The Randomized Communication Complexity of Randomized Auctions

Aviad Rubinstein*
Stanford University
U.S.A.

Junyao Zhao†
Stanford University
U.S.A.

ABSTRACT

We study the communication complexity of incentive compatible auction-protocols between a monopolist seller and a single buyer with a combinatorial valuation function over n items. Motivated by the fact that revenue-optimal auctions are randomized (as well as by an open problem of Babaioff, Gonczarowski, and Nisan), we focus on the *randomized* communication complexity of this problem (in contrast to most prior work on deterministic communication).

We design simple, incentive compatible, and revenue-optimal auction-protocols whose expected communication complexity is much (in fact infinitely) more efficient than their deterministic counterparts.

We also give nearly matching lower bounds on the expected communication complexity of approximately-revenue-optimal auctions. These results follow from a simple characterization of incentive compatible auction-protocols that allows us to prove lower bounds against randomized auction-protocols. In particular, our lower bounds give the first approximation-resistant, exponential separation between communication complexity of *incentivizing* vs *implementing* a Bayesian incentive compatible social choice rule, settling an open question of Fadel and Segal.

CCS CONCEPTS

• **Theory of computation → Algorithmic game theory**; **Communication complexity**.

KEYWORDS

Algorithmic game theory, Communication complexity

ACM Reference Format:
Aviad Rubinstein and Junyao Zhao. 2021. The Randomized Communication Complexity of Randomized Auctions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC '21), June 21–25, 2021, Virtual, Italy*. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3406325.3451111

*Supported by NSF CCF- 1954927, and a David and Lucile Packard Fellowship.
†Supported by NSF CCF-1954927.

STOC '21, June 21–25, 2021, Virtual, Italy

ACM ISBN 978-1-4503-8053-9/21/06…$15.00
https://doi.org/10.1145/3406325.3451111

1 INTRODUCTION

The central goal of Algorithmic Mechanism Design is to design mechanisms that guarantee good outcomes while taking into account both the selfish agents' incentives and the ever-increasing complexity of modern applications. A fundamental question to this field is whether simultaneously satisfying both the incentive and simplicity constraints is harder than satisfying each of them separately.

In this paper we focus on one of the simplest and most studied settings in the field: a monopolist, Bayesian, revenue-maximizing seller auctioning n items to a single risk-neutral buyer. An active line of work over the past two decades argues that even in this strategically-simple setting, and even for buyers with additive or unit-demand valuations[1], optimal mechanisms are inherently complex, e.g. they involve randomized lotteries [16, 50, 62, 66, 80] and are often computationally intractable [25, 26, 31].

One particularly influential measure of complexity of mechanisms is the *menu-size complexity* of [49]: by the taxation principle, a general incentive compatible mechanism can be canonically represented as a *menu*, where each *line* or option in the menu corresponds to a (possibly randomized) allocation and a payment. The menu-size complexity of a mechanism is then the number of lines in the corresponding menu. Perhaps the single most convincing evidence for the complexity of optimal mechanisms is an example due to [32], where the optimal mechanism for an additive buyer with two i.i.d. item valuations from a seemingly benign distribution ($\text{Beta}(1,2)$) requires an infinite and even *uncountable* menu-size complexity. We henceforth refer to this powerful example as the DDT example.

[32] and related complexity results for revenue-maximizing auctions have inspired fruitful lines of work that circumvent these barriers, e.g. by designing sub-optimal but simple mechanisms that approximate the optimal revenue (see discussion in Related work).

It is not a-priori clear, however, that the menu-size complexity by itself is an obstacle to using optimal mechanisms. For instance, the seller in the DDT example could in principle succinctly describe her[2] mechanism as "the-optimal-auction-for-$\text{Beta}(1,2)\times\text{Beta}(1,2)$" and even point the buyer to an explicit description in [32]. However, [5] recently observed that, once the mechanism is announced, the deterministic communication complexity to implement it is equal

[1] To circumvent some worst-case pathological examples, it is common in Algorithmic Mechanism Design to restrict the buyer's value distribution to *independent* (vs *correlated*) items, and/or restrict the combinatorial nature of buyer's value for bundles to one of the following classes:

additive, unit-demand ⊂ *gross-substitutes* ⊂ *submodular* ⊂ *XOS* ⊂ *subadditive.*

(See Section 2.1 for definitions.)
[2] Throughout the paper, we use feminine pronouns for the seller and masculine for the buyer.

(up to rounding) to the logarithm of the menu-size complexity. In the DDT example, for the buyer to deterministically specify his favorite line in the uncountable menu, he would need to send an infinite stream of bits. [5] left open the question of randomized communication complexity of optimal mechanisms. Indeed randomized communication is a natural complexity measure in this case since we already consider randomized allocations[3].

In this paper, inspired by [5]'s open question, we formulate a notion[4] of an *incentive compatible (IC) auction protocol*, which is a two-party (possibly randomized) interactive communication protocol between a seller and a buyer with an allocation and payment associated with every transcript of the protocol. Before presenting our results in this model, below we briefly discuss our modeling assumptions; a full definition appears in Section 2.

1.1 Brief Discussion of Modeling Assumptions

Per the discussion above, we assume that the protocol and auction format are public information. The buyer privately knows his true type (or valuations of items/bundles).

We mostly focus on the total expected communication complexity of the protocol. For our protocols, we bound the interim expectation, i.e. for every buyer's type, the communication complexity of the protocol is bounded, in expectation[5] over the protocol's randomness[6]. Our lower bounds hold even for ex ante expectation, i.e. even if we allowed that some buyers may know in advance that they are expected to participate in a prohibitively long protocol.

The seller in our model has *no private information* and is *not strategic*. At the end of the communication protocol she must know the allocation and payment.

We model the buyer's strategic aspect as a complete information single-player extensive-form game with buyer's nodes and nodes of Chance; each leaf is associated with an allocation and a payment. In practice, nodes of Chance could be implemented by a trusted seller (e.g. when the seller is an auditable firm), a trusted intermediary, a cryptographic protocol for coin tossing[7], or a publicly observable, renewable[8] external source of randomness.

As is common in the aforementioned literature on randomized mechanisms, we assume that the buyer is risk-neutral. In particular, we require that the protocol is interim individually rational. In direct revelation mechanisms, it is possible to transform interim to ex-post individual rationality by correlating the payment with the randomized allocation. Similarly, at the cost of a bounded increase in the communication complexity, it is possible to transform our protocols to become ex-post approximately individually rational (see the full version for details).

While we make little restrictions on buyer valuations, we do generally assume that the buyer's valuation is capped at some arbitrarily large value U. The complexity of our protocols [9] does not depend on U, e.g. U can be all the money in the universe (typically much smaller).

1.2 Our IC Auction Protocols

We design IC auction protocols that are simple, surprisingly efficient, and are *exactly* revenue-optimal. For instance, in Theorem 3.2 we give a revenue-optimal IC auction protocol for the DDT example where the *buyer sends less than two bits in expectation*. (In contrast, for a deterministic auction selling two items separately, merely specifying the allocation requires the buyer to send two bits!)

Main positive result. Our main positive result is a generic transformation of an arbitrary (revenue-optimal or otherwise) IC and IR mechanism for additive, unit-demand, or general combinatorial valuations to an IC auction protocol that uses $O(n\log(n))$, $O(n\log(n))$, $O(2^n n)$ bits in expectation respectively. We note that our protocols work for correlated prior distributions, and even for non-monotone and negative valuations[10].

Theorem (See Theorems 3.1 and 4.1)**.**
For any prior $\mathcal{D}$ of buyer's (additive/unit-demand/combinatorial) valuations over n items bounded by maximum valuation U, and any IC mechanism $\mathcal{M}$, there is an IC auction protocol with the same expected payment and allocation, using ($O(n\log n)$/$O(n\log n)$/ $O(2^n n)$) bits of communication in expectation.

Trading off revenue for even better communication efficiency. We obtain an exponentially more efficient protocol for the special case of unit-demand with *independent items*. Specifically, at the cost of an ε-fraction loss in revenue, we obtain an IC auction protocol that uses only polylog(n) communication.

Theorem (See Theorem 5.1)**.** Let $\mathcal{D}$ be a distribution of independent unit-demand valuations over n items bounded by maximum valuation U. Then, for any constant $\varepsilon > 0$, there is a $(1-\varepsilon)$-approximately revenue-optimal IC auction protocol using polylog(n) bits of communication in expectation.

Exhibiting the richness of our IC auction protocol model, this protocol is substantially different from the generic transformation in our main result, and builds on the recent *symmetric menu-size complexity* of [57].

Remark 1.1. For simplicity of presentation we focus on the expected communication complexity. Here we briefly remark that

[3] Different applications have different simplicity desiderata. (E.g. highly regulated FCC auctions *vs* very fast ad auctions with automated bidders *vs* smart contracts that require costly documentation of transaction details on a blockchain.) Ultimately, there is no universal "right" measures of complexity, and studying a variety gives us a more complete understanding.

[4] Technically, our definition of IC auction protocol is a special case of *Bayesian incentive compatibility (BIC)-incentivizable binary dynamic mechanism (BDM)* [40]. We discuss this connection further in Related work.

[5] In expectation vs high probability: We remark that by Markov's inequality in expectation *upper* bounds on the communication complexity imply similar upper bounds w.h.p.; e.g. if the expected complexity is at most C, then it is at most C/α w.p. $\geq 1-\alpha$.

[6] In fact, all our protocols happen to satisfy a slightly stronger desideratum where all the communication complexity bounds that we prove also (approximately) hold for the communication complexity of the future of any prefix of the protocol. I.e. for any setting where we bound the expected communication complexity by C, it is also true that, conditioning on any history of the protocol (possibly much longer than C), the remaining expected communication complexity is $O(C)$. This means that the buyer and seller *always* expect -for every run of the protocol, and at any point during the execution- that the protocol will end soon.

[7] We're mostly interested in mechanisms that are *exactly* revenue-optimal, while the security of cryptographic protocols always has a negligible but non-zero chance of being broken even by naive brute-force algorithms. In theory, this small chance of cheating on the coin tosses would violate the buyer's exact incentive constraints.

[8] By "renewable" we mean that at each step of the protocol the parties have access to fresh random bits not predictable in previous iterations; for example, they could look at the weather each day.

[9] Except the ex-post approximately individually rational protocols in the full version.

[10] We assume for simplicity that all payments are non-negative.

our protocols also have desirable properties in terms of round- and random-coin-complexities. For round complexity, our protocols use $O(\log(n))$ rounds in expectation ($O(n)$ for general combinatorial valuations). For the protocols in Theorems 3.1 and 4.1, it will be easy to see how (using trivial batching) one can further compress the number of rounds: at the cost of a constant factor increase in the communication complexity, these protocols can be compressed to $1+\varepsilon$ rounds in expectation. In terms of random coins, our protocols can be implemented with $O(\log(n))$ coins in expectation ($O(n)$ for general combinatorial valuations).

1.3 Communication Complexity Lower Bounds

We show that beyond the (important) special case covered by Theorem 5.1, the communication complexity of our protocols is almost the best possible, in the following strong sense:

Theorem (See Theorems 6.3, 6.7, and 7.1)**.** For revenue maximization with n items, any incentive compatible auction protocol that achieves any constant factor approximation of the optimal revenue must use at least:

- $\Omega(n)$ communication for unit-demand valuations;
- $2^{\Omega(n^{1/3})}$ communication for gross substitutes valuations;
- $2^{\Omega(n)}$ for XOS valuations.

Furthermore, any incentive compatible auction protocol obtaining more than 80% of the optimal revenue must use at least:

- $2^{\Omega(n)}$ communication for XOS valuations over *independent items*.

To place the result for independent items in the greater context of Algorithmic Mechanism Design, contrast it with simple-but-approximately-optimal mechanism independent subadditive valuations: [74] showed that a constant fraction of revenue can be guaranteed by simple mechanisms; this constant has been improved in followup works [17, 18, 23], but no non-trivial upper bound on the best approximation factor were known[11]. Assuming that efficient randomized communication is a *necessary* desideratum for "simple mechanism", our result for independent items implies that the optimal approximation factor is bounded away from 1 – even for the special case of XOS valuations.

Note also that our upper and lower bounds for correlated valuations are nearly tight in the following ways:

- For unit-demand and combinatorial valuations, our upper and lower bounds nearly match (up to logarithmic factors), even though the lower bounds hold for *arbitrary (constant) approximation factor* vs *exactly* revenue-optimal in upper bounds. Furthermore the combinatorial upper bound holds for *arbitrary* combinatorial valuations, which are much more general than XOS valuations used in the lower bound.
- The correlation in our unit-demand lower bound is necessary by Theorem 5.1.

We remark that for one interesting case an exponential gap remains:

Open Question 1.2. What is the randomized communication complexity of exactly revenue optimal IC auction protocols for unit demand valuations over independent items?

Our lower bound for unit-demand requires correlated items (and this is an inherent limitation of our technique). On the other hand, our protocol for unit-demand with independent items (Theorem 5.1) does not guarantee exact revenue optimality.

1.4 Separating the Complexity of Implementing and Incentivizing

Our results also have implications for a question of Fadel and Segal [40]. They study, for any fixed social choice rule, the *communication cost of selfishness*, i.e. the difference in communication complexity between (i) implementing it, and (ii) implementing it in a Bayesian incentive compatible protocol. They give examples where the communication cost of selfishness is exponential, but those examples are very brittle in the sense that they rely on agents' utilities to have unbounded (or at least exponential) precision. They ask whether the communication cost of selfishness on any (possibly contrived) social choice rule can be reduced substantially if agents' utilities have a bounded precision [40, Open Question 3]. Our source of hardness is inherently different from the instances in [40]: we harness the combinatorial structure of the valuations rather than exploiting the long representation of high-precision numbers.

In more detail, in our constructions the buyer's utility only requires constant precision[12] for any outcome (and the seller is not strategic, i.e. she has constant utility zero). Furthermore, for our hard instances of unit-demand valuations, we show (Remark 6.6) that the exactly revenue-optimal IC mechanism can be implemented by a randomized (non-IC) protocol using $O(\log(n))$ communication even in the worst case, hence resolving [40]'s open question on the negative[13]. We remark that by [40, Corollary 3], this exponential separation is tight.

Corollary 1.3 (See Remark 6.6). *There exists a randomized protocol for a revenue maximization instance, in which the buyer's valuation has constant precision, such that there is an exponential separation between the communication complexity of its approximately Bayesian IC implementation and that of its non-IC implementation.*

Remark 1.4 (Separations for deterministic vs randomized protocols)**.**

Formally, [40] phrase their open question for deterministic protocols. To view Corollary 1.3 in this context, note that in our model the seller is not strategic; hence one can consider an equivalent deterministic social choice rule in a slightly different setting where the random seed (only $O(\log(n)$ bits are necessary) to the revenue-optimal auction is replaced by a seller's type. The requirements from the protocol in this setting is only stricter, so the communication lower bound on IC auction protocols trivially extends. On the other hand, for the non-IC auction protocol the seller can just send the buyer her type (aka the random seed).

[11] Note that this is a maximization problem, so *upper bound* on the approximation factor refers to an impossibility result.

[12] We require constant precision marginal contribution per item. For unit-demand, this translates to constant precision for any outcome. For gross substitutes, etc. this translates to $O(\log(n))$ bits to represent outcome utilities, which is still negligible.

[13] Note that it was an open question to obtain such a separation for *any* social choice rule, let alone a natural and important one like revenue-maximizing auctions.

Interestingly, this separation between the communication complexity of implementing and incentivizing optimal auctions holds in a more general sense (albeit for expected communication in randomized protocols): In the full version we show a *non-IC* auction protocol [14] that for *any* buyer with unit-demand (resp. combinatorial) valuations, the exactly optimal IC mechanism can be implemented by a randomized (non-IC) protocol using $O(\log(n))$ (resp. $O(n)$) communication.

1.5 Technical Highlights: Infinitely More Efficient Auction Protocols

Abstracting away the game theory and other detail, we explain the simple idea which is at the core of our main positive result (Theorems 3.1 and 4.1). Simplifying further, consider a randomized auction of just a single item: our goal is to compress the infinite deterministic communication complexity of a protocol where the buyer tells the seller exactly with what probability he expects to receive the item. Denote this probability of allocation by p. Given p, one way to allocate with probability p using unbiased coin tosses is to generate a uniformly random number $\tau \in [0,1]$ (whose binary representation is a uniformly random stream of bits after the decimal point), and to allocate the item iff $p > \tau$[15].

The key insight: for any fixed p, we don't actually need to know τ to infinite precision - we only need to know the prefix of τ's binary representation until the first bit on which it differs from p. Similarly, for a fixed τ, we only need to know p to the same precision. So here is our core protocol: draw[16] $\tau \in [0,1]$ uniformly at random, and ask the buyer to stream the binary representation of p - only with enough precision to determine whether $p > \tau$. Each time the buyer sends a bit from the binary representation of p it differs from the corresponding bit of τ with probability $1/2$; i.e. the protocol terminates with probability $1/2$ after each round. Hence we reduced the infinite deterministic protocol to one where the buyer only sends 2 bits in expectation.

What happens when we bring back incentives? It's not too hard to show that the protocol remains incentive-compatible as long as the buyer doesn't learn anything about τ until the end of the protocol. This is actually too good to be true, since the protocol length must depend on τ (otherwise it would be deterministic - and hence infinite), and the buyer must know whether the protocol is continuing in order to participate. Fortunately we can argue that if the only thing the buyer learns about τ is that the protocol is continuing, this information cannot help him cheat. Intuitively, he has already committed to the prefix of the protocol, and the extension of his strategy for the rest of the protocol is optimal conditioned on actually being asked to use it.

[14]The non-IC auction protocol is closer to [40]'s notion of *implementing* (as opposed to incentivizing) a mechanism, or to [5]'s definition of randomized communication complexity of auctions.

[15]For historical context, we remark that the setup up to this point is similar to the 1-bit public-coin protocol for single-item auctions in [5].

[16]Here and in all our protocols, τ can be drawn on the fly so the expected number of random bits is also bounded.

1.6 Technical Highlights: A Characterization of Randomized IC Auction Protocols

It is natural to try to prove communication lower bounds of IC auction protocols via a modular approach of: (i) use Game Theory to define a restricted communication problem that we have to solve in order to obtain near-optimal revenue; and then (ii) use standard techniques from Communication Complexity (e.g. a reduction from Set Disjointness). This approach has worked successfully in other applications of communication complexity to game theory (e.g. [34, 46, 55, 65]). However, our non-IC auction protocol in the full version formally precludes such a modular approach because there *is* an efficient communication protocol that exactly solves the game theoretic problem we are after. (In other words, the modular approach cannot separate the communication complexity of incentivizing and implementing a social choice rule.) Instead we need to simultaneously consider the complexity and incentives constraints, in particular we need to consider the joint evolution of the buyer's prior and incentives in an arbitrary randomized protocol.

Our main novel insight is the following simple characterization of incentive compatible communication protocols: In a general communication protocol, each buyer's node can partition the buyer's types in an arbitrary way. But for IC protocols, the buyer's next bit is fully determined by his respective value for the expected allocations conditioned on sending "0" or "1"; this means that it can only partition the buyer's types into halfspaces in valuation space (see Figure 2). Thus IC mechanisms are much less expressive.

The second part of the proof combines tools from Auction Theory and Error Correction Codes to construct, for each class of valuations, a family of priors whose (approximately) optimal mechanisms are all different. Finally, a simple counting argument shows that the total number of short IC protocols that satisfy our characterization is too small to cover all the different mechanisms.

1.7 Related Work

For general social choice settings, [40] define *binary dynamic mechanism (BDM)*, which formalizes the notion of communication protocol between multiple agents with outcomes and payments associated with the protocol-tree leaves. They contrast the communication complexity of *incentivizable* vs *implementable* BDMs. Our notion of IC auction protocols is equivalent to requiring that the protocol is incentivizable.

One subtle difference between our model and [40] is that the latter define BDMs as deterministic, while we focus on randomized protocols. In our context we can encode the seller's random number source as her type[17]. In this sense, our IC auction protocol is a special case of their *Bayesian incentive compatibility (BIC)-incentivizable BDM*. But this view misses the distinction between trusting a Bayesian prior about other agents valuations and behavior and merely trusting the source of randomness.

Our paper resolves an open question from [40] of separating the communication complexity of incentivizing and implementing *Bayesian* incentive compatible social choice rules. Very recently, [37, 73] resolved a different open question from the same paper about separating the communication complexity of incentivizing and

[17]Formally they only define finite BDMs, but they also discuss the natural infinite variant [40, Appendix B.1].

implementing *ex-post* incentive compatible social choice rules. [73] also separate the communication complexity of ex-post vs dominant strategy incentive compatibility.

Our paper exhibits a strong separation between the communication complexity associated with direct revelation and general mechanisms. Related separations have been shown before by [29] and [34] for social-welfare maximization with two or more strategic buyers. Specifically, [29] show an exponential gap between the communication complexity of direct revelation versus interactive mechanisms. [34] shows that in several important settings, the "taxation complexity" of deterministic mechanisms is approximately equivalent to the communication complexity, but exhibits an exponential gap between the two for truthful-in-expectation mechanisms. In contrast, we consider revenue maximization with a single strategic buyer and as few as two items. Arguably, the separation for a single strategic buyer in our settings is more surprising since he communicates with a seller who doesn't receive exogenous private information. More generally, communication complexity of (approximate) social welfare maximization in auctions with multiple strategic buyers has been extensively studied for combinatorial auctions [2, 3, 14, 15, 33, 34, 36, 39, 64, 77] and related settings [4, 11, 12, 35, 65].

Our paper is inspired by a discussion in [5] about the communication complexity of revenue-maximizing auctions. They prove that in general the deterministic communication complexity is equivalent (up to rounding) to the logarithm of the menu-size complexity. They also define a measure of randomized communication complexity of an auction, which is most closely related to [40]'s weaker notion of implementable protocols. They give a randomized protocol for implementing *any*[18] incentive-compatible auction for selling a single item using 1 bit of communication and (possibly infinitely many) public random coins.

Our protocols circumvent the intractability of exactly communicating payments (to infinite precision) by replacing them with random payments while preserving expectation. Related ideas have been used before in algorithmic mechanism design, e.g. by [1, 7].

The study of communication complexity in economics has its roots in classic works of [9] and [53]. Early mathematical formulations of the question were given by [54, 63, 69]. Outside of auctions, communication complexity has also been considered in AGT in the context of voting rules [19, 30, 68, 79] equilibrium computation [8, 28, 42, 43, 46, 47, 71] fair division [13, 67, 78], interdomain routing [59], and stable matching [45].

Since the seminal [49], menu-size complexity has been further studied by [5, 32, 44, 57, 75]. For a buyer with additive valuations over independent items, [5] prove[19] an $\log^{O(n)}(n)$ upper bound on the menu-size complexity of approximately-optimal mechanisms. In this special case, this translates to an upper bound of $O(n \log\log(n))$ on the deterministic communication complexity - slightly more efficient than our $O(n\log(n))$ upper bound on randomized communication complexity[20]. Our proof is arguably much simpler[21]. [44] explores the dependence on ε in the menu complexity of mechanisms with additive-ε-suboptimality in revenue; his main result, combined with [49], implies a $\Theta(\log(1/\varepsilon))$ bound on the deterministic communication complexity with two items.

For a buyer with unit-demand valuations over independent items [57] define a related notion of symmetric menu-size complexity which counts the number of lines up to symmetries, and prove an $n^{\text{polylog}(n)}$ upper bound on the symmetric menu-size complexity. We use a slightly stronger notion of partition-symmetric menu-size complexity (see Definition 5.2); the bound of [57] also holds for this stronger definition. We use this result for our nearly-revenue-optimal IC auction protocol. This provides further evidence that the relatively new notion of (partition)-symmetric menu complexity is a natural complexity measure for auctions.

Over the past decade, computational and menu-size complexity results of optimal auctions have motivated the design of suboptimal but simple mechanisms that approximate the optimal revenue [6, 18, 20–24, 27, 48, 51, 52, 56, 60, 72, 74, 74] or require resource augmentation [10, 38, 41, 61, 70]. Our results suggest that, in some cases, even strong menu-size complexity lower bounds do not preclude efficient optimal mechanisms.

2 MODEL AND DEFINITIONS

Our main notion in this paper is that of IC auction protocols:

Definition 2.1 ((IC) auction protocols)**.** An *auction protocol* consists of:

- A (possibly infinite) binary tree whose internal nodes are labeled either B (for Buyer) or C (for Chance).
- Each node of Chance has an associated probability distribution over its children.
- Each leaf node has an associated (non-negative) payment and (feasible) allocation.
- A suggested mapping from Buyer's types to Buyer's strategies, where a *Buyer's strategy* corresponds to a choice of child for each Buyer's node.

We say that an auction protocol is *finite* if it is guaranteed to terminate after a finite number of rounds with probability 1 for every Buyer's strategy. We say that an auction protocol is *individually rational (IR)* if the Buyer has a strategy that guarantees expected payment 0 and empty allocation. We say that an auction protocol is *IC* (in-expectation) if it is finite and IR, and if the Buyer weakly prefers the suggested Buyer's strategy corresponding to his type over any other strategy in the protocol.

The *expected communication complexity*, of an auction protocol is the expected depth of the leaf reached by a worst-case Buyer's strategy (and in expectation over nodes of Chance). Theorem 3.2 refers to the *expected Buyer's communication*, which only counts the number of Buyer's nodes on the path to the leaf.

[18] Note that the (revenue-)optimal auction for a single item is already deterministic and uses only 1 bit of communication.

[19] Theorem 1.2 of [5] states a slightly weaker bound of $n^{O(n)}$; the stronger bound is suggested in Footnote 3 of their paper.

[20] The results are incomparable: [5] uses deterministic communication, whereas our protocol gives *exact*-revenue-optimality and allows for *correlated* valuations. In particular, note that in our setting $O(n\log(n))$ is tight up to $O(\log n)$ factor, whereas for approximate revenue with independent valuations, the true answer (even for deterministic communication) is conjectured to be $O(\log(n))$ [5, Footnote 4].

[21] The main technical hurdle for [5] is a reduction to the case where the valuations are (almost) bounded by some large number $H = \text{poly}(n, \varepsilon)$ with only a negligible loss in revenue; we simply assume that the valuations are bounded by U, but it can be *arbitrarily large*. We remark that if we assume that the optimal mechanism obtains finite revenue (as is assumed in [5]; see Footnote 6 of their arXiv version), then it is easy to argue that for any $\varepsilon > 0$, capping the valuations by a sufficiently large $U(\varepsilon)$ preserves a $(1-\varepsilon)$-fraction of the revenue (see the full version for details).

Note that the buyer's strategy can be assumed wlog to be deterministic.

2.1 Valuation Classes

As is standard in the Algorithmic Mechanism Design literature, we consider buyers whose value for a bundle can be restricted to one of the following classes:

Definition 2.2 (Valuation classes)**.** A valuation function $v : 2^{[n]} \to \mathbb{R}_{\geq 0}$ may be restricted to one of the following classes:

Additive If it can be written as $v(S) = \sum_{i \in S} v_i$ for some item values v_i's.

Unit-demand If it can be written as $v(S) = \max_{i \in S} v_i$ for some item values v_i's.

Matroid-rank If, for some matroid M and item values v_i, it can be written as

$$v(S) = \max_{T \text{ is independent in } M} \sum_{i \in S \cap T} v_i.$$

XOS [22] If item value-vectors $\mathbf{v}_i$ of dimension d, it can be written as

$$v(S) = \max_{j \in \{1,\ldots,d\}} \sum_{i \in S} v_{i,j}.$$

The aforementioned classes are related to other well-studied classes like gross-substitutes, submodular, and subadditive in the following hierarchy:

$$\text{additive, unit-demand} \subset \text{matroid rank} \subset \text{gross substitutes} \subset \text{submodular} \subset \text{XOS} \subset \text{subadditive}.$$

The formal definition of gross substitutes, submodular, and subadditive is not important for our purposes; they are economically significant because they capture different natural notions of substitutability between items (see e.g. [58]).

In general, we are interested in any prior distribution over valuations of any of above-mentioned types. In particular, we also consider the notion of combinatorial valuations over independent items, which has been recently used by e.g. [17, 18, 23, 74].

Definition 2.3 (independent items [76])**.** A prior distribution $\mathcal{D}$ of valuations has a latent structure of independent items if there is a latent product distribution $\mathcal{D}_1 \times \mathcal{D}_2 \cdots \times \mathcal{D}_n$ with arbitrary support such that, a sample valuation v from $\mathcal{D}$ can be generated by first sampling a_i from $\mathcal{D}_i$ for all $i \in [n]$, and then for every $S \in [n]$, the value of $v(S)$ is uniquely determined by $\{a_i \mid i \in S\}$.

2.2 Menu-Size Complexity

Definition 2.4 (Menu-size complexity)**.** By the taxation principle, any mechanism can be canonically described by the expected allocation and payment for each type. This description induces a *menu*, or collection of *menu lines*, where each menu line is the expected allocation and payment for some type. The *menu-size complexity* of a mechanism is the number of distinct menu lines.

[22] XOS valuations are sometimes also called *fractionally subadditive*.

3 IC AUCTION PROTOCOLS FOR AN ADDITIVE BUYER

Theorem 3.1. *For any prior $\mathcal{D}$ of Buyer's additive valuations over n items bounded by maximum valuation U, and any truthful mechanism $\mathcal{M}$, there is an IC auction protocol with the same expected payment and allocation, using $O(n \log n)$ bits of communication.*

Proof. First, we convert $\mathcal{M}$ to a strategically-equivalent mechanism $\mathcal{M}'$ where the payment is always either zero or U. Note that by IR, the expected payment P in $\mathcal{M}$ for every type is always at most U; therefore for each type we can implement expected payment P by charging a payment of U with probability P/U (and zero otherwise). We henceforth identify each type of Buyer with the corresponding vector in $[0,1]^{n+1}$, which describes the probability that $\mathcal{M}'$ allocates each item to the Buyer, and the probability ($n+1$-th coordinate) that the Buyer pays U. We can further identify the mechanism $\mathcal{M}'$ with the set of allowed types/vectors in $[0,1]^{n+1}$.

Buyer's nodes and suggested strategy. Each Buyer's node[23] corresponds to a choice of $n+1$ bits. Given the Buyer's type and mechanism $\mathcal{M}'$, let $p_1, \ldots, p_n$ denote the probability that Buyer is allocated items $1, \ldots, n$, respectively, and let $p_{n+1} = P/U$ denote the probability that the Buyer pays U. The Buyer's suggested strategy is to send, for each round r and $i \in [n+1]$, the r-th bit in the binary representation[24] of p_i.

Correcting infeasible bits. We enforce that at any point in the protocol, the Buyer's messages are consistent with some type, i.e. with the prefix of probabilities corresponding to some feasible menu line in $\mathcal{M}'$. If only possible value for the Buyer's next bit would possibly be consistent with the protocol's history, the protocol continues assuming that the Buyer indeed sent this bit (formally we remove the Buyer's node from the protocol since it is redundant).

Nodes of Chance. The distribution over nodes of Chance is determined by an implicit parameter τ drawn uniformly at random from $[0,1]$. Before each node of Chance, we will already know that τ belongs to a particular measurable subset $S \subseteq [0,1]$. For a partition $S_L \cup S_R = S$ (to be specified below), each child of this node of Chance will correspond to τ falling in each of S_L or S_R, which induces the probability distribution on the children. While τ plays a crucial role in defining and analyzing the protocol, we stress that it is only implicit: in the actual protocol it is drawn on the fly, with increasing precision at each node of Chance along the path of the protocol.

To define the r-th node of Chance along a given path, consider, for each $i \in [n+1]$, the concatenation of the i-th bits across the Buyer's r messages, and compare it to the first r bits in the binary representation of τ. If for every i, at least one of the bits is different, the protocol is terminated at a leaf as follows (see Payment and Allocation). Otherwise, the protocol continues in a Buyer's node. Note that for each node of Chance, only one of its children is an internal (Buyer's) node.

[23] Here we slightly abuse notation: we defined the auction protocols for binary trees, so this technically corresponds to a sub-tree of depth $n+1$ with all Buyer's nodes.

[24] If p_i has two binary representations, using either one throughout the protocol will work.

Payment and Allocation. At the end of the protocol, for each $i \in [n+1]$, let $\widehat{p_i^r} \in [0,1)$ denote the number whose binary representation is the concatenation of the i-th bit in each of the r rounds of the protocol (after correcting infeasible bits). For $i \in [n]$, the i-th item is allocated iff $\widehat{p_i^r} > \tau$; the Buyer pays U iff $\widehat{p_{n+1}^r} > \tau$, and otherwise he pays zero.

IC. The key observation for incentive compatibility is that a Buyer's strategy is completely determined by the infinite stream of messages that it would send in the (zero-probability) event that the protocol never terminates. To see this, recall that each node of Chance has only one internal node child. Hence for any fixed Buyer's strategy there is a unique infinite path in the tree, and every finite run of the protocol corresponds to a prefix of this path, up to some node of Chance that deviates from the path to a leaf.

Let $\widehat{p_i}$ denote the number whose binary representation is the infinite sequence of Buyer's i-th bits in the (zero-probability) event that the protocol never terminates. Recall from the previous paragraph that a Buyer's strategy is completely determined by the vector of $\widehat{p_i}$'s. Note further that the i-th item is allocated at the end of the protocol iff $\widehat{p_i} > \tau$; similarly, the Buyer pays U iff $\widehat{p_{n+1}} > \tau$. Therefore, since τ is drawn uniformly from $[0,1]$, the probability that the Buyer is allocated item i (resp. pays U) is exactly $\widehat{p_i}$. Hence, by IC of $\mathcal{M}'$, the suggested strategy $\widehat{p} = p$ is optimal for the Buyer.

Communication complexity. At each round of communication, the Buyer sends $n+1$ bits. Also, at each round r of communication, there is probability exactly $1/2$ that the i-th bit in the Buyer's message (for each $i \in [n+1]$) disagrees with the r-th bit of τ. (This probability is independent across rounds, but correlated for different i's.) After $2\log(n)$ rounds, each i has probability $1/n^2$ of agreeing with all of τ's bits. We can take a union bound over all i's to obtain that except with probability $1/n$, the protocol has already terminated. In the unlikely event that the protocol continues, we can re-apply the same analysis from scratch.

Let r_{ub} denote an upper bound on the expected number of rounds in the protocol, corresponding to the worst case where the above union bound is tight. Then we have that

$$r_{\text{ub}} \leq 2\log(n) + r_{\text{ub}}/n. \tag{1}$$

Solving the recurrence relation for r_{ub}, we have that $r_{\text{ub}} = O(\log(n))$. Since the Buyer sends $n+1$ bits in each round, the total communication complexity is $O(n\log(n))$. □

Special case: a protocol for the [32] example. We can also prove a concrete (non-asymptotic) bound on the expected number of bits that the buyer sends in the DDT example. Beyond the historical importance of this specific example, our result demonstrates that our protocols are communication-efficient not only in the asymptotic sense, especially if we take advantage of the particular features of a specific distribution. We in particular highlight the fact that the Buyer in this protocol sends *strictly less* bits[25] than he would with a simple deterministic auction selling each item separately.

Theorem 3.2. *Consider the case of $n = 2$ items and the Buyer drawing his valuations i.i.d. from* Beta$(1,2)$ *(i.e. the distribution on $[0,1]$ with density function $f(x) = 2(1-x)$). Then there is an IC auction protocol obtaining the maximum possible revenue where the Buyer sends less than two bits in expectation.*

The proof is deferred to the full paper.

[25] Here we only count communication from the Buyer and not the random coin tosses. In many scenarios random bits are cheap but informative communication is costly.

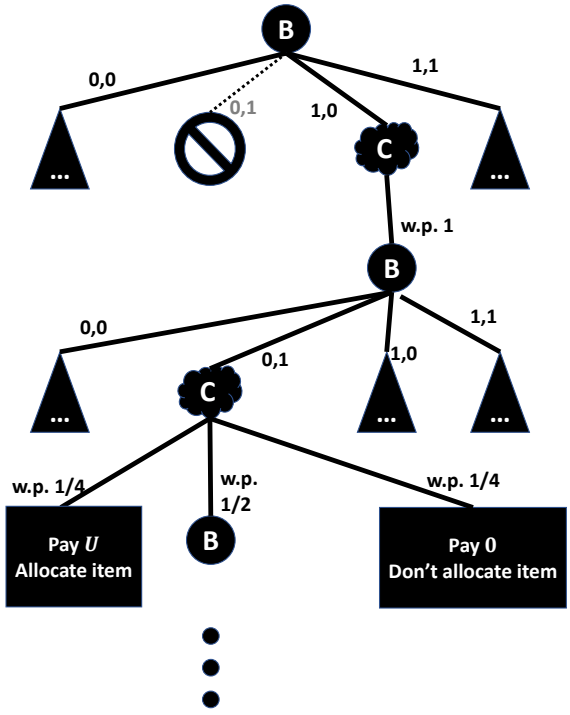

Figure 1: Example protocol
This figure depicts the first two iterations in an example protocol with one item, where the Buyer's favorite menu line has payment probability $2/3$ ($.\overline{10}$ in binary) and item allocated with probability $1/3$ ($.\overline{01}$ in binary). Nodes marked with B (resp. C) correspond to Buyer (resp. Chance). Triangles correspond to sub-trees never visited for this particular Buyer's valuation. In the first iteration, the Buyer sends $1, 0$, corresponding to the first bit in the probability of payment,allocation. Notice that $0, 1$ is an infeasible prefix for the Buyer since it would violate IC constraints (lower probability of payment and higher probability of allocation). At the first node of Chance, τ cannot disagree with both bits, hence the protocol proceeds to the next Buyer's node with probability 1. In the next iteration the Buyer sends the second bit from each probability. Finally, in the second node of Chance:

- The Buyer pays U and receives the item w.p. $1/4$ ($\tau < 1/4 < 1/3, 2/3$).
- The Buyer pays nothing and receives nothing w.p. $1/4$ ($\tau > 3/4 > 1/3.2/3$).
- W.p. $1/2$ the protocol continues.

4 AN EXTENSION FOR GENERAL VALUATIONS

The following theorem is an analogue of Theorem 3.1 for general combinatorial valuations (not necessarily subadditive or monotone). The communication complexity upper bound is parameterized by B, the number of bundles ever assigned by the direct revelation mechanism. For example, for unit demand valuations, $B \le n+1$; for general valuations, $B \le 2^n$.

Theorem 4.1. *Let $\mathcal{D}$ be any prior over Buyer's combinatorial valuations over n items bounded by maximum valuation U, and any truthful mechanism $\mathcal{M}$. Suppose that for any type and realization of randomness, $\mathcal{M}$ only ever allocates one of B bundles. Then there is an IC auction protocol with the same expected payment and allocation using $O(B\log(B))$ bits of communication.*

Proof sketch. For any type, consider a partition of $[0,1]$ into B intervals, where the b-th interval is of length identical to the probability that $\mathcal{M}$ allocates Bundle b to the Buyer. The rest of the proof proceeds analogously to the proof of Theorem 3.1. First, we transform $\mathcal{M}$ into a mechanism $\mathcal{M}'$ with payment 0 or U. We henceforth identify between a type and the $B-1$ probabilities that define the partition, and the probability that the Buyer pays U. The nodes of Chance are parameterized by a threshold τ drawn uniformly at random from $[0,1]$. At each round of communication the Buyer (allegedly) sends the next bit in the binary representation of each of the B probabilities that define his type. The protocol terminates when it has received enough information to determine in which of the B intervals τ lies and whether τ is smaller than the probability of payment. The allocation is the bundle corresponding to this interval, and the payment is U if τ is smaller than the probability of payment (and zero otherwise). □

5 UNIT-DEMAND, INDEPENDENT ITEMS: TRADING OFF REVENUE AND COMMUNICATION

Theorem 5.1. *Let $\mathcal{D}$ be a distribution of independent unit-demand valuations over n items bounded by maximum valuation U. Then, for any constant $\varepsilon > 0$, there is a $(1-\varepsilon)$-approximately revenue-optimal IC auction protocol using $\mathrm{polylog}(n)$ bits of communication.*

Our proof uses a result of [57] for *Partition-symmetric menus* which we introduce in Section 5.1. The proof of Theorem 5.1 is given in Section 5.2.

5.1 Partition-Symmetric Menu-Size Complexity

Symmetries. The following is a slight strengthening of the symmetric menu-size complexity measure recently introduced by [57].

Definition 5.2 (Partition-symmetric menu-size complexity)**.** A *partition-symmetric menu line* consists of a payment, (randomized) allocation, and a partitioning of items into subsets $S_1,\dots,S_\sigma$. We say that a direct revelation mechanism $\mathcal{M}$ supports this partition-symmetric menu line if its menu contains a line with the same payment for any permutation of the allocation that respects the partition (i.e. permutation π such that $\pi(S_i)=S_i$ for all i). The *partition-symmetric menu-size complexity* of $\mathcal{M}$ is the smallest c such that $\mathcal{M}$ can be written as the union of c partition-symmetric menu lines.

The following theorem follows from [57]; the statement here is slightly stronger than the formulation of Theorem IV.5 in their paper in the sense that (i) we consider the specific symmetry group induced by a partition of the items; and (ii) we require that the allocation probabilities are rounded to a discrete set L_δ. Both desiderata follow from their proof [81].

Theorem 5.3 ([57]). *Let $\mathcal{D}$ be a distribution of independent unit-demand valuations over n items. Then, for any constant $\varepsilon > 0$, there exists a unit-demand mechanism[26] with partition-symmetric menu-size complexity at most $n^{\mathrm{polylog}(n)}$ which recovers at least $(1-\varepsilon)$-fraction of the optimal revenue. Furthermore, for some constant $\delta > 0$ that depends on ε, the probabilities that the mechanism allocates each item always belong to the discrete set $L_\delta := \left\{1, 1-\delta, (1-\delta)^2, \dots, (1-\delta)^{\frac{3}{\delta}\ln n}\right\} \cup \{0\}$; in particular there are only $O(\log n)$ possible probabilities.*

5.2 Proof of Theorem 5.1

Proof. We begin with the partition-symmetric mechanism of [57] (see Theorem 5.3). Denote its partition-symmetric menu-size complexity by C. In the first stage of the protocol, the Buyer chooses a partition-symmetric menu line among C options, and then a subset S_i is drawn by Chance from the $\sigma \le n$ subsets in the partition. (Each subset S_i is drawn with probability equal to the sum of probabilities of items in that subset.) This first stage uses $O(\log n + \log C)$ communication. We henceforth focus on implementing the mechanism restricted to S_i. I.e. a mechanism whose menu has a fixed payment P and the set of feasible allocations is symmetric with respect to any permutation of S_i.

Since the set of feasible allocations is symmetric, it suffices to consider the histogram of allocation probabilities. The Buyer may assign each probability from the histogram to any item in S_i. Recall also that by Theorem 5.3, all the probabilities in the histogram belong wlog to a discrete set L_δ of $O(\log(n))$ feasible probabilities. In particular, the histogram can be described by $O(\log^2(n))$ bits (since the count for each probability is an integer between 0 and $|S_i| \le n$).

The second stage of the protocol proceeds by recursively considering smaller subsets of S_i. The nodes of Chance are parameterized by a number τ draws uniformly at random from $[0,1]$. At the first iteration, the Buyer's suggested strategy is to send the histogram of probabilities for the lexicographically first half of items in S_i. (This is equivalent to sending the histogram for the second half of the items since the total histogram is known.) If the sum of probabilities in the first half is greater than τ, the protocol recurses on the first half; otherwise it recurses on the second half. After $O(\log|S_i|)$ iterations, only one item is left. The Buyer is allocated that item and pays P.

[26] We say that a mechanism is *unit-demand* if it never allocates more than one item to the Buyer. (This is wlog for direct revelation mechanisms with unit-demand buyers. But in general, for mechanisms where the Seller does not fully learn the Buyer's valuation, it is not obvious how to convert a mechanism where she allocates a bundle of items to a unit-demand mechanism without increasing the partition-symmetric menu-size complexity.)

IC. We prove that the second stage of the protocol is IC and also has the same expected allocation and payment as in the original mechanism; IC of the first stage then follows from IC of the original mechanism. To show second-stage IC, let $\widehat{p_j}$ denote the probability that the Buyer assigns item j in the last iteration when it is not the only remaining item. Observe that any Buyer's strategy for the second stage is fully determined by the vector of $\widehat{p_j}$'s. By reverse induction over the iterations of the protocol, observe that the histogram of all $\widehat{p_j}$'s is exactly equal to the histogram of feasible probabilities. Finally note that at the end of the protocol, the Buyer is allocated item j with probability $\widehat{p_j}$. Therefore, by IC of the original protocol, the Buyer's suggested strategy is optimal.

Communication complexity. The first stage of the protocol requires $O(\log n + \log C)$ communication. Each iteration of the second stage requires $O(\log^2(n))$ bits to describe the histogram, and there are at most $O(\log(n))$ iterations. Hence the total communication complexity is $O(\log^3 n + \log C) = \text{polylog}(n)$. □

6 COMMUNICATION LOWER BOUND FOR UNIT-DEMAND VALUATIONS

We consider revenue maximization with unit-demand valuations as an example to demonstrate our proof technique. Our framework for constructing hard instances will rely on the design properties of a set system and a vector family, which are presented in the two lemmata in the following subsection.

6.1 Combinatorial Designs

Lemma 6.1. *For any constant $\varepsilon, \delta > 0$, there exists a family of size-εn subsets $\mathcal{X}_{n,\varepsilon,\delta} \subset \{0,1\}^n$ such that $|\mathcal{X}_{n,\varepsilon,\delta}| = 2^{\Omega(n)}$, and the intersection between any two distinct subsets $x_1, x_2 \in \mathcal{X}_{n,\varepsilon,\delta}$ has size at most $(1+\delta)\varepsilon^2 n$.*

Proof. The probabilistic proof is folklore and provided in the full paper. □

Lemma 6.2. *For any constant $\varepsilon > 0$ and large integer constant ℓ, let $\mathcal{R}_{\ell,\varepsilon}$ be the discrete distribution supported on $\{\varepsilon^{\ell-1}, \varepsilon^{\ell-2}, \ldots, 1\}$ such that $p^{(i)} \propto \varepsilon^{\ell-i}$, where we denote $p^{(i)} := \Pr[\varepsilon^{i-1}]$ (this is approximately the "equal-revenue distribution"). Then, for any constant $\eta > 0$, there exists a family of vectors $C_{N,\ell,\varepsilon,\eta} \subset \{\varepsilon^{\ell-1}, \varepsilon^{\ell-2}, \ldots, 1\}^N$ such that*

- *$|C_{N,\ell,\varepsilon,\eta}| = 2^{\Omega(N)}$,*
- *and moreover, for any $m = \omega(1)$ distinct vectors in $C_{N,\ell,\varepsilon,\eta}$, for all but η fraction of $j \in [N]$, for any $i \in [\ell]$, there are $(1 \pm \eta)p^{(i)}$ fraction of these m vectors whose j-th coordinates are ε^{i-1}.*

Proof. We construct $C_{N,\ell,\varepsilon,\eta}$ simply by independently sampling $2^{\delta N}$ vectors from product distribution $\mathcal{R}_{\ell,\varepsilon}^N$ for arbitrarily small constant $\delta > 0$, and we show that the desired properties hold with high probability. First, the probability that two random vectors have the same value at j-th coordinate is $p := \sum_{i\in[\ell]} p^{(i)} \cdot p^{(i)}$ for any j, and therefore, the probability that the two random vectors are exactly the same is p^N. For $\delta < \log(1/p)/2$, by a union bound over all the pairs of random vectors, every vector is distinct with high probability. Second, for any m random vectors, for any $i \in [\ell]$, $j \in [N]$, let $m_{i,j}$ be the number of vectors whose j-th coordinates are ε^{i-1} among the m random vectors, then by Chernoff bound,

$$\Pr[|m_{i,j} - p^{(i)}m| \geq \eta \cdot p^{(i)}m] \leq e^{-\eta^2 \cdot p^{(i)} m/3}.$$

By a union bound, the probability that there exists $i \in [\ell]$ such that $m_{i,j}$ is not within $(1 \pm \eta)p^{(i)}m$ is at most $\ell \cdot e^{-\eta^2 \cdot p^{(i)} m/3}$. It follows that for any fixed η fraction of $j \in [N]$, the probability that there exists $i \in [\ell]$ such that $m_{i,j}$ is not within $(1 \pm \eta)p^{(i)}$ for all j among the η fraction is at most $(\ell \cdot e^{-\eta^2 \cdot p^{(1)} m/3})^{\eta N}$ (notice that $p^{(1)}$ is the smallest among all $p^{(i)}$'s). By another union bound over all possible η fraction of $j \in [N]$, the probability that the second property in the statement is violated for m random vectors is at most

$$\binom{N}{\eta N} \cdot (\ell \cdot e^{-\eta^2 \cdot p^{(1)} m/3})^{\eta N} \leq (e/\eta)^{\eta N} \cdot (\ell \cdot e^{-\eta^2 \cdot p^{(1)} m/3})^{\eta N}$$
$$= ((e/\eta) \cdot \ell \cdot e^{-\eta^2 \cdot p^{(1)} m/3})^{\eta N},$$

which is $e^{-\theta m N}$ for some constant θ that does not depend on δ. Since there are $\binom{2^{\delta N}}{m} \leq (e \cdot 2^{\delta N}/m)^m \leq e^{\delta m N}$ distinct subsets of m random vectors of $C_{N,\ell,\varepsilon,\eta}$, by union bound, for $\delta < \theta$, for any fixed m, the second property in the statement is violated with probability at most $e^{-(\theta-\delta)mN}$. Finally, the proof finishes by taking a union bound over all $m = \omega(1)$, namely, $\sum_{m=\omega(1)} e^{-(\theta-\delta)mN} = o(1)$. □

6.2 The Main Lower Bound Result

Now we prove the following lower bound result for communication complexity of approximate revenue maximization with unit-demand valuations. Specifically, we construct a family of priors and show that most priors are hard for all low-communication (almost) truthful-in-expectation randomized protocols to approximately maximize revenue.

Theorem 6.3. *For every constant $\tau > 0$, any τ-approximate (almost) truthful-in-expectation protocol for revenue maximization, where the seller has n items, and the buyers have unit-demand valuations, requires $\Omega(n)$ bits of communication in expectation.*

Proof. We first construct a family of prior distributions of the buyers' valuations and then argue that in order to achieve any constant approximation, a protocol tree (which we will elaborate shortly) can not be shared by many prior distributions, which implies the communication complexity lower bound by a counting argument.

Construction. For arbitrarily tiny constants $\varepsilon_1, \varepsilon_2, \delta_1, \eta > 0$ and large integer constant ℓ such that $\eta, \varepsilon_1(1+\delta_1) \ll \varepsilon_2^\ell$, we take the set family $\mathcal{X}_{n,\varepsilon_1,\delta_1}$ from Lemma 6.1 and let $N := |\mathcal{X}_{n,\varepsilon_1,\delta_1}| = 2^{\Omega(n)}$, and then, we take the vector family $C_{N,\ell,\varepsilon_2,\eta}$ from Lemma 6.2 with $|C_{N,\ell,\varepsilon_2,\eta}| = 2^{\Omega(N)} = 2^{2^{\Omega(n)}}$. We let each $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$ represent a subset of items. Notice that we can fix a one-to-one mapping between the coordinates of a vector in $C_{N,\ell,\varepsilon_2,\eta}$ and all the sets in $\mathcal{X}_{n,\varepsilon_1,\delta_1}$, and therefore, for any vector $c \in C_{N,\ell,\varepsilon_2,\eta}$, $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$, we can denote $c(x)$ as c's value at the coordinate that corresponds to x.

For each vector $c \in C_{N,\ell,\varepsilon_2,\eta}$, we construct a prior distribution $\mathcal{D}_c$ of the buyers' valuations as follows — First, for each $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$,

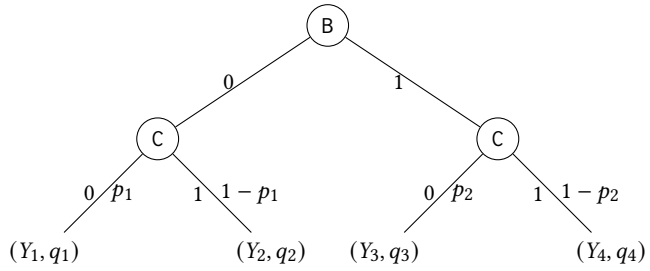

Figure 2: A depth-2 protocol tree.

we define a unit-demand valuation $v_c^x : 2^{[n]} \to \mathbb{R}_{\geq 0}$ as follows:

$$v_c^x(S) := \begin{cases} 0 & x \cap S = \emptyset \\ c(x) & \text{otherwise.} \end{cases}$$

Then, we let $\mathcal{D}_c$ be the uniform distribution over v_c^x's for all $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$. Finally, the family of prior distributions is $\mathcal{F} = \{\mathcal{D}_c \mid c \in C_{N,\ell,\varepsilon_2,\eta}\}$.

Interpretation. The following interpretations might be helpful for reading the proof. Each $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$ corresponds to a set of items which are (equally) valuable to the buyer with valuation v_c^x. Each vector $c \in C_{N,\ell,\varepsilon_2,\eta}$ specifies for each $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$ how valuable such an item is to the buyer with valuation v_c^x. By the design property of $\mathcal{X}_{n,\varepsilon_1,\delta_1}$, every $v_c^{x_1}, v_c^{x_2}$ with distinct x_1, x_2 are interested in mostly different items. By the design property of $C_{N,\ell,\varepsilon_2,\eta}$, for a large number of valuations v_c^x's with distinct c's but the same x, the values of an item in x to these valuations are distributed roughly according to the "equal revenue distribution" $\mathcal{R}_{\ell,\varepsilon_2}$ defined in Lemma 6.2.

An optimal truthful-in-expectation protocol for the hard instances. The first step for proving the lower bound is to show that there is a truthful-in-expectation protocol that extracts the full welfare using $O(n)$ bits of communication for the family of Bayesian instances constructed above. The protocol is as follows: the buyer sends the set x that corresponds to his valuation v_c^x to the seller, which takes n bits, and then, if $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$ (otherwise the seller stops), the seller samples an item i from set x uniformly at random and gives the item i to the buyer and charges him $c(x)$, where c corresponds to the prior $\mathcal{D}_c$. This protocol is obviously individual rational and revenue maximizing if the buyer tells the truth. To show truthfulness in expectation, suppose the buyer's true set of interest is x; without loss of generality, we can assume that the buyer sends some $x' \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$, because otherwise, the seller stops, and the buyer gets net utility 0, which is not better than telling the true x. Moreover, if the buyer sends $x' \neq x$, by the design property of $\mathcal{X}_{n,\varepsilon_1,\delta_1}$, he receives an item in x with probability at most $\varepsilon_1(1+\delta_1)$. Hence in expectation, the net utility is at most $\varepsilon_1(1+\delta_1)c(x) - c(x') \leq \varepsilon_1(1+\delta_1) - \varepsilon_2^{\ell-1} < 0$, where the first inequality is due to $c(x) \leq 1$ and $c(x') \geq \varepsilon_2^{\ell-1}$, and the second is due to our choice of parameters. Thus, sending x instead of x' is strictly better in expectation.

Representing a protocol as a protocol tree per prior distribution. Observe that once the prior distribution is fixed, a protocol can be viewed as a protocol tree. See Figure 2 for example. Without loss of generality, the protocol tree starts with the root B representing the buyer's round and then alternates between the buyer B and the seller C (Chance). At each round, represented by a node, the buyer or the seller can choose to send a bit 0, represented by left edge, or bit 1, represented by right edge, to the other. At a leaf, both players agree on a set of items Y allocated to the buyer and a payment q to the seller. The protocol is possibly randomized, and hence, at a seller's round, the seller[27] can send bit 0 with probability p and send bit 1 with probability $1-p$, which are represented by the weights on the edges. At a buyer's round, the buyer's strategy depends on his valuation, but we can assume without loss of generality that the buyer always deterministically chooses a bit to send, because the buyer is strategic and hence sending the bit that has better net utility in expectation (sending the bit that maximizes the seller's revenue if both choices are (almost) equal, and sending bit 0 if it is still a tie) is a (almost) dominant strategy for the buyer that maximizes the seller's revenue among all (almost) dominant strategies. Therefore, the buyer's prescribed (almost) dominant strategy can be deterministically decided by the protocol tree and his valuation.

To make the proof easier, we show that we can without loss of generality assume some nice properties for the protocol trees, and we will only consider such protocol trees afterwards.

Claim 6.4. *Any (almost) truthful-in-expectation protocol with $O(k)$ communication in expectation for our hard instance can be changed (with arbitrarily small loss of the approximation factor) such that*

- *the protocol tree has $O(k)$ depth,*
- *and moreover, the payment at any leaf of the protocol tree is $2^{O(k)}$.*

Suppose a protocol uses αk bits of communication in expectation where α is a positive constant. For an arbitrarily large constant β, by Markov's inequality, the protocol takes $\geq \beta k$ communication with probability at most $\gamma := \alpha/\beta$. Observe that if we trim all the nodes at level $\geq \beta k$ of the protocol tree $\mathcal{T}$, the buyer's expected utility (before payment) is at least $1-\gamma$ fraction of that for $\mathcal{T}$ (for our instance, the loss is at most γ). If we further trim every node that is reached with probability $\leq 4^{-\beta k}$ for any buyer, the buyer's expected utility loses at most another $2^{-\beta k}$, because there are at most $2^{\beta k}$ nodes left after the first trimming step. Since we introduce new leaves after trimming, we need to specify the allocation and the payment for each of them. For each new leaf, we simply let its allocation be the empty set, and we let its payment be the least possible expected payment at this node in $\mathcal{T}$ (that is, the minimum expected payment achieved by the worst possible buyer's responses in the subtree rooted at this node in $\mathcal{T}$).

After the above changes, the first property obviously holds for the new protocol tree $\mathcal{T}'$, and the second also holds, because if any leaf has payment larger than $4^{\beta k}$, then the probability of reaching that leaf (or node) in the original $\mathcal{T}$ for any buyer is at most $4^{-\beta k}$ (otherwise the expected payment is greater than 1 for a buyer that reaches this node with probability $> 4^{-\beta k}$, which exceeds the largest possible buyer's value and hence violates individual

[27] We assume that the seller is not strategic in the private-coin model. In the public-coin model, the seller can not be strategic, because his responses can be inferred from the public randomness and the pre-specified protocol tree, and thus, he can keep silent unless he observes that the buyer is cheating.

rationality), and this leaf should have been trimmed. It remains to show that the approximation factor is decreased arbitrarily little by the above changes.

To prove this, consider the buyer's (almost) dominant strategy s^* in $\mathcal{T}$, we change the strategy in the way that the buyer makes the same response as s^* at every node that will not be trimmed by the above steps and makes the worst possible responses (which minimize the expected payment) in the subtree rooted at every node that will be trimmed. This results in a new strategy s that gives the buyer almost the same expected utility as s^* (as we have shown, the loss is at most $\gamma + 2^{-\beta k}$). It follows that the expected payment for s can only be $\gamma + 2^{-\beta k}$ (plus another negligible error if the original protocol is only almost truthful) less than that for s^*, since otherwise the expected net utility of s is significantly better than s^*. Moreover, the expected payment for s can not be more than that for s^* by definition of s, and hence, the expected net utility of s is same as that of s^* up to negligible error. Furthermore, observe that s (ignore s's responses at trimmed nodes) gets the same expected utility and payment for the buyer in $\mathcal{T}'$. If s is an almost dominant strategy in $\mathcal{T}'$ (which indeed is as we will show), then we are done because we have shown the expected payment for s in $\mathcal{T}'$ (or $\mathcal{T}$) is same (up to negligible error) as that for s^* in $\mathcal{T}$.

To see s is an almost dominant strategy in $\mathcal{T}'$, suppose for contradiction there is another strategy s' with an non-negligible improvement of expected net utility over s in $\mathcal{T}'$. We extend s' to a strategy for $\mathcal{T}$ by letting it make worst possible response (which minimize the expected payment) for the nodes that will be trimmed in $\mathcal{T}'$. Note that the extended s' has the same expected net utility in $\mathcal{T}$ as that in $\mathcal{T}'$, which is significantly better than s's expected net utility in $\mathcal{T}'$ (and hence s or s^*'s expected net utility in $\mathcal{T}$). This contradicts that s^* is a (almost) dominant strategy in $\mathcal{T}$.

One protocol tree can not be shared by many priors. Now we show the main claim that leads to the lower bound result.

Claim 6.5. *For any constant $\tau > 0$ and any $m = \omega(1)$, any single protocol tree can only achieve τ approximation on $\le m$ priors in $\mathcal{F}$.*

Assume for contradiction that there are $m = w(1)$ priors $\mathcal{D}_{c_1}, \ldots, \mathcal{D}_{c_m}$ in $\mathcal{F}$ sharing the same protocol tree. By Lemma 6.2, for all but η fraction of $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$, the empirical distribution of $c_i(x)$'s for $i \in [m]$ is close to $\mathcal{R}_{\ell,\varepsilon_2}$ defined in Lemma 6.2, namely, the number of i's such that $c_i(x) = \varepsilon_2^{t-1}$ is $(1 \pm \eta)p^{(t)}m$, where $p^{(t)} \propto \varepsilon_2^{\ell-t}$. In the rest of the proof of Claim 6.5, we show that for any such x, the average revenue over valuations $v^x_{c_i}$ for all $i \in [m]$ achieved by the protocol tree is at most $\theta := \frac{(\varepsilon_2+1/\ell)(1+\eta)}{1-\eta}$ fraction of the optimum. Notice that θ is a constant that we can make arbitrarily small. This will finish the proof of the claim, because for at least one of $\mathcal{D}_{c_1}, \ldots, \mathcal{D}_{c_m}$, the protocol tree achieves no more than the average of the expected revenues for $\mathcal{D}_{c_1}, \ldots, \mathcal{D}_{c_m}$, which is at most $\tau = \theta + \frac{\eta\varepsilon_2^{1-\ell}}{1-\eta}$ fraction of the optimal revenue (we generously assume that it achieves full revenue on the η fraction of $x \in \mathcal{X}_{n,\varepsilon_1,\delta_1}$ that is excluded from the above analysis, and the full revenue for any x from this η fraction is at most 1, which is at most $\varepsilon_2^{1-\ell}$ times the full revenue of any x' from the other $1-\eta$ fraction), and $\frac{\eta\varepsilon_2^{1-\ell}}{1-\eta}$ is arbitrarily small by our choice of parameters.

Now consider any such x that the empirical distribution of $c_i(x)$'s for $i \in [m]$ is close to $\mathcal{R}_{\ell,\varepsilon_2}$, and let C_t be the set of c_i's with $c_i(x) = \varepsilon_2^{t-1}$. Without loss of generality, the buyers with valuation $v^x_{c_{j_t}}$ for all $c_{j_t} \in C_t$ use the same dominant strategy. Moreover, consider any $c_{j_t} \in C_t$ and any $c_{j_{t+1}} \in C_{t+1}$, we denote the expected utility and payment achieved by the prescribed dominant strategy for $v^x_{c_{j_t}}$ by u_t and q_t, respectively, and analogously, we denote u_{t+1} and q_{t+1} for $v^x_{c_{j_{t+1}}}$. If the buyer with valuation $v^x_{c_{j_t}}$ plays the strategy for $v^x_{c_{j_{t+1}}}$ instead, he will get expected utility u_{t+1}/ε_2 and payment q_{t+1}, because by definition $v^x_{c_{j_t}} = v^x_{c_{j_{t+1}}}/\varepsilon_2$. By definition of (almost) dominant strategy, we have the following inequality (the inequality holds approximately when we consider almost truthful-in-expectation protocols, and the error is negligible to the later derivations)

$$\frac{u_{t+1}}{\varepsilon_2} - q_{t+1} \le u_t - q_t. \tag{2}$$

Moreover, by individual rationality,

$$q_{t+1} \le u_{t+1}, \tag{3}$$

and it follows that

$$\begin{aligned} q_t &\le u_t - \frac{u_{t+1}}{\varepsilon_2} + q_{t+1} && \text{(Rearranging Eq. (2))} \\ &\le u_t - \frac{u_{t+1}}{\varepsilon_2} + u_{t+1} && \text{(By Eq. (3))} \\ &= u_t - u_{t+1}\left(\frac{1}{\varepsilon_2} - 1\right). \end{aligned} \tag{4}$$

Furthermore, because $c_i(x)$'s for $i \in [m]$ are distributed like $\mathcal{R}_{\ell,\varepsilon_2}$, the sum of the revenues obtained from the $v^x_{c_i}$'s for all $i \in [m]$ is at most (up to a $(1+\eta)$ multiplicative error)

$$\begin{aligned} \sum_{t=1}^{\ell} mp^{(t)}q_t &\le \sum_{t=1}^{\ell-1} mp^{(t)}\left(u_t - u_{t+1}\left(\frac{1}{\varepsilon_2} - 1\right)\right) + mp^{(\ell)}u_\ell \\ &\quad \text{(By Eq. (4) and Eq. (3))} \\ &= mp^{(1)}u_1 + m\sum_{t=2}^{\ell} u_t\left(p^{(t)} - \frac{p^{(t-1)}}{\varepsilon_2} + p^{(t-1)}\right) \\ &\quad \text{(Rearranging the sum)} \\ &= mp^{(1)}u_1 + m\sum_{t=2}^{\ell} u_t p^{(t-1)} \\ &\quad \text{(By definition of } p^{(t)}\text{)} \\ &= mp^{(1)}u_1 + m\varepsilon_2\sum_{t=2}^{\ell} u_t p^{(t)} \\ &\le mp^{(1)} + m\varepsilon_2\sum_{t=2}^{\ell} p^{(t)}\varepsilon_2^{t-1}, \\ &\quad \text{(By } u_t \le \varepsilon_2^{t-1}\text{)} \end{aligned}$$

which is at most ε_2 fraction of $\sum_{t=1}^{\ell} mp^{(t)}\varepsilon_2^{t-1}$ plus $mp^{(1)}$, but $mp^{(1)}$ is only $1/\ell$ fraction of $\sum_{t=1}^{\ell} mp^{(t)}\varepsilon_2^{t-1}$ by its definition. Because $c_i(x)$'s for $i \in [m]$ are distributed like $\mathcal{R}_{\ell,\varepsilon_2}$, the optimal total revenue we can get from all the $v^x_{c_i}$ for $i \in [m]$ (which is equal to their total value) is at least $(1-\eta)\sum_{t=1}^{\ell} mp^{(t)}\varepsilon_2^{t-1}$, and hence, the

average revenue achieved by the protocol tree on valuations $v^x_{c_i}$ for $i \in [m]$ is at most $\frac{(\varepsilon_2+1/\ell)(1+\eta)}{1-\eta}$ fraction of the optimum.

Finishing the proof by a counting argument. For any constant $\tau > 0$, suppose that the communication complexity of a τ-approximate truthful-in-expectation protocol is $k = o(n)$, and without loss of generality we assume that the protocol always uses up k bits. We count how many protocol trees we can have. Note that a protocol tree is determined by the (Y, q) pairs on the leaves and the probabilities on the edges. Without loss of generality, we can assume that the payments and the probabilities have finite precision, namely, the probabilities are rounded to $\{i/4^n \mid i = 0, 1, \ldots, 4^n\}$, and the payments are rounded to $\{i/4^n \mid i = 0, 1, \ldots, 2^{O(n)}\}$. To see this, first observe that rounding can only change the payment at any leaf by at most $1/4^n$, and similarly, it can only change the probability of reaching any leaf by $O(1/4^n)$, and therefore, it only changes the expected utility and the expected payment for the buyer by at most $O(2^k/4^n) = O(1/2^n)$. As we have noted along the proof, the analysis works for almost truthful-in-expectation protocols, which tolerates this extra $O(1/2^n)$ error.

Therefore, there are at most 2^n choices of Y and at most $2^{O(n)}$ choices of q, which implies at most $2^{O(n)}$ choices of (Y, q) at each leaf, and there are at most 4^n choices of the probability on each edge. Since the depth of the protocol tree is no more than k, there are 2^k leaves and 2^{k+1} edges at most. Altogether, there are at most $(2^{O(n)})^{2^k} \cdot (4^n)^{2^{k+1}} = 2^{2^{k+o(n)}}$ possible protocol trees. Furthermore, by Claim 6.5, these protocol trees can only beat τ-approximation on at most $2^{2^{k+o(n)}} \cdot m$ priors in total for any $m = \omega(1)$, but there are $2^{2^{\Omega(n)}}$ priors in $\mathcal{F}$. Hence, most priors in $\mathcal{F}$ are hard for all the $o(n)$-communication protocols.

□

6.3 Separating the Complexity of Implementing and Incentivizing

Remark 6.6. There is an $O(\log n)$-communication implementation of the optimal protocol for our hard instances. Combining with the lower bound, this shows an exponential separation between communication complexity of almost truthful-in-expectation implementation and that of non-truthful implementation for this protocol, even when the buyer's valuation has constant precision.

Proof. A more communication-efficient non-truthful implementation is that the buyer randomly chooses an item i of interest and sends i and $c(x)$ to the seller, and then the seller gives the item i to the buyer and charges the buyer $c(x)$, which only uses $O(\log n)$ bits of communication. □

6.4 Extending to Gross-Substitutes Valuations

Our technique for proving lower bound for unit-demand valuations can be applied to gross-substitutes (and XOS) valuations. We state the result below, the proof of which can be found in the full paper.

Theorem 6.7. *For every constant $\tau > 0$, any τ-approximate (almost) truthful-in-expectation protocol for revenue maximization, where the seller has n items, and the buyers have gross substitutes valuations, requires $2^{\Omega(n^{1/3})}$ bits of communication in expectation. For XOS valuations with n items, the communication complexity lower bound for any constant approximation can be improved to $2^{\Omega(n)}$.*

7 COMMUNICATION LOWER BOUND FOR XOS VALUATIONS WITH INDEPENDENT ITEMS

In this section, we sketch the construction of hard instances to show that beating 4/5 approximation for XOS valuations with independent items requires exponential communication. The complete proof can be found in the full paper. Note that constant-factor approximation is known (e.g., [74]) for more general subadditive valuations with independent items.

Theorem 7.1. *For every constant $\tau > 0$, any $(\frac{4}{5} + \tau)$-approximate (almost) truthful-in-expectation protocol for revenue maximization, where the seller has n items, and the buyers have XOS valuations with independent items, requires $2^{\Omega(n)}$ bits of communication in expectation.*

Proof sketch. The proof follows the same strategy as the proof of the previous lower bounds. First, we construct a family of prior distributions of XOS valuations with independent items. We focus on the following special case of prior distributions of XOS valuations with independent items — Given any integer b, for each item $i \in [n]$ there is a distribution $\mathcal{D}_i$ over $\mathbb{R}^b_{\geq 0}$, an XOS valuation v is generated by first sampling a vector $a^{(i)}$ from each $\mathcal{D}_i$ and then defined as

$$v(S) = \max_{j \in [b]} \sum_{i \in S} a^{(i)}_j.$$

In this case, $\mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$ specifies a prior distribution of XOS valuations.

Construction. Let $\varepsilon_0, \varepsilon_1, \delta_1, \eta, \gamma > 0$ be arbitrarily tiny constants such that $\varepsilon_1(1 + \delta_1) + \varepsilon_0(1 + \delta_0) < \frac{1}{2-\gamma} - \frac{1}{2}$. Given a set family $\mathcal{X}_{n-1,\varepsilon_0,\delta_0}$ from Lemma 6.1, we let $b := |\mathcal{X}_{n-1,\varepsilon_0,\delta_0}| = 2^{\Omega(n)}$. We can think of each set in $\mathcal{X}_{n-1,\varepsilon_0,\delta_0} = \{x^{(1)}, x^{(2)}, \ldots, x^{(b)}\}$ as a binary vector. For each $i \in [n-1]$, we let $\mathcal{D}_i$ be the trivial distribution with singleton support $\{a^{(i)}\}$, where $a^{(i)} \in \mathbb{R}^b_{\geq 0}$ is defined as $a^{(i)}_j = \frac{x^{(j)}_i}{(2-\gamma)\varepsilon_0(n-1)}$ for all $j \in [b]$. Now we take another set family $\mathcal{X}_{b,\varepsilon_1,\delta_1} = \{y^{(1)}, y^{(2)}, \ldots, y^{(N)}\}$ from Lemma 6.1 and a vector family $C_{N,2,\frac{1}{2},\eta} = \{c^{(1)}, c^{(2)}, \ldots, c^{(M)}\}$ from Lemma 6.2, where $N := |\mathcal{X}_{b,\varepsilon_1,\delta_1}| = 2^{\Omega(b)}$ and $M := |C_{N,2,\frac{1}{2},\eta}| = 2^{\Omega(N)}$. For each $c^{(i)}$, we let $\mathcal{D}^{c^{(i)}}_n$ be the uniform distribution over $\{\frac{c^{(i)}_j+1}{2} \cdot y^{(j)} \mid j \in [N]\}$. The family of prior distributions is $\mathcal{F} = \{\mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}^{c^{(i)}}_n \mid i \in [M]\}$. For each prior, a valuation is sampled according to the procedure described in the previous paragraph, and specifically, a valuation function $v^{y^{(j)}}_{c^{(i)}}$, determined by $c^{(i)}$ and $y^{(j)}$, is given as follows

$$v^{y^{(j)}}_{c^{(i)}}(S) = \max_{t \in [b]} \mathbb{1}\{n \in S\} \cdot \frac{c^{(i)}_j + 1}{2} \cdot y^{(j)}_t + \sum_{r \in S \setminus \{n\}} \frac{x^{(t)}_r}{(2-\gamma)\varepsilon_0(n-1)}.$$

Interpretation. In this instance, any valuation $v^{y^{(j)}}_{c^{(i)}}$, when restricted to items $[n-1]$, becomes a single scaled binary XOS valuation in which the clauses correspond to the scaled binary vectors $x^{(1)}, x^{(2)}, \ldots, x^{(b)}$ (they represent pairwise nearly disjoint subsets (of items in $[n-1]$) that are equally valuable to every buyer), and each of these clause has total value $\frac{1}{2-\gamma}$. Each binary vector $y^{(j)}$ then decides which of these clauses $x^{(1)}, x^{(2)}, \ldots, x^{(b)}$ interact with the item n, i.e., the item n has positive contribution to the clause $x^{(t)}$ in the valuation $v^{y^{(j)}}_{c^{(i)}}$ iff $y^{(j)}_t = 1$. (Distinct $y^{(j)}$'s define almost completely different interactions.) Each binary vector $c^{(i)}$ then specifies for each $y^{(j)}$ how large the contribution of n is for each clause where, according to $y^{(j)}$, the item n has positive contribution, i.e., the item n contributes value 1 to every clause $x^{(t)}$ it interacts with (i.e., for which $y^{(j)}_t = 1$) in the valuation $v^{y^{(j)}}_{c^{(i)}}$ if $c^{(i)}_j = 1$ and contributes value 0 if otherwise. For a large number of valuations $v^{y^{(j)}}_{c^{(i)}}$'s with distinct $c^{(i)}$'s but the same $y^{(j)}$, the contributions of the item n in these valuations (to every clause it interacts with) are distributed roughly according to the "equal revenue distribution" $\mathcal{R}_{2,1/2}$. □

REFERENCES

[1] Aaron Archer, Christos H. Papadimitriou, Kunal Talwar, and Éva Tardos. 2003. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. *Internet Math.* 1, 2 (2003), 129–150. https://doi.org/10.1080/15427951.2004.10129086

[2] Sepehr Assadi. 2017. Combinatorial Auctions Do Need Modest Interaction. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017.* 145–162. https://doi.org/10.1145/3033274.3085121

[3] Sepehr Assadi, Hrishikesh Khandeparkar, Raghuvansh R Saxena, and S Matthew Weinberg. 2020. Separating the communication complexity of truthful and non-truthful combinatorial auctions. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing.* 1073–1085.

[4] Moshe Babaioff, Liad Blumrosen, and Michael Schapira. 2013. The communication burden of payment determination. *Games Econ. Behav.* 77, 1 (2013), 153–167. https://doi.org/10.1016/j.geb.2012.08.007

[5] Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. 2017. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017.* 869–877. https://doi.org/10.1145/3055399.3055426

[6] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2014. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014.* 21–30. https://doi.org/10.1109/FOCS.2014.11

[7] Moshe Babaioff, Robert D. Kleinberg, and Aleksandrs Slivkins. 2015. Truthful Mechanisms with Implicit Payment Computation. *J. ACM* 62, 2 (2015), 10:1–10:37. https://doi.org/10.1145/2724705

[8] Yakov Babichenko and Aviad Rubinstein. 2017. Communication complexity of approximate Nash equilibria. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017.* 878–889. https://doi.org/10.1145/3055399.3055407

[9] Chester Irving Barnard. 1938. *The Functions of the Executive.* Harvard University Press.

[10] Hedyeh Beyhaghi and S. Matthew Weinberg. 2019. Optimal (and Benchmark-Optimal) Competition Complexity for Additive Buyers over Independent Items. In *Proceedings of the 51st ACM Symposium on Theory of Computing Conference (STOC).*

[11] Liad Blumrosen and Michal Feldman. 2013. Mechanism design with a restricted action space. *Games Econ. Behav.* 82 (2013), 424–443. https://doi.org/10.1016/j.geb.2013.03.005

[12] Liad Blumrosen, Noam Nisan, and Ilya Segal. 2007. Auctions with Severely Bounded Communication. *J. Artif. Intell. Res.* 28 (2007), 233–266. https://doi.org/10.1613/jair.2081

[13] Simina Brânzei and Noam Nisan. 2019. Communication Complexity of Cake Cutting. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019.* 525. https://doi.org/10.1145/3328526.3329644

[14] Mark Braverman, Jieming Mao, and S. Matthew Weinberg. 2016. Interpolating Between Truthful and non-Truthful Mechanisms for Combinatorial Auctions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016.* 1444–1457. https://doi.org/10.1137/1.9781611974331.ch99

[15] Mark Braverman, Jieming Mao, and S. Matthew Weinberg. 2018. On Simultaneous Two-player Combinatorial Auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018.* 2256–2273. https://doi.org/10.1137/1.9781611975031.146

[16] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2010. Pricing Randomized Allocations. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010.* 585–597. https://doi.org/10.1137/1.9781611973075.49

[17] Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. 2016. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the 48th ACM Conference on Theory of Computation(STOC).*

[18] Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017.* 170–183. https://doi.org/10.1145/3055399.3055465

[19] Ioannis Caragiannis and Ariel D. Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artif. Intell.* 175, 9-10 (2011), 1655–1671. https://doi.org/10.1016/j.artint.2011.03.005

[20] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC).*

[21] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC).*

[22] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. 2015. The power of randomness in Bayesian optimal mechanism design. *Games and Economic Behavior* 91 (2015), 297–317. https://doi.org/10.1016/j.geb.2012.08.010

[23] Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016.* 579–596. https://doi.org/10.1145/2940716.2940756

[24] Shuchi Chawla, Yifeng Teng, and Christos Tzamos. 2019. Buy-Many Mechanisms are Not Much Better than Item Pricing. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019.* 237–238. https://doi.org/10.1145/3328526.3329583

[25] Xi Chen, Ilias Diakonikolas, Anthi Orfanou, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. 2015. On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015.* 1464–1479. https://doi.org/10.1109/FOCS.2015.93

[26] Xi Chen, Ilias Diakonikolas, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. 2014. The Complexity of Optimal Multidimensional Pricing. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014.* 1319–1328. https://doi.org/10.1137/1.9781611973402.97

[27] Yu Cheng, Nick Gravin, Kamesh Munagala, and Kangning Wang. 2018. A Simple Mechanism for a Budget-Constrained Buyer. In *Web and Internet Economics - 14th International Conference, WINE 2018, Oxford, UK, December 15-17, 2018, Proceedings.* 96–110. https://doi.org/10.1007/978-3-030-04612-5_7

[28] Vincent Conitzer and Tuomas Sandholm. 2004. Communication complexity as a lower bound for learning in games. In *Proceedings of the twenty-first international conference on Machine learning.* ACM, 24.

[29] Vincent Conitzer and Tuomas Sandholm. 2004. Computational criticisms of the revelation principle. In *Proceedings 5th ACM Conference on Electronic Commerce (EC-2004), New York, NY, USA, May 17-20, 2004.* 262–263. https://doi.org/10.1145/988772.988824

[30] Vincent Conitzer and Tuomas Sandholm. 2005. Communication complexity of common voting rules. In *Proceedings 6th ACM Conference on Electronic Commerce (EC-2005), Vancouver, BC, Canada, June 5-8, 2005.* 78–87. https://doi.org/10.1145/1064009.1064018

[31] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2014. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA).*

[32] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2017. Strong Duality for a Multiple-Good Monopolist. *Econometrica* 85, 3 (2017), 735–767.

[33] Shahar Dobzinski. 2016. Breaking the Logarithmic Barrier for Truthful Combinatorial Auctions with Submodular Bidders. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing* (Cambridge, MA, USA) *(STOC 2016).* ACM, New York, NY, USA, 940–948. https://doi.org/10.1145/2897518.2897569

[34] Shahar Dobzinski. 2016. Computational Efficiency Requires Simple Taxation. In *FOCS.*

[35] Shahar Dobzinski and Shaddin Dughmi. 2013. On the Power of Randomization in Algorithmic Mechanism Design. *SIAM J. Comput.* 42, 6 (2013), 2287–2304. https://doi.org/10.1137/090780146

[36] Shahar Dobzinski, Noam Nisan, and Sigal Oren. 2014. Economic Efficiency Requires Interaction. In *the 46th annual ACM symposium on Theory of computing (STOC)*.

[37] Shahar Dobzinski and Shiri Ron. 2020. The Communication Complexity of Payment Computation. *CoRR* abs/2012.14623 (2020). arXiv:2012.14623 https://arxiv.org/abs/2012.14623

[38] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 343. https://doi.org/10.1145/3033274.3085115

[39] Tomer Ezra, Michal Feldman, Eric Neyman, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2019. Settling the Communication Complexity of Combinatorial Auctions with Two Subadditive Buyers. In *the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.

[40] Ronald Fadel and Ilya Segal. 2009. The communication cost of selfishness. *J. Econ. Theory* 144, 5 (2009), 1895–1920. https://doi.org/10.1016/j.jet.2007.09.015

[41] Michal Feldman, Ophir Friedler, and Aviad Rubinstein. 2018. 99% Revenue via Enhanced Competition. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*. 443–460. https://doi.org/10.1145/3219166.3219202

[42] Anat Ganor and Karthik C. S. 2018. Communication Complexity of Correlated Equilibrium with Small Support. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*. 12:1–12:16. https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2018.12

[43] Anat Ganor, Karthik C. S., and Dömötör Pálvölgyi. 2019. On Communication Complexity of Fixed Point Computation. *CoRR* abs/1909.10958 (2019). arXiv:1909.10958 http://arxiv.org/abs/1909.10958

[44] Yannai A. Gonczarowski. 2018. Bounding the menu-size of approximately optimal auctions via optimal-transport duality. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*. 123–131. https://doi.org/10.1145/3188745.3188786

[45] Yannai A. Gonczarowski, Noam Nisan, Rafail Ostrovsky, and Will Rosenbaum. 2019. A stable marriage requires communication. *Games Econ. Behav.* 118 (2019), 626–647. https://doi.org/10.1016/j.geb.2018.10.013

[46] Mika Göös and Aviad Rubinstein. 2018. Near-Optimal Communication Lower Bounds for Approximate Nash Equilibria. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*. 397–403. https://doi.org/10.1109/FOCS.2018.00045

[47] Sergiu Hart and Yishay Mansour. 2010. How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games and Economic Behavior* 69, 1 (2010), 107–126.

[48] Sergiu Hart and Noam Nisan. 2017. Approximate revenue maximization with multiple items. *J. Economic Theory* 172 (2017), 313–347. https://doi.org/10.1016/j.jet.2017.09.001

[49] Sergiu Hart and Noam Nisan. 2019. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *J. Economic Theory* 183 (2019), 991–1029. https://doi.org/10.1016/j.jet.2019.07.006

[50] Sergiu Hart and Philip J. Reny. 2015. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics* 10, 3 (2015), 893–922.

[51] Sergiu Hart and Philip J. Reny. 2017. The Better Half of Selling Separately. *CoRR* abs/1712.08973 (2017). arXiv:1712.08973 http://arxiv.org/abs/1712.08973

[52] Jason D. Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*. 225–234.

[53] F. A. Hayek. 1945. The Use of Knowledge in Society. *The American Economic Review* 35, 4 (1945), 519–530. http://www.jstor.org/stable/1809376

[54] L. Hurwicz. 1960. *Optimality and Informational Efficiency in Resource Allocation Processes*. Stanford University Press. https://books.google.com/books?id=p4fwtgAACAAJ

[55] Nicole Immorlica, Brendan Lucier, Jieming Mao, Vasilis Syrgkanis, and Christos Tzamos. 2018. Combinatorial Assortment Optimization. In *Web and Internet Economics - 14th International Conference, WINE 2018, Oxford, UK, December 15-17, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 11316)*, George Christodoulou and Tobias Harks (Eds.). Springer, 218–231. https://doi.org/10.1007/978-3-030-04612-5_15

[56] Robert Kleinberg and S. Matthew Weinberg. 2012. Matroid prophet inequalities. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*. 123–136. https://doi.org/10.1145/2213977.2213991

[57] Pravesh Kothari, Divyarthi Mohan, Ariel Schvartzman, Sahil Singla, and S. Matthew Weinberg. 2019. Approximation Schemes for a Buyer with Independent Items via Symmetries. In *the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.

[58] Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games Econ. Behav.* 55, 2 (2006), 270–296. https://doi.org/10.1016/j.geb.2005.02.006

[59] Hagay Levin, Michael Schapira, and Aviv Zohar. 2011. Interdomain Routing and Games. *SIAM J. Comput.* 40, 6 (2011), 1892–1912. https://doi.org/10.1137/080734017

[60] Xinye Li and Andrew Chi-Chih Yao. 2013. On Revenue Maximization for Selling Multiple Independently Distributed Items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.

[61] Siqi Liu and Christos-Alexandros Psomas. 2018. On the Competition Complexity of Dynamic Mechanism Design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*. 2008–2025. https://doi.org/10.1137/1.9781611975031.131

[62] A. M. Manelli and D. R. Vincent. 2010. Bayesian and Dominant-Strrategy Implementation in the Independent Private-Values Model. *Econometrica* 78, 6 (2010), 1905–1938.

[63] Kenneth Mount and Stanley Reiter. 1974. The informational size of message spaces. *Journal of Economic Theory* 8, 2 (1974), 161 – 192. https://doi.org/10.1016/0022-0531(74)90012-X

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

[65] Christos H. Papadimitriou, Michael Schapira, and Yaron Singer. 2008. On the Hardness of Being Truthful. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.

[66] Gregory Pavlov. 2011. Optimal Mechanism for Selling Two Goods. *The B.E. Journal of Theoretical Economics* 11, 3 (2011).

[67] Benjamin Plaut and Tim Roughgarden. 2019. Communication Complexity of Discrete Fair Division. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*. 2014–2033. https://doi.org/10.1137/1.9781611975482.122

[68] Ariel D. Procaccia and Jeffrey S. Rosenschein. 2006. The communication complexity of coalition formation among autonomous agents. In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8-12, 2006*. 505–512. https://doi.org/10.1145/1160633.1160727

[69] Stefan Reichelstein. 1984. Incentive compatibility and informational requirements. *Journal of Economic Theory* 34, 1 (1984), 32 – 51. https://doi.org/10.1016/0022-0531(84)90160-1

[70] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. 2012. Supply-limiting mechanisms. In *13th ACM Conference on Electronic Commerce (EC)*.

[71] Tim Roughgarden and Omri Weinstein. 2016. On the Communication Complexity of Approximate Fixed Points.. In *Electronic Colloquium on Computational Complexity (ECCC)*, Vol. 23. 55.

[72] Aviad Rubinstein. 2016. On the Computational Complexity of Optimal Simple Mechanisms. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*. 21–28. https://doi.org/10.1145/2840728.2840736

[73] Aviad Rubinstein, Raghuvansh R. Saxena, Clayton Thomas, S. Matthew Weinberg, and Junyao Zhao. 2020. Exponential Communication Separations between Notions of Selfishness. *CoRR* abs/2012.14898 (2020). arXiv:2012.14898 https://arxiv.org/abs/2012.14898

[74] Aviad Rubinstein and S. Matthew Weinberg. 2018. Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity. *ACM Trans. Economics and Comput.* 6, 3-4 (2018), 19:1–19:25. https://doi.org/10.1145/3105448

[75] Raghuvansh R. Saxena, Ariel Schvartzman, and S. Matthew Weinberg. 2018. The menu complexity of "one-and-a-half-dimensional" mechanism design. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*. 2026–2035. https://doi.org/10.1137/1.9781611975031.132

[76] Gideon Schechtman. 2003. Chapter 37 - Concentration, Results and Applications. Handbook of the Geometry of Banach Spaces, Vol. 2. Elsevier Science B.V., 1603 – 1634. http://www.sciencedirect.com/science/article/pii/S187458490380044X

[77] Ilya Segal. 2007. The communication requirements of social choice rules and supporting budget sets. *J. Econ. Theory* 136, 1 (2007), 341–378. https://doi.org/10.1016/j.jet.2006.09.011

[78] Ilya R. Segal. 2010. Nash implementation with little communication. *Theoretical Economics* 5, 1 (2010), 51–71. https://doi.org/10.3982/TE576 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE576

[79] Travis C. Service and Julie A. Adams. 2012. Communication complexity of approximating voting rules. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*. 593–602. http://dl.acm.org/citation.cfm?id=2343781

[80] John Thanassoulis. 2004. Haggling over substitutes. *Journal of Economic Theory* 117 (2004), 217–245.

[81] S. Matthew Weinberg. 2020. Personal communication.