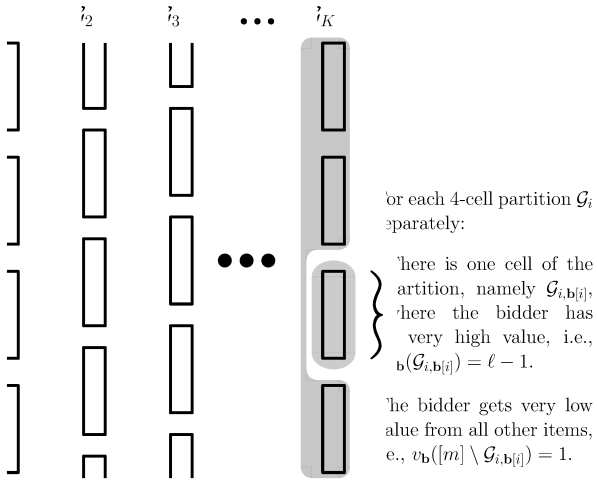

Fig. 1: Illustration of valuations in Proposition III.1.

Theorem III.2 ([EFN+19]). *For any communication protocol $\mathcal{P}$ that outputs an allocation that gives approximation better than $\frac{1}{2}$ for two bidders with valuations in $\mathcal{V}_{4\text{-cell}}$, the communication cost of $\mathcal{P}$ is $\exp(m)$ bits.*

For deterministic protocols (our focus), one can prove this result directly by reducing to the equality problem on the strings $\mathbf{b}$; for details and further results for randomized protocols (which we do not need), see [EFN+19]. Note that Theorem III.2 shows hardness even for non-truthful protocols, and not just for truthful auction rules.

3) (Failed proof that) getting good welfare implies "no bad 4-tuple": We now proceed to argue that "sufficiently different" elements of the valuation family given by Proposition III.1 cannot possibly present the same menu in a truthful, communication-efficient and approximately optimal auction. Note, however, that this part of the proof has a gap, as we explain when this gap occurs, and later correct in the full version of the proof. Recall that, when we discuss three-bidder auctions, we refer to the three bidders as Alice, Bob, and Charlie (abbreviated A, B, and C).

We now describe the notion of "bad 4-tuples". For some fixed auction rule, we say that a 4-tuple $(\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3, \mathbf{b}_4)$ of strings in $[4]^K$ is *bad* if it satisfies the following two properties:

- The strings are all pairwise different on some single index i_*, i.e., we have $\{\mathbf{b}_1[i_*], \mathbf{b}_2[i_*], \mathbf{b}_3[i_*], \mathbf{b}_4[i_*]\} = \{1, 2, 3, 4\}$.
- Bob and Charlie present the same menu to Alice when their valuations are $(v_{\mathbf{b}_1}, v_{\mathbf{b}_2})$ or when they are $(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$, i.e., $\mathsf{Menu}_A(v_{\mathbf{b}_1}, v_{\mathbf{b}_2}) = \mathsf{Menu}_A(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$.

We will next give a false proof that any communication-efficient auction rule achieving a good welfare approximation must avoid all bad 4-tuples. Intuitively, the key will be to show that a bad 4-tuple implies that the welfare is sometimes bad, by considering Alice with a single-minded valuation such that the optimal allocation is very different on $(v_{\mathbf{b}_1}, v_{\mathbf{b}_2})$ versus on $(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$. The gap in the proof will come from the way we apply Theorem III.2, as we explain below.

(Failed) Proposition III.3. *Consider an auction rule $\mathcal{A}$ which is truthful for three bidders with valuations in $\mathsf{SubAdd} \cup \mathsf{SingleM}$, has a $\mathrm{poly}(m)$-bit communication protocol, and gets a $(\sqrt{3}-1)/2 + \xi \approx 0.366 + \xi$ approximation to the optimal welfare for some constant $\xi > 0$. Now, suppose $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3, \mathbf{b}_4 \in [4]^K$ and $i_* \in [K]$ are such that $\{\mathbf{b}_1[i_*], \mathbf{b}_2[i_*], \mathbf{b}_3[i_*], \mathbf{b}_4[i_*]\} = \{1, 2, 3, 4\}$. Then, $\mathsf{Menu}^{\mathcal{A}}_A(v_{\mathbf{b}_1}, v_{\mathbf{b}_2}) \neq \mathsf{Menu}^{\mathcal{A}}_A(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$.*

(Failed) Proof.. Fix a truthful auction rule $\mathcal{A}$, and four vectors $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3, \mathbf{b}_4 \in [4]^K$ such that $\{\mathbf{b}_1[i_*], \mathbf{b}_2[i_*], \mathbf{b}_3[i_*], \mathbf{b}_4[i_*]\} = \{1, 2, 3, 4\}$. Without loss of generality, assume that $\mathbf{b}_j[i^*] = j$ for $j \in [4]$. Informally speaking, what this means is that for the "special" partition $\mathcal{G}_{i_*,1} \cup \mathcal{G}_{i_*,2} \cup \mathcal{G}_{i_*,3} \cup \mathcal{G}_{i_*,4} = [m]$, we have that for $j = 1, \ldots, 4$, valuation $v_{\mathbf{b}_j}$ "loves" set $\mathcal{G}_{i_*,j}$ and "hates" the rest of the items, $[m] \setminus \mathcal{G}_{i_*,j}$. Now, suppose for contradiction that $\mathsf{Menu}^{\mathcal{A}}_A(v_{\mathbf{b}_1}, v_{\mathbf{b}_2}) = \mathsf{Menu}^{\mathcal{A}}_A(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$, and denote this menu with M for simplicity.

Consider the case when Alice has a single-minded valuation v such that she values the set $S = \mathcal{G}_{i_*,1} \cup \mathcal{G}_{i_*,2}$ with some value x (and thus Alice values all sets containing S at x as well, and all sets not containing S at 0). We will tune the precise value of x later; for now, suppose only that $x < \ell$, and that $M(S) \neq x$ (which is without loss of generality, since we can always slightly perturb the value of x after fixing $\mathcal{A}$ and $\mathbf{b}_j$ for $j = 1, \ldots, 4$).

Now, we examine two cases based on whether the mechanism $\mathcal{A}$ allocates S to Alice or not.

- Suppose $M(S) < x$. In this case, Alice must receive (a bundle containing) $S = \mathcal{G}_{i_*,1} \cup \mathcal{G}_{i_*,2}$ whenever the menu is M. Suppose that Bob's and Charlie's valuations are $(v_{\mathbf{b}_1}, v_{\mathbf{b}_2})$.
  By the definition of $v_{\mathbf{b}}$ and the fact that $\mathbf{b}_1[i^*], \mathbf{b}_2[i^*] \in \{1, 2\}$, we have that $v_{\mathbf{b}_1}([m] \setminus S), v_{\mathbf{b}_2}([m] \setminus S) \leq 1$. Thus, $\mathcal{A}$ gets welfare at most $x + 2$.
  On the other hand, consider the allocation that gives nothing to Alice, gives $\mathcal{G}_{i_*,1}$ to Bob, and gives $\mathcal{G}_{i_*,2}$ to Charlie. The welfare of this allocation is $2(\ell - 1)$. Thus, the welfare approximation ratio of $\mathcal{A}$ in this case is at most $(x + 2)/\big(2(\ell - 1)\big)$.
- Suppose $M(S) > x$. In this case, Alice cannot receive a bundle containing $S = \mathcal{G}_{i_*,1} \cup \mathcal{G}_{i_*,2}$ when the menu is M. Suppose that Bob's and Charlie's valuations are $(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$.
  False Argument: Now, one may *intuitively* think that since Alice gets no items in this case, we are just solving the allocation problem of awarding items between Bob and Charlie, who are two bidders with valuations drawn from the class of valuations specified in Proposition III.1. On these instances, giving $\mathcal{G}_{i_*,3}$ to Bob and $\mathcal{G}_{i_*,4}$ to Charlie gets welfare $2(\ell - 1)$. However, *intuitively*, by Theorem III.2, any communication-efficient auction rule should sometimes get welfare at most ℓ,

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

reflecting the $(1/2)$-approximation hardness result of [EFN+19]. This argument is *not actually correct*, because the communication-efficient protocol for $\mathcal{A}$ can actually use information from Alice's valuation to focus on the specific partition $\mathcal{G}_{i_*}$.[13] However, to illustrate the idea, let us pretend that this is true for the remainder of the proof, and that $\mathcal{A}$ gets welfare at most ℓ in this case.

On the other hand, consider the allocation that gives S to Alice, gives $\mathcal{G}_{i,3}$ to Bob, and gives $\mathcal{G}_{i,4}$ to Charlie. The welfare of this allocation is $x + 2(\ell - 1)$, and thus the optimal welfare on this instance is at least this large. Thus, the welfare approximation ratio in this case is at most $\ell/\big(x + 2(\ell - 1)\big)$.

Therefore, under our assumption that the menus for Alice with respect to $(v_{\mathbf{b}_1}, v_{\mathbf{b}_2})$ and $(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$ are identical, the welfare approximation ratio of $\mathcal{A}$ can be at most

$$\max\left\{\frac{x+2}{2(\ell-1)}, \frac{\ell}{x+2(\ell-1)}\right\}.$$

Taking $x = (\sqrt{3}-1)\ell$, the ratio above converges to $\frac{\sqrt{3}-1}{2} \approx .366$ as $\ell \to \infty$.[14] □

4) "No bad 4-tuple" implies high communication: (Failed) Proposition III.3 implies in particular that if the auction rule has any bad 4-tuples, then the auction cannot get a good approximation to the optimal welfare. We will show next that for any protocol which avoids all bad 4-tuples, the communication cost of the protocol must be high. For clarity and generality, we state this part of the argument for general communication problems.

Lemma III.4 (4-tuple rectangle argument)**.** *For a function $f : \mathcal{T} \times \mathcal{T} \to \mathcal{R}$, let $\mathcal{S} = \{v_{\mathbf{b}}\}_{\mathbf{b} \in [4]^K} \subseteq \mathcal{T}$ be a set of distinct inputs to f parameterized by a K-length string $\mathbf{b} \in [4]^K$. Suppose that for every 4-tuple $(\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3, \mathbf{b}_4) \in \mathcal{S}^4$ that differs on some index $i \in [K]$ (i.e., $\{\mathbf{b}_1[i], \mathbf{b}_2[i], \mathbf{b}_3[i], \mathbf{b}_4[i]\} = \{1,2,3,4\}$), we have $f(v_{\mathbf{b}_1}, v_{\mathbf{b}_2}) \neq f(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$. Then, the communication complexity of f is at least*

$$\log_2\left(4^K/3^K\right) = \Omega(K).$$

Proof. Consider any communication protocol for f, and for any inputs $v_B, v_C \in \mathcal{T}$, let $\tau(v_B, v_C)$ be the transcript of the protocol on inputs v_B, v_C. Recall that, as in standard rectangle arguments, if $f(v_B, v_C) \neq f(v'_B, v'_C)$ then $\tau(v_B, v_C) \neq \tau(v'_B, v'_C)$. Also, if $\tau(v_B, v_C) = \tau(v'_B, v'_C)$ then $\tau(v'_B, v_C) = \tau(v_B, v'_C) = \tau(v_B, v_C)$.

To show our lower bound, we focus on the transcripts used when both players have the same input in $\mathcal{S}$. For any transcript T of the protocol, let $\mathcal{I}(T) \subseteq \mathcal{S}$ denote the (possibly empty) set of all $v_{\mathbf{b}} \in \mathcal{S}$ such that the transcript given the input $(v_{\mathbf{b}}, v_{\mathbf{b}})$ is T.

We will now show that for all transcripts T, we have $|\mathcal{I}(T)| \leq 3^K$. For that, we claim that for every transcript T and for every $i \in [K]$, we have $\{\mathbf{b}[i] \mid \mathbf{b} \in \mathcal{I}(T)\} \neq \{1,2,3,4\}$. To see this, assume for contradiction that there is such a transcript T and an index i such that $\tau(v_{\mathbf{b}_j}, v_{\mathbf{b}_j}) = T$ for each $j = 1, \ldots, 4$. Then, by a standard rectangle argument, we have $T = \tau(v_{\mathbf{b}_1}, v_{\mathbf{b}_2}) = \tau(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$ as well. However, by the assumption of the lemma, we have $f(v_{\mathbf{b}_1}, v_{\mathbf{b}_2}) \neq f(v_{\mathbf{b}_3}, v_{\mathbf{b}_4})$, which is a contradiction. Thus, since the string corresponding to elements in $\mathcal{I}(T)$ can take at most 3 values on each index i, we know that $|\mathcal{I}(T)| \leq 3^K$.

Now, this implies that any protocol correctly computing f requires at least $|\mathcal{S}|/3^K = 4^K/3^K$ distinct transcripts in total. Thus, the communication complexity of the protocol has to be $\Omega(K)$.

□

If the failed Proposition III.3 had actually been true, then combining it with Lemma III.4 would immediately imply that the communication complexity of the two-player function $f = \mathsf{Menu}_A(\cdot,\cdot)$ is $\Omega(K) = \exp(m)$. Furthermore, by the techniques of [Dob16b] recalled in our Theorem II.7, we could then conclude that the auction rule itself has communication complexity at least $\exp(m)$.

B. Correcting the gap in this proof

We now briefly discuss the main idea we use below to correct the gap in the failed proof of Proposition III.3. For an argument like the above to work, we need to have a construction such that it will not only be hard (communication-wise) to decide whether we want to allocate to Alice or not, but even conditional on our decision regarding Alice, it will be communication-hard for us to find a "good enough" partition of the items of Bob and Charlie. For that, we create a two-layered construction, which we describe in detail in Section IV.[15] This two-layered construction is more involved that the one discussed in Section III-A2 and Theorem III.2, but its proof is a straightforward extension of the techniques used, i.e., those of [EFN+19].

After making this adjustment, the (false) Proposition III.3 remains nearly identical in how we argue about the welfare achieved by the auction with a "bad 4-tuple". However, the definition of a "bad 4-tuple" becomes considerably more involved. For instance, a "bad 4-tuple" is defined by a condition both about protocols for calculating the menu, and about protocols for calculating the auction rule. This technical fact precludes using a clean and self-contained "generalized

[13] Put another way, Theorem III.2 directly implies the following: if a poly-communication auction rule never allocates any items to Alice, then there must exist some inputs where the auction rule gets welfare ℓ, while the optimal welfare Bob and Charlie could get amongst themselves is at least $2(\ell - 1)$. Completing this proof would require a stronger property: for *any* poly-communication auction rule, there exists some case where the auction rule gets welfare ℓ, but the optimal total welfare is in fact $x + 2(\ell - 1)$.

[14] This specific value of x arises from assuming $x = \rho\ell$ for some $\rho \in (0,1)$, then making the two terms in the above maximum equal to each other while ignoring lower-order factors. In more detail, assuming that $\rho\ell/(2\ell) = \ell/(\rho\ell + 2\ell)$ for $\rho > 0$ implies that $\rho^2 + 2\rho - 2 = 0$, and thus that $\rho = \sqrt{3} - 1$. Then, the above maximum goes to $\rho/2$ as $\ell \to \infty$.

[15] Note that we do not overrule that it is possible to prove Theorem IV.1 using the construction specified in this failed attempt, i.e., we do not know whether the exists a $\mathrm{poly}(m)$-communication truthful auction for the classes of valuations considered above.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

rectangle argument" such as Lemma III.4; still, our transcript-counting arguments follow much of the intuition provided by Lemma III.4. For the full argument, see Section IV.

IV. Main Result: A Separation for Three Bidders

We now state and prove the main results of our paper: a separation for three-bidder truthful auction vs. non-truthful protocols. First, we state the lower bound for poly-communication truthful auctions with bidders in $\mathsf{SubAdd} \cup \mathsf{SingleM}$:

Theorem IV.1. *Every deterministic truthful mechanism $\mathcal{A}$ for three bidders with valuations in $\mathsf{SubAdd}_m \cup \mathsf{SingleM}_m$ that guarantees a $(\frac{\sqrt{3}-1}{2} + \frac{12}{\log m})$-approximation to the optimal social welfare requires $\exp\left(\Omega(\frac{\sqrt{m}}{\log m})\right)$ bits of communication.*

We prove this lower bound below. To establish a separation, we must also give an upper bound for non-truthful protocols. We prove this next, i.e., we construct a $\mathrm{poly}(m)$-communication protocol getting a good welfare approximation. This protocol is a simple extension of a result of [Fei09]:

Theorem IV.2 ([Fei09])**.** *For any number of bidders n and items m, there is a deterministic, $\mathrm{poly}(m,n)$-communication (non-truthful) protocol that guarantees a $\frac{1}{2}$-approximation to the optimal social welfare for bidders with valuations in SubAdd_m.*

Applying the reduction in Lemma II.3 to the communication efficient protocol in Theorem IV.2 gives the following:

Corollary IV.3. *For any number of bidders $n = O(\log m)$, there is a deterministic $\mathrm{poly}(m)$-communication (non-truthful) protocol that guarantees a $\frac{1}{2}$-approximation to the optimal social welfare for bidders with valuations in $\mathsf{SubAdd}_m \cup \mathsf{SingleM}_m$.*

Together, Theorem IV.1 and Corollary IV.3 give the main result of our paper: truthful mechanisms are provably less powerful than non-truthful protocols for bidders with valuations in $\mathsf{SubAdd} \cup \mathsf{SingleM}$. Observe that the upper bound, and thus the separation, holds for every constant $n \geq 3$, and indeed even for every number of bidders that is logarithmic in the number of items.[16]

We are now ready to prove Theorem IV.1; this proof occupies the remainder of Section IV.

A. Proof Outline

As discussed in Section III, our proof proceeds as follows. To begin, we construct a class of subadditive valuations, generalizing [EFN+19] and our construction in Proposition III.1. Then, the proof of our communication lower bound proceeds in two major steps. First, we show that getting a good welfare approximation implies that the auction must always avoid situations we term "bad 4-tuples" (see below for a definition). Second, we show that any protocol which always avoids bad 4-tuples requires high communication.

B. Description of the class of valuations

Our construction is a significant generalization the construction of [EFN+19]. Our goal is to have a set of subadditive valuations of two bidders, whom we name Bob and Charlie, and a collection of subsets $\mathcal{G} = \{G_1, \ldots, G_k\} \subseteq 2^{[m]}$ for a large enough k such that for each $i \in [k]$:

(1) For every $i \in [k]$, we can find valuations of Bob and Charlie that jointly have high value for G_i and low value for $\overline{G}_i$. Alternatively, we can find valuations of Bob and Charlie that have low value for G_i and high value for $\overline{G}_i$.
(2) For every $i \in [k]$ and $S \in \{G_i, \overline{G}_i\}$, even if we know that Bob and Charlie have high joint value for S, dividing the items in S among Bob and Charlie in an approximately optimal way takes exponentially-many bits of communication.

The construction in Proposition III.1 meets condition (1).[17] For condition (2) to hold, we build the following two-layered collection of subsets:

Definition IV.4. A *k-width-family* $\mathcal{F} = (\mathcal{G}, \{\mathcal{H}_i^{(t)}\}_{t\in\{0,1\}, i\in[k]})$ of $[m]$ consists of $2k+1$ collections $\mathcal{G}, \mathcal{H}_1^{(0)}, \ldots, \mathcal{H}_k^{(0)}, \mathcal{H}_1^{(1)}, \ldots, \mathcal{H}_k^{(1)} \subseteq 2^{[m]}$ of size k each, where we denote $\mathcal{G} = \{G_1, \ldots, G_k\}$ and $\mathcal{H}_i^{(t)} = \{H_{i,1}^{(t)}, \ldots, H_{i,k}^{(t)}\}$ for each $t \in \{0,1\}, i \in [k]$ such that for all $i, j \in [k]$, $H_{i,j}^{(0)} \subseteq \overline{G_i}$ and $H_{i,j}^{(1)} \subseteq G_i$.

For every k-width-family $\mathcal{F}$, a vector $\boldsymbol{b} \in \{0,1\}^k$, and a $k \times k$ binary matrix $\boldsymbol{C} \in \{0,1\}^{k\times k}$, denote with $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}]$ the collection of k^2 subsets $\{S_{i,j}\}_{i,j\in[k]}$, where for each $i, j \in [k]$,

$$S_{i,j} = \begin{cases} G_i \cup H_{i,j}^{(0)} & \boldsymbol{b}[i] = 0, \boldsymbol{C}[i][j] = 0, \\ G_i \cup (\overline{G_i} \setminus H_{i,j}^{(0)}) & \boldsymbol{b}[i] = 0, \boldsymbol{C}[i][j] = 1, \\ \overline{G_i} \cup H_{i,j}^{(1)} & \boldsymbol{b}[i] = 1, \boldsymbol{C}[i][j] = 0, \\ \overline{G_i} \cup (G_i \setminus H_{i,j}^{(1)}) & \boldsymbol{b}[i] = 1, \boldsymbol{C}[i][j] = 1. \end{cases}$$

Recall the following property for a collection (as defined in [EFN+19]):

Definition IV.5. A collection $\mathcal{S}$ of subsets in $[m]$ is called *ℓ-sparse* if for all $T_1, T_2, \ldots, T_{\ell-1} \in \mathcal{S}$, $\bigcup_{j=1}^{\ell-1} T_j \neq [m]$.

In words, an ℓ-sparse collection requires at least ℓ elements to cover all of $[m]$. We will now use ℓ-sparseness to define another property (mildly generalizing the corresponding property in [EFN+19]):

[16] Observe that a trivial second-price auction on the grand bundle among the three bidders in $\mathsf{SubAdd} \cup \mathsf{SingleM}$ that gives at least $\frac{1}{3}$ of the optimal welfare, is truthful and communication-efficient, whilst by Theorem IV.1, no truthful and communication efficient has approximation better than ≈ 0.366. Hence, there remains a small gap in understanding the achievable approximation ratio for truthful mechanisms in this class.

[17] In more detail, Proposition III.1 constructs a large set of 4-cell partitions $\{\{\mathcal{G}_{i,1}, \mathcal{G}_{i,2}, \mathcal{G}_{i,3}, \mathcal{G}_{i,4}\} \mid i \in [k]\}$, and a set of valuations of Bob and Charlie, such that the set family $\{\mathcal{G}_{i,1} \cup \mathcal{G}_{i,2} \mid i \in [k]\}$ meets condition (1) above. Our construction extends the one in Proposition III.1 by replacing the index $j \in [4]$ with another family of sets $\{H_j \mid j \in [k]\}$ in order to meet condition (2).

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

Definition IV.6. A k-width-family is *ℓ-independent* if for all $\boldsymbol{b} \in \{0,1\}^k, \boldsymbol{C} \in \{0,1\}^{k \times k}$, the collection $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}]$ is ℓ-sparse.

For exposition, we write an example of a 2-independent 2-width-family in Section C. By using the probabilistic method, we show that there exists such a family for large enough values of k and ℓ:

Lemma IV.7. *For sufficiently large m, there exists an ℓ-independent k-width-family, where $\ell = \frac{1}{4}\log m$ and $k = \exp\left(\frac{2\sqrt{m}}{\log m}\right)$.*

The proof of Lemma IV.7 is in Section A; the proof uses a (somewhat involved, but standard) application of the probabilistic method. In accordance with [EFN+19], we translate ℓ-sparse collections into modified set-cover valuations, denoted as $v^{\ell}_{\mathcal{S}}(\cdot)$. This translation is exactly as in [EFN+19], but we include it for completeness. Given an ℓ-sparse collection $\mathcal{S} = \{S_1, \ldots, S_d\}$, for every $X \subseteq M$:

$$\sigma_{\mathcal{S}}(X) = \begin{cases} \min\left\{|Y| \;\middle|\; Y \subseteq [d],\ X \subseteq \bigcup_{i \in Y} S_i\right\} & \text{if } X \subseteq \bigcup_{i \in [d]} S_i \\ \max\{\ell, d\} & \text{otherwise.} \end{cases}$$

In words, $\sigma_{\mathcal{S}}(X)$ is the minimal number of subsets of $\mathcal{S}$ that cover X, if it exists, and otherwise it is a "large" number. We now use $\sigma_{\mathcal{S}}$ to define a valuation function $v^{\ell}_{\mathcal{S}}$. For every subset of items $X \subseteq M$, if $\phi_{\mathcal{S}}(X) < \frac{\ell}{2}$, set $v^{\ell}_{\mathcal{S}}(X) = \sigma_{\mathcal{S}}(X)$ and $v^{\ell}_{\mathcal{S}}(\overline{X}) = \ell - \sigma_{\mathcal{S}}(X)$. For every subset that remains undefined, let $v^{\ell}_{\mathcal{S}}(\overline{X}) = \frac{\ell}{2}$. In principle, it is possible that $v^{\ell}_{\mathcal{S}}$ is not well defined. However, [EFN+19] shows that:

Lemma IV.8 ([EFN+19])**.** *For any ℓ-sparse collection $\mathcal{S}$ and $v^{\ell}_{\mathcal{S}}(\cdot)$:*

1) *$v^{\ell}_{\mathcal{S}}(\cdot)$ is well defined.*
2) *$v^{\ell}_{\mathcal{S}}(\cdot)$ is monotone, normalized and subadditive.*
3) *For every $S \subseteq [m]$, $v^{\ell}_{\mathcal{S}}(S) + v^{\ell}_{\mathcal{S}}(\overline{S}) = \ell$.*
4) *For every $S \in \mathcal{S}$, $v^{\ell}_{\mathcal{S}}(S) = 1$ and $v^{\ell}_{\mathcal{S}}(\overline{S}) = \ell - 1$.*

We note one counter-intuitive aspect of the valuations defined by $\mathcal{S}$: the valuation $v^{\ell}_{\mathcal{S}}$ has "low" value for the bundles S that are in the collection $\mathcal{S}$.

We are now ready to define the class of valuations for which we show hardness. Given an ℓ-independent k-width-family $\mathcal{F} = (\mathcal{G}, \{\mathcal{H}^{(t)}_i\}_{t \in \{0,1\}, i \in [k]})$ and denote $\mathcal{G} = \{G_1, \ldots, G_k\}$ and $\mathcal{H}^{(t)}_i = \{H^{(t)}_{i,1}, \ldots, H^{(t)}_{i,k}\}$ for each $t \in \{0,1\}, i \in [k]$, we define the following sub-class of subadditive valuations:

$$\mathsf{SubAdd}^* = \left\{v^{\ell}_{\mathcal{F}[\boldsymbol{b},\boldsymbol{C}]}(\cdot) \mid \boldsymbol{b} \in \{0,1\}^k, \boldsymbol{C} \in \{0,1\}^{k \times k}\right\} \subseteq \mathsf{SubAdd}_m.$$

The following claim lists several properties of the valuation functions in SubAdd^*, which we will use in Section IV-C:

Claim IV.9. *Fix a vector $\boldsymbol{b} \in \{0,1\}^k$ and a matrix $\boldsymbol{C} \in \{0,1\}^{k \times k}$, and let $v(\cdot)$ be the valuation parameterized by the collection $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}]$. Then, for every pair of indices $i, j \in [k]$:*

1) *If $\boldsymbol{b}[i] = 0$, then $v(G_i) = 1$.*
2) *If $\boldsymbol{b}[i] = 0$ and $\boldsymbol{C}[i][j] = 0$, then $v(\overline{G}_i \setminus H^{(0)}_{i,j}) = \ell - 1$.*
3) *If $\boldsymbol{b}[i] = 0$ and $\boldsymbol{C}[i][j] = 1$, then $v(H^{(0)}_{i,j}) = \ell - 1$.*
4) *If $\boldsymbol{b}[i] = 1$ and $\boldsymbol{C}[i][j] = 0$, then $v(G_i \setminus H^{(1)}_{i,j}) = \ell - 1$.*
5) *If $\boldsymbol{b}[i] = 1$ and $\boldsymbol{C}[i][j] = 1$, then $v(H^{(1)}_{i,j}) = \ell - 1$.*

We skip the proof of Claim IV.9, as it directly follows from the construction and Lemma IV.8.

The class of single-minded valuations $\mathsf{SingleM}^* \subseteq \mathsf{SingleM}_m$ satisfies that each valuation is parameterized by an index $i \in [k]$ and $\delta \in \{0,1\}$ such that:

$$\forall X \subseteq M, \quad v_{i,\delta}(X) = \begin{cases} (\sqrt{3}-1)\ell + \delta & \overline{G}_i \subseteq X, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

C. Getting good welfare implies "no bad 4-tuple"

Now, fix a truthful communication-efficient three-player mechanism $\mathcal{A}$ that gives an approximation strictly better than $\frac{\sqrt{3}-1}{2} + \frac{3}{\ell}$ when all the bidders' valuation classes are equal to $\mathsf{SubAdd}_m \cup \mathsf{SingleM}_m$, where $\ell = \frac{1}{4}\log m$. For convenience, we name the bidders Alice, Bob and Charlie.

The auction rule $\mathcal{A}$ defines two other communication problems which we reason about and use to define a notion of "bad 4-tuples". These two communication problems are to calculate the social welfare of $\mathcal{A}$ and to calculate the menu in $\mathcal{A}$. Let $\mathcal{SW}(v_A, v_B, v_C)$ be the communication problem between all three bidders that computes the social welfare that $\mathcal{A}$ gets with valuations (v_A, v_B, v_C), and denote the transcript function of its most efficient protocol by $\tau_{\mathcal{SW}}(v_A, v_B, v_C)$. Note that $\mathrm{cc}(\mathcal{SW}) \leq \mathrm{cc}(\mathcal{A}) + \mathrm{poly}(m)$.[18] Also, let $\mathcal{P}(v_B, v_C)$ be the communication problem between Bob and Charlie that finds the menu $\mathsf{Menu}_A(v_B, v_C)$ for Alice defined by the auction $\mathcal{A}$, and denote the transcript function of its most efficient protocol by $\tau_{\mathcal{P}}(v_B, v_C)$.

We now define 4-bad tuples and show that they cannot exist when $\mathcal{A}$ gets a good approximation to the optimal welfare.

Proposition IV.10. *Let $v^{(1)}, v^{(2)}, v^{(3)}, v^{(4)}$ be valuations in SubAdd^*, parameterized by the vectors $\boldsymbol{b}^{(1)}, \boldsymbol{b}^{(2)}, \boldsymbol{b}^{(3)}, \boldsymbol{b}^{(4)} \in \{0,1\}^k$ and the matrices $\boldsymbol{C}^{(1)}, \boldsymbol{C}^{(2)}, \boldsymbol{C}^{(3)}, \boldsymbol{C}^{(4)} \in \{0,1\}^{k \times k}$, respectively. We say that $\{(\boldsymbol{b}^{(j)}, \boldsymbol{C}^{(j)})\}_{j \in [4]}$ is a* bad 4-tuple *if there exist $i^*, j^*_1, j^*_2 \in [k]$ such that the five following conditions hold simultaneously:*

1) *$\boldsymbol{b}^{(1)}[i^*] = \boldsymbol{b}^{(2)}[i^*] = 0$ and $\boldsymbol{b}^{(3)}[i^*] = \boldsymbol{b}^{(4)}[i^*] = 1$.*
2) *$\boldsymbol{C}^{(1)}[i^*][j^*_1] = 0$ and $\boldsymbol{C}^{(2)}[i^*][j^*_1] = 1$.*
3) *$\boldsymbol{C}^{(3)}[i^*][j^*_2] = 0$ and $\boldsymbol{C}^{(4)}[i^*][j^*_2] = 1$.*
4) *$\tau_{\mathcal{P}}(v^{(1)}, v^{(1)}) = \tau_{\mathcal{P}}(v^{(2)}, v^{(2)}) = \tau_{\mathcal{P}}(v^{(3)}, v^{(3)}) = \tau_{\mathcal{P}}(v^{(4)}, v^{(4)})$, i.e. the transcript of the protocol that computes the menu is identical for all these pairs.*
5) *$\tau_{\mathcal{SW}}(v_{i^*,0}, v^{(3)}, x^{(3)}) = \tau_{\mathcal{SW}}(v_{i^*,0}, v^{(4)}, x^{(4)})$, i.e. the transcript of the protocol that computes the social welfare of $\mathcal{A}$ is identical for both of these valuation profiles.*[19]

[18] This follows because one can always calculate $\mathcal{SW}$ by simply calculating $\mathcal{A}$, then having each bidder send their value for the set of items they receive (and we assume that bidders' valuations use precision $\mathrm{poly}(m)$, as discussed in Section II).

[19] We remind that $v_{i^*,0}$ is the single-minded valuation defined in Equation 1.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

Then, there cannot exist any bad 4-tuples (i.e., for all such $\{(\boldsymbol{b}^{(j)}, \boldsymbol{C}^{(j)})\}_{j\in[4]}$, there cannot exist $i^*, j_1^*, j_2^* \in [k]$ such that all five above conditions hold).

Informally, this means that for every "sufficiently different" elements of the valuation family SubAdd^*, their transcript either in the protocol of menu computation of $\mathcal{A}$ or the protocol that computes the welfare of $\mathcal{A}$ should differ.

Proof of Proposition IV.10. We will show that if all conditions hold simultaneously for some indices i^*, j_1^*, j_2^*, then the approximation ratio must be strictly smaller than $\frac{\sqrt{3}-1}{2} + \frac{3}{\ell}$, which leads to a contradiction. We do that by showing that in particular, the mechanism outputs a bad approximation for one of the valuation profiles $(v_{i^*,1}, v^{(1)}, v^{(2)})$ or $(v_{i^*,0}, v^{(3)}, v^{(4)})$. Note that by condition 4 and by the rectangle argument, $\tau_{\mathcal{P}}(v^{(1)}, v^{(2)}) = \tau_{\mathcal{P}}(v^{(3)}, v^{(4)})$, so in particular the menus presented to Alice given the valuations $(v^{(1)}, v^{(2)})$ and $(v^{(3)}, v^{(4)})$ are identical. Denote this menu with $M_A = \mathsf{Menu}_A^{\mathcal{A}}(v^{(1)}, v^{(2)}) = \mathsf{Menu}_A^{\mathcal{A}}(v^{(3)}, v^{(4)})$ and recall that $M_A(S)$ is the price in the menu for any bundle $S \subseteq [m]$. Note that by Proposition II.6, we can assume that all prices are non-negative. We complete the proof by case analysis.

Note that if $M_A(\overline{G}_{i^*}) \leq (\sqrt{3}-1)\ell$, then given the valuation profile $(v_{i^*,1}, v^{(1)}, v^{(2)})$, Alice wins $\overline{G}_{i^*}$ because all prices are non-negative, so this is the only valuable bundle for her. However, since both valuations $v^{(1)}, v^{(2)}$ are parameterized by $\boldsymbol{b}^{(1)}, \boldsymbol{b}^{(2)}$ such that $\boldsymbol{b}^{(1)}[i^*] = \boldsymbol{b}^{(2)}[i^*]$, by Property 1 in Claim IV.9, we have that $v^{(1)}(G_{i^*}) = v^{(2)}(G_{i^*}) = 1$. As a result, the total welfare of the allocation of $\mathcal{A}$ is at most $((\sqrt{3}-1)\ell + 1) + 2$. However, the optimal welfare is at least $2(\ell-1)$, obtained by giving $\overline{G}_{i^*} \setminus H^{(0)}_{i^*,j_1^*}$ to Bob and $H^{(0)}_{i^*,j_1^*}$ to Charlie (according to Property 2 and 3 of Claim IV.9).

However, if $M_A(\overline{G}_{i^*}) > (\sqrt{3}-1)\ell$, we claim it still implies a failure on another valuation profile $(v_{i^*,0}, v^{(3)}, v^{(4)})$. For that, we analyze the output of the mechanism $\mathcal{A}$ on $(v_{i^*,0}, v^{(3)}, v^{(3)})$.

Since the price of $\overline{G}_{i^*}$ that Bob and Charlie present to Alice is is $M_A(\overline{G}_{i^*})$ given $(v^{(3)}, v^{(3)})$, she does not win $\overline{G}_{i^*}$, resulting a welfare of zero. In addition, since Bob and Charlie's valuations are identical, the total welfare from them is at most ℓ by Condition 3 of Lemma IV.8. Therefore, the social welfare that $\mathcal{A}$ obtains on $(v_{i^*,0}, v^{(3)}, v^{(3)})$ is at most ℓ. Recall that by assumption, $(v_{i^*,0}, v^{(3)}, v^{(3)})$ and $(v_{i^*,0}, v^{(4)}, v^{(4)})$ have the same transcript in the protocol that computes the social welfare. By the rectangle argument, $(v_{i^*,0}, v^{(3)}, v^{(4)})$ must also have the same transcript, and thus the welfare of the allocation from $\mathcal{A}$ in this case is also at most ℓ. However, the optimal welfare from $(v_{i^*,0}, v^{(3)}, v^{(4)})$ is at least $(\sqrt{3}-1)\ell + 2(\ell-1)$, obtained by giving $\overline{G}_{i^*}$ to Alice, $G_{i^*} \setminus H^{(1)}_{i^*,j_2^*}$ to Bob, and $H^{(1)}_{i^*,j_2^*}$ to Charlie.

Thus, the best possible approximation ratio of the mechanism $\mathcal{A}$ is at most:

$$\max\left\{\frac{((\sqrt{3}-1)\ell+1)+2}{2(\ell-1)}, \frac{\ell}{(\sqrt{3}-1)\ell+2(\ell-1)}\right\} < \frac{\sqrt{3}-1}{2} + \frac{3}{\ell}. \quad \text{(for } \ell \geq 3\text{)}$$

By that, we get a contradiction, which completes the proof. □

D. "No bad 4-tuple" implies high communication

We are now ready to finalize the proof of Theorem IV.1. Note that by Proposition IV.10, a communication-efficient truthful mechanism with approximation ratio no worse than $\frac{\sqrt{3}-1}{2} + \frac{3}{\ell}$ cannot have a "bad 4-tuple". We will now show $e^{\Omega(\frac{\sqrt{m}}{\log m})}$ bits of communication is necessary for the mechanism $\mathcal{A}$ to avoid such "bad 4-tuples".

For any fixed transcript T for $\mathcal{P}$, let $C(T)$ be the set of pairs $(\boldsymbol{b}, \boldsymbol{C}) \in \{0,1\}^k \times \{0,1\}^{k\times k}$ such that $\mathcal{P}$ uses transcript T on $(v_{\boldsymbol{b},\boldsymbol{C}}, v_{\boldsymbol{b},\boldsymbol{C}})$, i.e., $C(T) = \{(\boldsymbol{b}, \boldsymbol{C}) : \tau_{\mathcal{P}}(v_{\boldsymbol{b},\boldsymbol{C}}, v_{\boldsymbol{b},\boldsymbol{C}}) = T\}$. For any index $i \in [k]$ and bit $z \in \{0,1\}$, we further define for each transcript the following subsets of length-k bit-vectors:

$$C_{i,z}(T) = \{\boldsymbol{C}[i] : (\boldsymbol{b}, \boldsymbol{C}) \in C(T), \boldsymbol{b}[i] = z\} \subseteq \{0,1\}^k.$$

In other words, for every transcript T of the the menu computation protocol $\mathcal{P}$, the subset $C_{i,z}(T)$ contains for every $(\boldsymbol{b}, \boldsymbol{C}) \in C(T)$, the i'th rows of $\boldsymbol{C}$ only for cases in which $\boldsymbol{b}[i] = z$. Note that $C_{i,z}(T)$ may be empty.

We now claim that for every transcript T of the menu computation protocol, and for every index i, it holds that $|C_{i,0}(T)| \leq 1$ or $|C_{i,1}(T)| \leq 2^{\mathrm{cc}(\mathcal{SW})}$. The reason for it is that if the converse holds, i.e. there exists an index i^* such that both $|C_{i^*,0}(T)| > 1$ and $|C_{i^*,1}(T)| > 2^{\mathrm{cc}(\mathcal{SW})}$, then there is a bad 4-tuple.[20]

Note that $|C_{i,0}(T)| \leq 1$ or $|C_{i,1}(T)| \leq 2^{\mathrm{cc}(\mathcal{SW})}$ implies that:

$$|C_{i,0}(T)| + |C_{i,1}(T)| \leq 2^k + 2^{\mathrm{cc}(\mathcal{SW})}. \tag{2}$$

The explanation for it is by a simple case analysis. If $|C_{i,0}(T)| \leq 1$, then:

$$|C_{i,0}(T)| + |C_{i,1}(T)| \leq 1 + 2^k \leq 2^{\mathrm{cc}(\mathcal{SW})} + 2^k.$$

If $|C_{i,1}(T)| \leq 2^{\mathrm{cc}(\mathcal{SW})}$, then similarly:

$$|C_{i,0}(T)| + |C_{i,1}(T)| \leq 2^k + 2^{\mathrm{cc}(\mathcal{SW})}.$$

By definition and by plugging in Equation 2, we have that:

$$|C(T)| \leq \prod_{i\in[k]} (|C_{i,0}(T)| + |C_{i,1}(T)|) \leq (2^k + 2^{\mathrm{cc}(\mathcal{SW})})^k. \tag{3}$$

[20] To see why, observe that by the pigeonhole principle, if both $|C_{i^*,0}(T)|$ and $|C_{i^*,1}(T)|$ are greater than 1, then there exist $(\boldsymbol{b}^{(1)}, \boldsymbol{C}^{(1)}), (\boldsymbol{b}^{(2)}, \boldsymbol{C}^{(2)}), (\boldsymbol{b}^{(3)}, \boldsymbol{C}^{(3)}), (\boldsymbol{b}^{(4)}, \boldsymbol{C}^{(4)})$ that satisfy Property 1, 2, 3, and 4 of Proposition IV.10 for this i^* and for some j_1^*, j_2^*, since they all share transcript T. Furthermore, the fact that $|C_{i^*,1}(T)|$ is greater than the number of transcripts for the protocol that computes the social welfare, implies that there are two vectors in $C_{i^*,1}(T)$ that also share the same transcript for the protocol $\mathcal{SW}$.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

Denote with $\mathcal{T}$ the subset of transcripts in the menu protocol, i.e. $\mathcal{T} = \{\tau(v_B, v_C) : (v_B, v_C) \in \mathsf{SubAdd}^*)\}$. Observe that:

$$\begin{aligned} 2^{k+k^2} &= \sum_{T \in \mathcal{T}} C(T) \\ &\qquad \text{(the total number of pairs } (\boldsymbol{b}, \boldsymbol{C}) \text{ is } 2^{k+k^2}) \\ &\leq |\mathcal{T}| \max_{T \in \mathcal{T}} C(T) \\ &\leq 2^{\mathrm{cc}(\mathcal{P})} \cdot (2^k + 2^{\mathrm{cc}(\mathcal{SW})})^k. \qquad \text{(by Equation 3)} \end{aligned}$$

Therefore, we have:

$$\begin{aligned} 2^{\mathrm{cc}(\mathcal{P})} &\geq \frac{2^{k+k^2}}{(2^k + 2^{\mathrm{cc}(\mathcal{SW})})^k} \\ &= \left(\frac{2^{1+k}}{2^k + 2^{\mathrm{cc}(\mathcal{SW})}}\right)^k = \left(\frac{2}{1 + 2^{\mathrm{cc}(\mathcal{SW})-k}}\right)^k. \end{aligned} \tag{4}$$

To conclude the proof, it suffices to show that $\frac{2}{1+2^{\mathrm{cc}(\mathcal{SW})-k}} > B$ for some constant $B > 1$. We can show this as follows. Consider we use an ℓ-independent and k-width family with $k = e^{\frac{2\sqrt{m}}{\log m}}, \ell = \frac{1}{4}\log m$, which is guaranteed to exist by Lemma IV.7. Let $d = \mathrm{cc}(\mathcal{SW}) - \mathrm{cc}(\mathcal{A})$. We remind that d is at most polynomial in m. Observe that if $\mathrm{cc}(\mathcal{A}) \geq k - d - 1$, then $\mathrm{cc}(\mathcal{A}) = \exp(\Omega(\frac{2\sqrt{m}}{\log m}))$, and we are done. Thus, we can assume that $\mathrm{cc}(\mathcal{A}) < k - d - 1$, which implies that $\mathrm{cc}(\mathcal{SW}) \leq k - 1$. Therefore, a simple computation gives that:

$$\frac{2}{1 + 2^{\mathrm{cc}(\mathcal{SW})-k}} \geq \frac{4}{3}. \tag{5}$$

Combining Equation 4 and Equation 5 gives that $\mathrm{cc}(\mathcal{P}) \geq k \log \frac{4}{3}$. Thus, applying Theorem II.7 gives that $\mathrm{poly}(\mathrm{cc}(\mathcal{A}), m) \geq k \log \frac{4}{3}$, so $\mathrm{cc}(\mathcal{A}) \geq e^{\Omega(\frac{\sqrt{m}}{\log m})}$. This completes the proof that $\mathrm{cc}(\mathcal{A})$ is exponentially high.

Remark IV.11. A discerning reader might have observed that in the construction utilized in the proof of Theorem IV.1, Alice's valuation set comprises solely single-minded valuations, whereas Bob's and Charlie's valuation sets consist exclusively of subadditive valuations. Consequently, one might be inclined to argue that we demonstrate an impossibility for the scenario featuring one single-minded bidder and two subadditive bidders, i.e., for the valuation class $\mathsf{SingleM} \times \mathsf{SubAdd} \times \mathsf{SubAdd}$. However, this assertion is not accurate, and the lower bound and gap that we show are for the class $(\mathsf{SingleM} \cup \mathsf{SubAdd})^3$. The rationale behind this lies in the taxation framework which, as we restate in Theorem II.7, necessitates all bidders' valuation classes to be "sufficiently rich" (in this case, the classes must include subadditive valuations).

V. Supplemental Results: Lower Bounds on Two-Bidder Mechanisms

In this section, we give an additional impossibility result for two bidders which holds when bidders might be single-minded. In fact, the results of this section served as stepping stones towards our approach in Sections III and IV.

In contrast to the three-bidder hardness in Sections III and IV that requires showing hardness of computing the menu which was more involved, here we use the taxation framework in a more direct way. We do not need to show that computing the menu is hard.[21] We prove our impossibility result by showing that the number of menus in every truthful auction has to be high. Interestingly, our bound can also be achieved by arguing about the hardness of simultaneous protocols. We believe that due to the relative simplicity of the construction, this section serves as a way to learn about the structure of mechanisms for bidders both with and without complementarities. We opted to defer this discussion until this section to keep Sections III and IV self-contained.

We show that for two bidders with valuations in $\mathsf{XOS} \cup \mathsf{SingleM}$, any deterministic truthful mechanism that gives an approximation asymptotically better than $\frac{\sqrt{5}-1}{2}$ requires exponential communication.Formally:

Theorem V.1. *Every deterministic truthful mechanism $\mathcal{A}$ for two bidders with valuations in $\mathsf{XOS}_m \cup \mathsf{SingleM}_m$ that guarantees a $(\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}})$-approximation to the optimal social welfare requires $\exp(\Omega(m^{\frac{1}{3}}))$ bits of communication.*

Note that the deterministic protocol that $\frac{3}{4}$-approximates the optimal welfare for XOS [Fei09, DNS10] can be generalized to work for $\mathsf{XOS} \cup \mathsf{SingleM}$ by using Lemma II.3. Therefore, Theorem V.7 indeed implies a separation.

The proof of Theorem V.7 is based on a direct taxation argument, i.e. we show that the number of menus of every truthful mechanism that gives a "good" approximation has to be doubly exponential in m. To the best of our knowledge, this is not easy to do with the lower bound construction of [AKSW20].[22] We begin by describing the class of valuations that we use to show the separation (Section V-A). We then prove Theorem V.7 in Section V-B. In Section V-C, we explore the challenges associated with directly generalizing Theorem V.7 to a scenario involving three bidders. By that, we illustrate the challenges encountered when proving Theorem IV.1.

En route, we explore the structure of payments of approximately optimal and truthful mechanisms, showing that as the approximation guarantee of a mechanism goes to 1, then its payments approach VCG payments. We state this formally in Section B.

A. Description of the Class of Valuations

Recall the definition of XOS valuations:

Definition V.2 (XOS valuations)**.** A valuation function $v : 2^{[m]} \to \mathbb{R}$ is *XOS* if there exists a collection of additive clauses $\mathcal{C} \subseteq \mathbb{R}^m$ such that for all $S \in 2^{[m]}$, $v(S) = \max_{\boldsymbol{c} \in \mathcal{C}} \sum_{i \in S} c_i$. We denote the family of all XOS valuation functions over m items by $\mathsf{XOS} = \mathsf{XOS}_m$.

[21] Actually, there are only two players, so computing the menu is a one player problem which cannot possibly be hard.

[22] Theorem V.7 can also be proven by showing a lower bound on simultaneous protocols, which is the same argument used in [AKSW20]. See Remark V.9.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

An XOS valuation is *binary* if all its clauses are 0/1 valued, i.e., for every $c \in \mathcal{C}$, $c_i \in \{0,1\}$. Equivalently, binary XOS valuations can be defined by a collection of subsets of $[m]$:

Definition V.3 (Binary XOS valuations)**.** A valuation function $v : 2^{[m]} \to \mathbb{R}$ is *binary XOS* if there exists a collection $\mathcal{G} \subseteq 2^{[m]}$ such that $v(S) = \max_{G \in \mathcal{G}} \{|S \cap G|\}$ for all $S \in 2^{[m]}$. Denote the family of all binary XOS valuation functions over m items by $\mathsf{BXOS} = \mathsf{BXOS}_m \subset \mathsf{XOS}_m$.

To prove Theorem V.1, we will consider a special subset of XOS valuations, which takes sets from a collection that satisfies the *average intersection property*:

Definition V.4. A collection $\mathcal{G} \subseteq 2^{[m]}$ satisfies the *b-average intersection property* for some $b \in (0,1)$ if every $G \in \mathcal{G}$ is of size bm, and for every $G_1 \neq G_2 \in \mathcal{G}$, $|G_1 \cap G_2| \leq 2b^2 m$.

By the probabilistic method, such collections always exist and can be made exponentially large. Formally:

Lemma V.5. *For any $b \in (0,1)$, there exists a collection $\mathcal{G} \subseteq 2^{[m]}$ of size $\exp(b^2 m/6)$ that satisfies the b-average intersection property.*

The proof of Lemma V.5 can be found in Section A4. Let $b = m^{-\frac{1}{3}}$ and $\mathcal{G}$ be such a collection containing $\exp(m^{\frac{1}{3}}/6)$ subsets of $[m]$ that satisfies $m^{-\frac{1}{3}}$-average intersection property. We define a special set of binary XOS valuations $\mathsf{XOS}^* = \{v_{\mathcal{H}} : \mathcal{H} \subseteq \mathcal{G}\} \subseteq \mathsf{BXOS}_m$ where

$$v_{\mathcal{H}}(S) = \max_{H \in \mathcal{H}} |S \cap H| \quad \text{for all } S \in 2^{[m]}.$$

Meanwhile, we consider the class of single-minded valuations $\mathsf{SingleM}^* \subseteq \mathsf{SingleM}_m$ containing single-minded valuations with constant value α, i.e., $\mathsf{SingleM}^* = \{v_{T,\delta} : T \subseteq [m], \delta \in \{0,1\}\}$ where

$$v_{T,\delta}(S) = (\alpha + \delta) \cdot \mathrm{I}[S \supseteq T] \quad \text{for all } S \in 2^{[m]}.$$

and α is a constant that we fix later.

Later, the valuation profiles $\{(v_A, v_B) : v_A \in \mathsf{SingleM}^*, v_B \in \mathsf{XOS}^*\}$ will be used to show a lower bound on the communication complexity.

B. Proof of Theorem V.1 via Taxation Complexity

In this section, we will present a proof of Theorem V.1 based on the Taxation Complexity framework [Dob16b].

Definition V.6 (Taxation complexity)**.** Let $\mathcal{A} = (f, p_1, \ldots, p_n) : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \Sigma \times \mathbb{R}^n$ be any deterministic truthful mechanism. The *taxation complexity* $\mathrm{tax}(\mathcal{A})$ of $\mathcal{A}$ is defined as the number of bits needed to represent the index of a specific menu, i.e., $\mathrm{tax}(\mathcal{A}) = \max_{i \in [n]} \log |\{\mathsf{Menu}_i^{\mathcal{A}}(v_{-i}) : v_{-i} \in \mathcal{V}_{-i}\}|$.

Theorem V.7 ([Dob16b])**.** *Let $\mathcal{A} = (f, p_1, \ldots, p_n) : \mathcal{V}_1 \times \cdots \times \mathcal{V}_n \to \Sigma \times \mathbb{R}^n$ be any deterministic truthful mechanism, where all domains $\mathcal{V}_1, \ldots, \mathcal{V}_n$ contain XOS_m. Then, $\mathrm{cc}(\mathcal{A}) \geq \frac{\mathrm{tax}(\mathcal{A})}{m} - 1$.*

Observe that Theorem V.7 only requires that the domain of valuations contains XOS, and it remains applicable for mechanisms that are truthful for a larger domain, e.g., $\mathsf{XOS} \cup \mathsf{SingleM}$. Therefore, our first step in the proof Theorem V.1 will be to show a lower bound on the taxation complexity of every mechanism that gives a "good" approximation. Formally:

Proposition V.8. *Let $\mathcal{V} = \mathsf{XOS}_m \cup \mathsf{SingleM}_m$ and $\mathcal{A} = (f, p_A, p_B) : \mathcal{V} \times \mathcal{V} \to \Sigma \times \mathbb{R}^2$ be any deterministic truthful mechanism. If $\mathcal{A}$ obtains a $(\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}})$-approximation to the optimal welfare, then $\mathrm{tax}(\mathcal{A}) \geq \exp(m^{\frac{1}{3}}/6)$.*

Proof of Proposition V.8. Suppose that the mechanism $\mathcal{A}$ indeed obtains an approximation to the welfare of at least $(\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}})$. Consider the case when Alice's valuation is $v_A \in \mathsf{SingleM}^*$ and Bob's valuation $v_B \in \mathsf{XOS}^*$. We want to show that for every distinct pair of valuations of Bob $v_B, v_B' \in \mathsf{XOS}^*$, the menus that are presented to Alice must be different, i.e., $\mathsf{Menu}_A^{\mathcal{A}}(v_B) \neq \mathsf{Menu}_A^{\mathcal{A}}(v_B')$. Then, $|\{\mathsf{Menu}_A^{\mathcal{A}}(v_B) : v_B \in \mathsf{XOS}^*\}| = |\mathsf{XOS}^*| = 2^{|\mathcal{G}|}$ and the proposition follows. For simplicity, from now on we denote $\mathsf{Menu}_A^{\mathcal{A}}(\cdot)$ with $\mathsf{Menu}(\cdot)$.

For any distinct $v_B, v_B' \in \mathsf{XOS}^*$, denote with $\mathcal{H}$ and with $\mathcal{H}'$ respectively the combinations that are associated with v_B and with v_B' respectively. Since $v_B \neq v_B'$, there exists at least one subset H such that (without loss of generality) $H \in \mathcal{H}$, whereas $H \notin \mathcal{H}'$. Suppose $\mathsf{Menu}(v_B) = \mathsf{Menu}(v_B') = M$ by contradiction. There are two possibilities:

- If $M(\overline{H}) \leq \alpha$, then given the valuation profile $(v_A = v_{\overline{H},1}, v_B')$, Alice always gets $\overline{H}$. The actual welfare the mechanism $\mathcal{A}$ gets is at most $v_A(\overline{H}) + v_B'(H) \leq (\alpha+1) + 2b^2 m$, while the optimal welfare being at least $v_B'([m]) = bm$.
- If $M(\overline{H}) > \alpha$, then given the valuation profile $(v_A = v_{\overline{H},0}, v_B)$, Alice does not get $\overline{H}$. The actual welfare the mechanism $\mathcal{A}$ gets is at most $v_B([m]) = bm$, while the optimal welfare being $v_A(\overline{H}) + v_B(H) = \alpha + bm$.

To summarize, whenever $\mathsf{Menu}(v_B) = \mathsf{Menu}(v_B')$ for some distinct $v_B, v_B' \in \mathsf{XOS}^*$, we know that the approximation ratio of the mechanism to the optimal welfare can be at most

$$\max\left\{\frac{(\alpha+1) + 2b^2 m}{bm}, \frac{bm}{\alpha + bm}\right\}. \tag{6}$$

Recall $b = m^{-\frac{1}{3}}$ and let $\alpha = \frac{\sqrt{5}-1}{2} bm = \frac{\sqrt{5}-1}{2} m^{\frac{2}{3}}$. Then, the ratio above is always strictly smaller than $\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}}$. Therefore, to get a $(\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}})$-approximation, $\mathcal{A}$ must have different menus for all $v_B \in \mathsf{XOS}^*$. As a result, $\mathrm{tax}(\mathcal{A}) \geq \log |\mathsf{XOS}^*| = |\mathcal{G}| = \exp(m^{\frac{1}{3}}/6)$. □

Having Proposition V.8, Theorem V.1 follows directly via the Taxation Complexity framework.

Proof of Theorem V.1. Since $\mathsf{XOS}_m \cup \mathsf{SingleM}_m \supseteq \mathsf{XOS}_m$, we have $\mathrm{cc}(\mathcal{A}) \geq \frac{\mathrm{tax}(\mathcal{A})}{m} - 1$ by Theorem V.7. Further by Proposition V.8, we have $\mathrm{cc}(\mathcal{A}) \geq \exp(\Omega(m^{\frac{1}{3}}))$. □

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

Remark V.9. Theorem V.1 can be also proven via a communication lower bound for simultaneous protocols combining with the two-bidder corollary from Taxation Complexity framework [Dob16b], as used in [AKSW20]. The proof follows a similar structure, and we include it in Section A5 for completeness.

Remark V.10. Similarly to Remark IV.11, one might think that the construction implies that our impossibility holds for a single-minded bidders and a bidder with XOS valuations, i.e. for the class $\mathsf{SingleM} \times \mathsf{XOS}$. However, our use of Theorem V.7 implies that our the lower bound applies solely to the valuation class $(\mathsf{SingleM} \cup \mathsf{XOS})^2$. In particular, the VCG mechanism can be implemented with $\mathrm{poly}(m)$ communication for $\mathsf{SingleM} \times \mathsf{XOS}$.[23]

C. Difficulties for a Stronger Three-Bidder Lower Bound

Naturally, one might wonder whether the class of valuations defined in Section V-A can be generalized to a stronger (i.e., better than $\frac{\sqrt{5}-1}{2}$) lower bound beyond two bidders, since the Taxation Complexity framework extends beyond two bidders. We now describe an attempt to generalize our construction and briefly discuss why it is not likely to work.

Recall that in Section V-B, the subset of valuations $\{(v_A, v_B) : v_A \in \mathsf{SingleM}^*, v_B \in \mathsf{XOS}^*\}$ was used as hard instances for the taxation complexity lower bound, where XOS^* is defined by some collection $\mathcal{G}$ that satisfies the b-average intersection property. To get a stronger three-bidder lower bound, we add another bidder named Charlie, and consider the valuations $\{(v_A, v_B, v_C) : v_A \in \mathsf{SingleM}^*, v_B, v_C \in \mathsf{XOS}^*\}$. Our hope is to get a stronger lower bound by *the extra communication hardness from welfare maximization between these two XOS bidders*.

Our plan is to use a similar argument to the one used in the proof of Proposition V.8. For convenience, let us assume $\mathcal{G}$ consists of exponentially many random sets of size $bm/2$ for some constant $b \in (0, 1)$. Similar to the proof of Lemma V.5, the collection $\mathcal{G}$ satisfies $(b/2)$-average intersection property with high probability and its size is $\exp(m)$. Consider four valuations $v^{(1)}, v^{(2)}, v^{(3)}, v^{(4)}$ sampled uniformly at random from XOS^* and associated with the subcollections $\mathcal{H}^{(1)}, \mathcal{H}^{(2)}, \mathcal{H}^{(3)}, \mathcal{H}^{(4)} \subseteq \mathcal{G}$, respectively.

We argue without proof that with high probability, there exist subsets H_1, H_2, H_3, H_4 such that for every $i \in [4]$:

1) $H_i \in \mathcal{H}^{(i)}$.
2) For every $j \neq i$, $H_i \notin \mathcal{H}^{(j)}$.
3) $H_1 \cap H_2 = H_3 \cap H_4 = \emptyset$.

As a result:

- Given the valuation profile $(v_B = v^{(1)}, v_C = v^{(2)})$, the maximum welfare bm is obtained by allocating $H_1 \cup H_2$ to them.
- Given the valuation profile $(v_B = v^{(3)}, v_C = v^{(4)})$, the welfare obtained from allocating $H_1 \cup H_2$ to Bob and Charlie is at most $2b^2 m$, due to the $(b/2)$-average intersection property. Maximum welfare of bm can be obtained from allocating $H_3 \cup H_4$ to them.

Let $v_A \in \mathsf{SingleM}^*$ be the single-minded valuation that has value α for the set $[m] \setminus (H_1 \cup H_2)$. Assume towards a contradiction that $\mathsf{Menu}(v^{(1)}, v^{(2)}) = \mathsf{Menu}(v^{(3)}, v^{(4)})$ and denote with P the price in both menus for the bundle $[m] \setminus (H_1 \cup H_2)$. Assume that $P \neq \alpha$ (which is without loss of generality, since we can always slightly perturb the value of α). Note that it must be the case that either:

- If $P < \alpha$, then Alice with value v_A gets the bundle $[m] \setminus (H_1 \cup H_2)$: then for $(v_A, v^{(3)}, v^{(4)})$, the mechanism gets welfare of at most $\alpha + 2b^2 m$, while the optimal welfare is $v_B^{(3)}(H_3) + v^{(4)}(H_4) = bm$.
- If $P > \alpha$, then Alice with valuation v_A does not get the bundle $[m] \setminus (H_1 \cup H_2)$: then for $(v_A, v^{(1)}, v^{(2)})$, the mechanism gets at most $(b - b^2/4)m$ *due to the communication hardness of allocating* $[m]$ *between Bob and Charlie*[24], while the optimal welfare is $\alpha + bm$.

Overall, the approximation ratio to the optimal welfare can be at most

$$\max\left\{\frac{\alpha + 2b^2 m}{bm}, \frac{(b - b^2/4)m}{\alpha + bm}\right\},$$

which is quite very similar to Equation 6 in the proof of Theorem V.7. The only difference is the numerator of the second term gets smaller (from bm to $(b - b^2/4)m$) which is a significant difference when b is close to 1). However, the minimum is still obtained when $b \to 0$ and $\alpha = \frac{\sqrt{5}-1}{2} bm$, resulting in no improvement at all.

As a result, the additional difficulty of allocating items among two bidders instead of one does not help, at least for this specific class of valuations. The reason for it is that for our original argument that gets $\frac{\sqrt{5}-1}{2}$ for two bidders, we want $b \to 0$ to ensure the intersection of a random clauses of the XOS bidder with the complement of the single-minded bidder to be as small as possible. However, to get a significant gap in communication hardness for two XOS bidders, we want $b \to 1$ since the intersection of a random clauses of them needs to be large, so that one cannot get a good welfare between them easily. Nevertheless, this idea turned out to be useful for the proof of Theorem IV.1, by using subadditive modified set-cover valuations.

Acknowledgements

The authors thank Shahar Dobzinski for the helpful discussions throughout the duration of this work.

References

[AKS21] Sepehr Assadi, Thomas Kesselheim, and Sahil Singla. Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier. In Dániel Marx, editor,

[23] This is because the optimal allocation can be computed with $\mathrm{poly}(m)$ communication: Alice simply sends her desired set S together with its value, and Bob replies whether the optimal allocation is $(S, \overline{S})$ or $(\emptyset, M)$.

[24] The welfare of $(b - b^2/4)m$ is obtained by giving any set $S \in \mathcal{H}^{(1)}$ to Bob and the rest of the items to Charlie. The $(b/2)$-average intersection property implies that this allocation has welfare of at least $(b - b^2/4)m$. The approximation ratio $\frac{(b-b^2/4)m}{bm}$ is smallest when $b = 1$, which is the most difficult case for welfare maximization between Bob and Charlie in general.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021, pages 653–661. SIAM, 2021. 1, 4

[AKSW20] Sepehr Assadi, Hrishikesh Khandeparkar, Raghuvansh R. Saxena, and S. Matthew Weinberg. Separating the communication complexity of truthful and non-truthful combinatorial auctions. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 1073–1085, New York, NY, USA, 2020. Association for Computing Machinery. 1, 2, 3, 5, 12, 14, 19

[AS19] Sepehr Assadi and Sahil Singla. Exponentially improved truthful combinatorial auctions with submodular bidders. In *Proceedings of the Sixtieth Annual IEEE Foundations of Computer Science (FOCS)*, 2019. 4

[BMW18] Mark Braverman, Jieming Mao, and S Matthew Weinberg. On simultaneous two-player combinatorial auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2256–2273. SIAM, 2018. 2

[Cla71] Edward H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11(1):17–33, 1971. 1

[CTW20] Linda Cai, Clayton Thomas, and S. Matthew Weinberg. Implementation in advised strategies: Welfare guarantees from posted-price mechanisms when demand queries are np-hard. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pages 61:1–61:32. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. 4

[DN11] Shahar Dobzinski and Noam Nisan. Limitations of vcg-based mechanisms. *Combinatorica*, 31(4):379–396, 2011. 3

[DN15] Shahar Dobzinski and Noam Nisan. Multi-unit auctions: Beyond roberts. *Journal of Economic Theory*, 156:14–44, 2015. Computer Science and Economic Theory. 1

[DNO14] Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. In *the 46th annual ACM symposium on Theory of computing (STOC)*, 2014. 3

[DNS10] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010. 1, 2, 3, 12

[DNS12] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. *J. Comput. Syst. Sci.*, 78(1):15–25, 2012. 4

[Dob07] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103, 2007. 4

[Dob11] Shahar Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 139–148, New York, NY, USA, 2011. Association for Computing Machinery. 2, 4

[Dob16a] Shahar Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2016, pages 940–948, New York, NY, USA, 2016. ACM. 4

[Dob16b] Shahar Dobzinski. Computational efficiency requires simple taxation. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 209–218. IEEE, 2016. 1, 2, 3, 4, 5, 6, 8, 13, 14, 16, 17, 19

[Dob23] Shahar Dobzinski. private communication, 2023. 19

[DRV22] Shahar Dobzinski, Shiri Ron, and Jan Vondrák. On the hardness of dominant strategy mechanism design. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 690–703. ACM, 2022. 3, 4

[DSS15] Amit Daniely, Michael Schapira, and Gal Shahaf. Inapproximability of truthful mechanisms via generalizations of the vc dimension. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, page 401–408, New York, NY, USA, 2015. Association for Computing Machinery. 1

[DV11] Shaddin Dughmi and Jan Vondrak. Limitations of Randomized Mechanisms for Combinatorial Auctions. In *52nd Annual Symposium on Foundations of Computer Science (FOCS)*, 2011. 4

[DV12a] Shahar Dobzinski and Jan Vondrák. From query complexity to computational complexity. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)*, 2012. 4

[DV12b] Shahar Dobzinski and Jan Vondrak. The Computational Complexity of Truthfulness in Combinatorial Auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 2012. 4

[DV13] Shahar Dobzinski and Jan Vondrák. Communication complexity of combinatorial auctions with submodular valuations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1205–1215. SIAM, 2013. 1

[DV16] Shahar Dobzinski and Jan Vondrák. Impossibility results for truthful combinatorial auctions with submodular valuations. *J. ACM*, 63(1):5:1–5:19, 2016. 4

[EFN+19] Tomer Ezra, Michal Feldman, Eric Neyman, Inbal Talgam-Cohen, and Matt Weinberg. Settling the communication complexity of combinatorial auctions with two subadditive buyers. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 249–272. IEEE, 2019. 2, 3, 6, 7, 8, 9, 10

[Fei09] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009. 1, 2, 3, 9, 12

[FV10] Uriel Feige and Jan Vondrák. The submodular welfare problem with demand queries. *Theory of Computing*, 6(1):247–290, 2010. https://theory.stanford.edu/~jvondrak/data/submod-improve-ToC.pdf. 1, 3

[Gro73] Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973. 1

[Gue81] Roger Guesnerie. On taxation and incentives: further remarks on the limits to redistribution. *University of Bonn*, 89, 1981. 5

[Ham79] Peter J Hammond. Straightforward individual incentive compatibility in large economies. *The Review of Economic Studies*, pages 263–282, 1979. 5

[HKMT04] Ron Holzman, Noa E. Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004. 3

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997. 4

[KV12] Piotr Krysta and Berthold Vöcking. Online mechanism design (randomized rounding on the fly). In *Automata, Languages, and Programming*, pages 636–647. Springer, 2012. 4

[LLN06] Benny Lehmann, Daniel J. Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006. 1, 4

[LMN03] R. Lavi, Ahuva Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 574–583, 2003. 1

[LOS02] Daniel Lehmann, Liadan O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602, 2002. 1, 3

[LS05] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 595–604, 2005. 1

[MSV08] Vahab S. Mirrokni, Michael Schapira, and Jan Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*, pages 70–77, 2008. 1, 4

[NS06] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006. 1

[QW24] Frederick Qiu and S Matthew Weinberg. Settling the communication complexity of vcg-based mechanisms for all approximation guarantees. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1192–1203, 2024. 1, 3

[Rob79] Kevin Roberts. The characterization of implementable choice rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences. Papers presented at the first European Summer Workshop of the Economic Society*, pages 321–349. North-Holland, 1979. 1

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

[RST+21] Aviad Rubinstein, Raghuvansh R. Saxena, Clayton Thomas, S. Matthew Weinberg, and Junyao Zhao. Exponential communication separations between notions of selfishness. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 947–960. ACM, 2021. 3, 4

[Vic61] William Vickrey. Counterspeculations, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961. 1

[Von07] Jan Vondrák. Submodularity in combinatorial optimization, 2007. Available at https://dspace.cuni.cz/bitstream/handle/20.500.11956/13738/140038775.pdf. 1

[Von08] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 67–74, 2008. 4

[Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *ACM Symposium on Theory of Computing*, pages 209–213, 1979. 2

APPENDIX

A. Missing Proofs

1) Proof of Lemma II.3: Let $\mathcal{P}$ be a protocol that achieves an α-approximation to the optimal welfare with bidders whose valuations are in class $\mathcal{C}$. We first describe the protocol $\mathcal{P}'$ for $\mathcal{C} \cup \mathsf{SingleM}$, and then prove that its approximation guarantee is also α.

a) Description of Protocol.: Briefly, the protocol operates by simply running $\mathcal{P}$ independently once for each possible method of allocating the single-minded bidders their preferred bundle.

In more detail, first, each bidder i sends a bit that specifies whether v_i is single-minded or belongs in $\mathcal{C}$ (if $v_i \in \mathcal{C} \cap \mathsf{SingleM}$, it is considered single-minded). Let $\mathsf{SM} \subseteq [n]$ be the set of bidders such that v_i is single-minded, and with Other the rest of the bidders. Then, we ask all the bidders in SM to send their (minimal) desired set and their value for it. We say that a subset of single-minded bidders $K \subseteq \mathsf{SM}$ is *feasible* if for every $i, j \in K$, $S_i \cap S_j = \emptyset$. Note that at this point in the protocol, we can identify all the feasible subsets of single-minded bidders.

Next, for every feasible subset of single-minded bidders $K \subseteq \mathsf{SM}$, we consider the following allocation. First, we allocate to the bidders in K their desired sets. The rest of the single-minded bidders are allocated with the empty bundle. Then, we allocate the remaining items to the bidders in Other, using the protocol $\mathcal{P}$. Afterwards, all bidders send their value for the bundle given this allocation, so its social welfare is known. We conclude by outputting the allocation with the maximum welfare. It is easy to see that the protocol $\mathcal{P}'$ requires at most $\mathrm{poly}(\mathrm{cc}(\mathcal{P}), m, 2^n)$ bits.

b) The Approximation Ratio: Informally, the approximation ratio α still holds for $\mathcal{P}'$ because we *exactly* optimize welfare for all single-minded bidders, and on the remaining bidders, we can still only be off by a factor of at most α.

In more detail, denote with $A = (A_1, \ldots, A_n)$ the allocation that the protocol outputs and with $O = (O_1, \ldots, O_n)$ an optimal allocation. Assume without loss of generality that $(O_1, \ldots, O_n)$ and $(A_1, \ldots, A_n)$ satisfy that no bidder is allocated any item that does not strictly increase her value. Denote with K^* the subset of single-minded bidders that get a valuable bundle in the optimal allocation, i.e. $K^* = \{i : v_i(O_i) > 0, i \in \mathsf{SM}\}$. Denote with $X = (X_1, \ldots, X_n)$ the allocation at the iteration where the feasible set of single-minded bidders is K. Note that by definition for every single-minded bidder i, $X_i = O_i$. Therefore,

$$\sum_{i \in \mathsf{SM}} v_i(X_i) = \sum_{i \in \mathsf{SM}} v_i(O_i) \tag{7}$$

Also, note that $\bigcup_{i \in \mathsf{SM}} X_i = \bigcup_{i \in \mathsf{SM}} O_i$, and we denote this subset of items with J_{SM}. Observe that since O is a welfare-maximizing allocation, it is in particular an optimal allocation of the items in $[m] \setminus J_{\mathsf{SM}}$ to the bidders in Other. We remind that the allocation of the bidders in Other is determined by the protocol $\mathcal{P}$. Due to the approximation guarantee, we get that:

$$\sum_{i \in \mathsf{Other}} v_i(X_i) \geq \alpha \cdot \Big(\sum_{i \in \mathsf{Other}} v_i(O_i) \Big) \tag{8}$$

Therefore, we can deduce that:

$$\begin{aligned}
\sum_{i \in [n]} v_i(A_i) &\geq \sum_{i \in [n]} v_i(X_i) && \text{(by construction)} \\
&= \sum_{i \in \mathsf{SM}} v_i(X_i) + \sum_{i \in \mathsf{Other}} v_i(X_i) \\
&\geq \sum_{i \in \mathsf{SM}} v_i(O_i) + \alpha \cdot \sum_{i \in \mathsf{Other}} v_i(O_i) && \text{(by (7) and (8))} \\
&\geq \alpha \cdot \sum_{i \in [n]} v_i(O_i)
\end{aligned}$$

which completes the proof.

2) Proof of Theorem II.7: We begin by redefining $\mathrm{price}(\mathcal{A})$ and $\mathrm{tax}(\mathcal{A})$. The definitions are identical to the ones in [Dob16b], and we write them for the sake of completeness. $\mathrm{price}(\mathcal{A})$ is the communication complexity of the $(n-1)$-bidder problem of computing the price for a specific bundle $S \subseteq M$ in the menu presented to player i in the mechanism $\mathcal{A}$ given the valuation profile v_{-i}. For the definition of $\mathrm{tax}(\mathcal{A})$, see Definition V.6.

Now, recall that by the menu reconstruction theorem [Dob16b, Theorem 3.1], the communication complexity of computing the menu presented to a player i given the valuations v_{-i} of the other players and the truthful mechanism $\mathcal{A}$ is $\mathrm{poly}(\mathrm{tax}(\mathcal{A}), \mathrm{price}(\mathcal{A}), m, n)$. This is true for every domain. In addition, by [Dob16b, Proposition F.1], it holds that $\mathrm{price}(\mathcal{A}) \leq \mathrm{cc}(\mathcal{A})$ for every truthful mechanism for a class of valuations that includes additive valuations. Similarly, [Dob16b, Proposition 2.3] implies that $\mathrm{tax}(\mathcal{A}) \leq \mathrm{cc}(\mathcal{A})$ for every truthful mechanism $\mathcal{A}$ that is truthful for a class of valuations that contains subadditive valuations.[25] The proof is obtained by combining these assertions.

[25] In fact, the statement in [Dob16b] is specifically for the class of subadditive valuations. However, the proof holds as-written for every class of valuations that contains subadditive valuations.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

3) Missing Proofs of Section IV:

Proof of Lemma IV.7. We will show that for sufficiently large m, there exists an ℓ-independent k-width-family, where $\ell = \frac{1}{4}\log m$ and $k = e^{\frac{2\sqrt{m}}{\log m}}$.

To obtain a ℓ-independent k-width-family $\mathcal{F} = (\mathcal{G}, \{\mathcal{H}_i^{(t)}\}_{t\in\{0,1\},i\in[k]})$, consider the following two-layered randomized construction:

1) *First Layer*: For every $i \in [k]$, let G_i be the random set that contains each item $e \in [m]$ independently with probability $\frac{1}{2}$.
2) *Second Layer*: For every $i, j \in [k]$, let $H_{i,j}^{(0)}$ be the random set that contains each $e \in \overline{G_i}$ independently with probability $\frac{1}{2}$. Similarly, let $H_{i,j}^{(1)}$ be the random set that contains each item $e \in G_i$ independently with probability $\frac{1}{2}$. Note that for every i and for every $t \in \{0,1\}$, the distribution of the items in $H_{i,j}^{(t)}$ is identical for all values of j.

Thus, by construction $\mathcal{F}$ is a k-width-family. Thus, it remains to prove that with a non-zero probability, $\mathcal{F}$ is ℓ-independent.

Recall that $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}] = \{S_{i,j}\}_{i,j\in[k]}$ is a collection of k^2 subsets. We can fix an index set $I \subset [k]\times[k]$ of size ℓ, which specifies a sub-collection $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I] = \{S_{i,j}\}_{(i,j)\in I}$ of size ℓ accordingly. Our plan is to show that for every choice of $\boldsymbol{b}, \boldsymbol{C}$ and I, it is very unlikely that the sub-collection $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I]$ covers $[m]$, and then the existence of a valid $\mathcal{F}$ follows from union bound.

Observe that for every item e and $S_{i,j}$, we have that $e \in S_{i,j}$ with probability $3/4$. In particular, it holds for all the subsets $S_{i,j}$ in $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I]$. Also, note that the events $\{e \in \cup\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I]\}_{e\in[m]}$ are mutually independent. Therefore, we have that:

$$\begin{aligned}
&\Pr\left[\bigcup \mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I] = [m]\right] \\
&= \prod_{e\in[m]} \left(1 - \Pr\left[e \notin \bigcup \mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I]\right]\right) \\
&= (1 - \frac{1}{4^\ell})^m \leq (1 - 2^{-2\ell})^m.
\end{aligned}$$

Finally by union bound:

$$\begin{aligned}
&\Pr\left[\exists \boldsymbol{b}, \boldsymbol{C}, I \text{ such that } \bigcup \mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I] = [m]\right] \\
&\leq \sum_{I\subset[k]\times[k],|I|=\ell} \Pr\left[\bigcup \mathcal{F}[\boldsymbol{b}, \boldsymbol{C}; I] = [m]\right] \\
&\leq \binom{k^2}{\ell} 2^{2\ell}(1 - 2^{-2\ell})^m \\
&< k^{2\ell}\exp(-2^{-2\ell}m) = \exp(2\ell\ln k - 2^{-2\ell}m) = 1.
\end{aligned}$$

Therefore, with non-zero probability, $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}]$ is ℓ-sparse for all $\boldsymbol{b}, \boldsymbol{C}$. □

4) Missing Proofs of Section V:

Proof of Lemma V.5. The proof is based on the probabilistic method. We construct the collection $\mathcal{G}$ by sampling uniformly at random $\exp(b^2m/6)$ subsets of $[m]$, where the size of each subset is bm. For any two subsets G_i, G_j of size bm sampled uniformly at random, note that the expected size of their intersection $\mathbb{E}[|G_i \cap G_j|] = b^2m$. By Chernoff bound, we have $\Pr[|G_i\cap G_j| > 2b^2m] \leq \exp(-b^2m/3)$. By applying the union bound, we conclude that:

$$\begin{aligned}
&\Pr[\exists i \neq j, |G_i \cap G_j| > 2b^2m] \\
&\leq \sum_{i\neq j} \Pr[|G_i \cap G_j| > 2b^2m] \\
&< |\mathcal{G}|^2 \cdot \exp(-b^2m/3) \leq 1.
\end{aligned}$$

□

5) Proof of Theorem V.1 by Hardness of Simultaneous Protocols: Recall the following two-bidder corollary from the Taxation Complexity framework:

Theorem A.1 ([Dob16b]). *Let $\mathcal{A} = (f, p_1, p_2) : \mathcal{V}_1 \times \mathcal{V}_2 \to \Sigma \times \mathbb{R}^n$ be any deterministic truthful mechanism that guarantees an α-approximation to the optimal welfare, and both domains $\mathcal{V}_1, \mathcal{V}_2$ contain XOS_m. Then, there exists a simultaneous protocol $f' : \mathcal{V}_1 \times \mathcal{V}_2 \to \Sigma$ that also guarantees an α-approximation with simultaneous communication complexity $\mathrm{cc}(f') \leq \mathrm{poly}(\mathrm{cc}(\mathcal{A}), m)$.*

Based on this result, it suffices to give a lower bound of the communication complexity for simultaneous algorithms to prove Theorem V.1.

Proposition A.2. *Let $f : \mathsf{SingleM} \times \mathsf{XOS} \to \Sigma$ be any deterministic simultaneous protocol. If f obtains a $(\frac{\sqrt{5}-1}{2}+3m^{-\frac{1}{3}})$-approximation to the optimal welfare, then its simultaneous communication complexity $\mathrm{cc}(f) \geq \exp(m^{\frac{1}{3}}/6)$.*

Proof. Let Alice be the single-minded bidder and Bob be the XOS bidder. Since f is a simultaneous algorithm, if Bob sends an identical message upon different valuations $v'_B \neq v_B$, then the resulting allocation will also be identical, i.e., $f(v_A, v_B) = f(v_A, v'_B)$ for all $v_A \in \mathsf{SingleM}$. We will show that it further leads to a contradiction with the approximation ratio.

For any distinct $v_B, v'_B \in \mathsf{XOS}^*$, denote with $\mathcal{H}$ and with $\mathcal{H}'$ respectively the combinations that are associated with v_B and with v'_B respectively. Since $v_B \neq v'_B$, there exists at least one subset H such that (without loss of generality) $H \in \mathcal{H}$, whereas $H \notin \mathcal{H}'$. Consider the case when Alice (with valuation v_A) gives a value of $\alpha < bm$ to the bundle $\overline{H} = [m] \setminus H$ (and zero otherwise). If by contradiction, $f(v_A, v_B) = f(v_A, v'_B)$, then Alice must be given a same bundle $A \subseteq [m]$ for both instances (v_A, v_B) and (v_A, v'_B). There are two possibilities:

- If $\overline{H} \subseteq A$ (i.e., Alice gets $\overline{H}$): for the instance (v_A, v'_B), the actual welfare f gets is at most $v_A(\overline{H}) + v'_B(H) \leq \alpha + 2b^2m$, while the optimal welfare being at least $v'_B([m]) = bm$.
- If $\overline{H} \not\subseteq A$ (i.e., Alice gets nothing valuable): for the instance (v_A, v_B), the actual welfare f gets is at most $v_B([m]) = bm$, while the optimal welfare being $v_A(\overline{H}) + v_B(H) = \alpha + bm$.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

Overall, we know that the approximation ratio of f can be at most

$$\max\left\{\frac{\alpha+2b^2m}{bm}, \frac{bm}{\alpha+bm}\right\}.$$

Recall $b = m^{-\frac{1}{3}}$ and let $\alpha = \frac{\sqrt{5}-1}{2}bm = \frac{\sqrt{5}-1}{2}m^{\frac{2}{3}}$. Then, the ratio above is always strictly smaller than $\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}}$. Therefore, to get a $(\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}})$-approximation, Bob must send different messages for all $v_B \in \mathsf{XOS}^*$. As a result, $\mathrm{cc}(f) \geq \log|\mathsf{XOS}^*| = |\mathcal{C}| = \exp(m^{\frac{1}{3}}/6)$. □

Proof of Theorem V.1. By Theorem A.1, there is some simultaneous protocol $f' : \mathcal{V} \times \mathcal{V} \to \Sigma$ that $(\frac{\sqrt{5}-1}{2} + 3m^{-\frac{1}{3}})$-approximates the optimal welfare with simultaneous communication complexity $\mathrm{cc}(f') \leq \mathrm{poly}(\mathrm{cc}(\mathcal{A}), m)$. However, by Proposition A.2, we know $\mathrm{cc}(f') \geq \exp(m^{\frac{1}{3}}/6)$. Therefore, we conclude that $\mathrm{cc}(\mathcal{A}) \geq \exp(\Omega(m^{\frac{1}{3}}))$. □

Remark A.3. We now show that simultaneous protocols for two bidders can achieve $\frac{\sqrt{5}-1}{2}$ of the welfare for the class $\mathsf{SingleM} \times \mathsf{XOS}$, meaning that Proposition A.2 is asymptotically tight. Denote the bundle that Alice (the single-minded bidder) wants by S. Consider some $f^* : \mathsf{SingleM} \times \mathsf{XOS} \to \Sigma$ that gives the bundle S to Alice and the bundle $[m] \setminus S$ to Bob when $v_1(S) \geq \frac{\sqrt{5}-1}{2}v_2([m])$, and otherwise allocates all items to Bob. f^* can be implemented by simultaneous communication and always gives a $\frac{\sqrt{5}-1}{2}$-approximation to welfare.

B. The Structure of Payments of Truthful Mechanisms

In this section, we show a formula for an upper bound and lower bound on the payments of approximately optimal and truthful mechanisms whose class of valuations includes single-minded bidders. An asymptotic interpretation of Proposition A.5, which might be of independent interest is that as the approximation guarantee α of a truthful mechanism goes to 1, then the payments of the mechanism are approach VCG payments. Moreover, Proposition A.5 in its current form applies to two-bidder mechanisms, but it can readily be extended to settings with an arbitrary number of bidders, as well as to the related setting of multi-unit auctions.

We start by defining a class of SingleM valuations which allow us to bound the prices that any two-player good-approximation mechanism must charge:

Definition A.4. Fix a valuation v, a nonempty subset of items $S \subseteq [m]$ and two constants $\alpha \in [0,1]$ and $\epsilon > 0$. We define a valuation $\mathrm{upper}_{v,\alpha,\epsilon,S}$ as follows: $\forall X \subseteq [m]$, $\mathrm{upper}_{v,\alpha,\epsilon,S}(X)$ equals to

$$\begin{cases} \frac{1}{\alpha}\cdot v([m]) - v([m]\setminus S) + \epsilon & X \supseteq S, \\ 0 & \text{otherwise.} \end{cases}$$

We define $\mathrm{lower}_{v,\alpha,\epsilon,S}$ similarly: $\forall X \subseteq [m]$, $\mathrm{lower}_{v,\alpha,\epsilon,S}(X)$ equals to

$$\begin{cases} \max\{\alpha\cdot v([m]) - v([m]\setminus S) - \epsilon, 0\} & X \supseteq S, \\ 0 & \text{otherwise.} \end{cases}$$

Now, we give the bound on prices enabled by Definition A.4:

Proposition A.5. *Let $\mathcal{A} = (f, p_1, p_2) : \mathcal{V}_1 \times \mathcal{V}_2 \to \Sigma \times \mathbb{R}^2$ be any deterministic truthful mechanism that gives an α-approximation to the optimal welfare. For some $\epsilon > 0$, if both $\mathrm{upper}_{v_2,\alpha,\epsilon,S}$ and $\mathrm{lower}_{v_2,\alpha,\epsilon,S}$ belong to $\mathcal{V}_1$, then for every valuation $v_2 \in \mathcal{V}_2$ and every bundle $S \subseteq [m]$,*

$$\begin{aligned} &\frac{1}{\alpha}\cdot v_2([m]) - v_2([m]\setminus S) + \epsilon \\ &\geq \mathsf{Menu}_1^{\mathcal{A}}(v_2)(S) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(\emptyset) \\ &\geq \alpha\cdot v_2([m]) - v_2([m]\setminus S) - \epsilon \end{aligned}$$

Proof. Denote with v_1^{up} and v_1^{low} the valuations $\mathrm{upper}_{v_2,\alpha,\epsilon,S}$ and $\mathrm{lower}_{v_2,\alpha,\epsilon,S}$ respectively. Consider the possible allocations that achieve an α approximation given the valuation profile (v_1^{up}, v_2).

Consider a bundle T that contains the desired bundle S. Using the fact that v_1^{up} is single-minded, we have that

$$v_1^{\mathrm{up}}(T) + v_2([m]\setminus T) \geq v_1^{\mathrm{up}}(S) + v_2([m]\setminus S) = \frac{1}{\alpha}v_2([m]) + \epsilon.$$

On the other hand, if a bundle T does not contain S, then:

$$v_1^{\mathrm{up}}(T) + v_2([m]\setminus T) \leq v_2([m]) < \alpha\cdot\left(\frac{1}{\alpha}v_2([m]) + \epsilon\right).$$

Thus, to get an α-approximation to the optimal welfare, we must allocate a bundle containing S to bidder 1 given the valuation profile (v_1^{up}, v_2).

Analogously, consider the possible allocations that achieve an α approximation given $(v_1^{\mathrm{low}}, v_2)$. We have $v_1^{\mathrm{low}}(S) + v_2([m]\setminus S) = \alpha\cdot v_2([m]) - \epsilon < \alpha\cdot v_2([m])$, which implies that in this case any allocation awarding S to bidder 1 cannot possibly be an α-approximation to the optimal welfare.

We now achieve the desired bounds on payments by using the fact that the mechanism is truthful, so given v_1^{up}, bidder 1 will not want to misreport the valuation v_1^{low} and vice-versa. Let X^{up} be the bundle that bidder 1 wins given (v_1^{up}, v_2). Using truthfulness and the monotonicity of the menu (Proposition II.6), we have that:

$$\begin{aligned} v_1^{\mathrm{up}}(S) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(S) &\geq v_1^{\mathrm{up}}(X^{\mathrm{up}}) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(X^{\mathrm{up}}) \\ &\geq v_1^{\mathrm{up}}(\emptyset) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(\emptyset). \end{aligned}$$

Rearranging this inequality and applying the definition of v_1^{up}:

$$\begin{aligned} \mathsf{Menu}_1^{\mathcal{A}}(v_2)(S) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(\emptyset) &\leq v_1^{\mathrm{up}}(S) - v_1^{\mathrm{up}}(\emptyset) \\ &= \frac{1}{\alpha}\cdot v_2([m]) - v_2([m]\setminus S) + \epsilon. \end{aligned}$$

Thus, we obtain the first part of the lemma.

For the second part, let X^{low} be the bundle that bidder 1 wins given $(v_1^{\mathrm{low}}, v_2)$. Since $v_1^{\mathrm{low}}(X^{\mathrm{low}}) = 0 = v_1^{\mathrm{low}}(\emptyset)$, applying truthfulness and menu monotonicity gives:

$$\begin{aligned} v_1^{\mathrm{low}}(\emptyset) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(\emptyset) &\geq v_1^{\mathrm{low}}(X^{\mathrm{low}}) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(X^{\mathrm{low}}) \\ &\geq v_1^{\mathrm{low}}(S) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(S), \end{aligned}$$

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

and again rearranging, we have

$$\begin{aligned}\mathsf{Menu}_1^{\mathcal{A}}(v_2)(S) - \mathsf{Menu}_1^{\mathcal{A}}(v_2)(\emptyset) &\geq v_1^{\text{low}}(S) - v_1^{\text{low}}(\emptyset) \\ &= \alpha \cdot v_2([m]) - v_2([m] \setminus S) - \epsilon,\end{aligned}$$

which finishes the proof. □

C. An Example of a 2-Width-Family

In this section, we give an example of a 2-width-family, to demystify our construction in Section IV. Consider the case where there are 6 items, which we denote with $\{1,2,3,4,5,6\}$. Consider the set family defined as follows:

$$\begin{aligned}\mathcal{G} &= \left\{G_1 = \{1,2\}, G_2 = \{3,5,6\}\right\} \\ \mathcal{H}_1^{(0)} &= \left\{H_{1,1}^{(0)} = \{3,4\}, H_{1,2}^{(0)} = \{5,6\}\right\} \\ \mathcal{H}_1^{(1)} &= \left\{H_{1,1}^{(1)} = \{1,2\}, H_{1,2}^{(1)} = \{1\}\right\} \\ \mathcal{H}_2^{(0)} &= \left\{H_{2,1}^{(0)} = \{2,4\}, H_{2,2}^{(0)} = \{1\}\right\} \\ \mathcal{H}_2^{(1)} &= \left\{H_{2,1}^{(1)} = \{3,5\}, H_{2,2}^{(1)} = \{5\}\right\}\end{aligned}$$

Now, we define:

$$\boldsymbol{b} = (0,1), \quad \boldsymbol{C} = \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}$$

Thus, the collection $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}]$ is composed of the subsets $S_{1,1}, S_{1,2}, S_{2,1}, S_{2,2}$ that are defined as follows:

$$\begin{aligned}b[1] = 0, C[1][1] = 0 \implies S_{1,1} &= G_1 \cup H_{1,1}^{(0)} \\ &= \{1,2,3,4\} \\ b[1] = 0, C[1][2] = 1 \implies S_{1,2} &= G_1 \cup (\overline{G}_1 \setminus H_{1,2}^{(0)}) \\ &= \{1,2,3,4\} \\ b[2] = 1, C[2][1] = 0 \implies S_{2,1} &= \overline{G}_2 \cup H_{2,1}^{(1)} \\ &= \{1,2,3,4,5\} \\ b[2] = 1, C[2][2] = 1 \implies S_{2,2} &= \overline{G}_2 \cup (G_2 \setminus H_{2,2}^{(1)}) \\ &= \{1,2,3,4,6\}\end{aligned}$$

Observe that by definition $\mathcal{F}[\boldsymbol{b}, \boldsymbol{C}]$ is 2-sparse, because there is no one subset among $S_{1,1}, S_{1,2}, S_{2,1}, S_{2,2}$ that covers $\{1,2,3,4,5,6\}$. However, it is not 3-sparse, because $S_{2,1} \cup S_{2,2} = \{1,2,3,4,5,6\}$, which implies that the family $\mathcal{F}$ is not 3-independent.

D. On The Precision of Mechanisms and Valuations

In this section, we dive deeper into an assumption of the taxation framework; in particular, we remark on the precision of the bidders' valuations vs. the mechanism's prices. Observe that when discussing the communication complexity of combinatorial auctions, it is necessary to assume that the valuation functions of each bidder has values that are bounded in some range. The reason for it is that otherwise, even the communication complexity of a single item auction is infinite, regardless of incentives.[26] Therefore, when formally considering a class of valuation functions, it is necessary to specify not only their properties, i.e., whether they are subadditive or single-minded, but also their precision.

Formally handling the precision of valuations has been discussed before, e.g. in [Dob16b, Appendix A.4.1]. As discussed in our footnote 11, the basic approach is to assume all valuations can be represented by k-bit numbers for some implicit parameter k, and to formally define a protocol for some valuation class as a *family* of protocols $(P_k)_k$ for different precision k (with the added assumption that when $k' > k$ and $P_{k'}$ is run on k-bit numbers, the output is the same as the output of P_k). To our knowledge, all prior upper and lower bounds in the algorithmic mechanism design literature hold in this formal model while implicitly hiding factors of $\text{poly}(k)$.

We will now address an additional assumption of the taxation framework, and show that it is in fact not necessary. Namely, [Dob16b] assumes that mechanisms that are "precision-aligned". A mechanism $\mathcal{A} : \mathcal{V} \to \Sigma \times \mathbb{R}^n$ is *precision-aligned* if the precision of the mechanism is equal to the precision of $\mathcal{V}$. More concretely, if every valuation in $\mathcal{V}$ has values in a certain range R, then the mechanism outputs prices that are in range R.

However, one can easily show that, in fact, assuming precision-aligned mechanisms is without loss of generality. Thus, the results of [Dob16b] (and hence, other results which build on these, such as ours and those of [AKSW20]) in fact hold for all mechanisms, not just for precision-aligned ones. Formally:

Proposition A.6 (Communicated by [Dob23])**.** *Let $\mathcal{A} : \mathcal{V} \to \Sigma \times \mathbb{R}^n$ be a truthful mechanism. Then, there exists a truthful precision-aligned mechanism $\mathcal{A}'$ such that:*

1) *For every valuation profile $v \in \mathcal{V}$, the mechanisms $\mathcal{A}$ and $\mathcal{A}'$ output the same allocation.*
2) *The communication complexity of $\mathcal{A}'$ is the same as the communication complexity of $\mathcal{A}$.*

Proof. We assume for simplicity that the range of the valuations in $\mathcal{V}$ is integer values in $\{0, \ldots, H\}$, though an analogous proof works for every fixed decimal precision.

The mechanism of $\mathcal{A}'$ that we propose is equivalent to the mechanism $\mathcal{A}$, except that the prices in each leaf are rounded down to the nearest integer. Thus, we get that the mechanisms $\mathcal{A}'$ and $\mathcal{A}$ have the same precision and the same communication complexity. It remains to show that $\mathcal{A}'$ is truthful. We remind that by the taxation principle, it suffices to show that for every player i, and every valuation profile (v_i, v_{-i}), the bundle that $\mathcal{A}'$ allocates to bidder i remains the most profitable given the new rounded prices.

Fix such player i and a valuation profile (v_i, v_{-i}). Let S be the most profitable bundle, and let T be some other bundle. Denote their prices given the (old) mechanism $\mathcal{A}$ with p_S and p_T, respectively. Since $\mathcal{A}$ is truthful, we have that $v_i(S) - p_S \geq v_i(T) - p_T$, and our goal is to show that: $v_i(S) - \lfloor p_S \rfloor \geq v_i(T) - \lfloor p_T \rfloor$.

[26] Consider a single-item auction with two bidders with values in the infinite range $R = \{0, \ell, \ell^2, \ell^3, \ldots\}$ for the item, where ℓ is arbitrarily large. Applying the fooling set argument for the pairs $(x,x)_{x \in R}$ gives that the communication complexity of every mechanism that gives an approximation better than ℓ for the optimal welfare requires at least one transcript for every element of R, so it requires infinite communication.

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.

To prove this, we represent p_S as an integer n_S and a fraction f_S, where $p_S = n_S + f_S$, where $n_S = \lfloor p_S \rfloor$ and $f_S \in [0, 1)$. We define n_T and f_T analogously. Therefore, we have that:

$$v_i(S) - n_S - f_S \geq v_i(T) - n_T - f_T \tag{9}$$

Thus, we now have that the integer part of the left side is $v_i(S) - n_S$ and the integer part of the right side is $v_i(T) - n_T$. Since both $f_S, f_T \in [0, 1)$, then their difference satisfies that $|f_T - f_S| < 1$, so for Equation 9 to hold, it has to be the case that $v_i(S) - n_S \geq v_i(T) - n_T$. We remind that $n_S = \lfloor p_S \rfloor$ and $n_T = \lfloor p_T \rfloor$, so we have that $v_i(S) - \lfloor p_S \rfloor \geq v_i(T) - \lfloor p_T \rfloor$. Moreover, n_S and n_T are in fact the menu prices in the new mechanism $\mathcal{A}'$, so this mechanism is in fact truthful, as needed. □

Authorized licensed use limited to: Princeton University. Downloaded on March 11,2025 at 01:26:01 UTC from IEEE Xplore. Restrictions apply.