The Randomized Communication Complexity of Revenue Maximization

AVIAD RUBINSTEIN
Stanford University
and
JUNYAO ZHAO
Stanford University

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 [Rubinstein and Zhao 2021]. Motivated by the fact that revenue-optimal auctions are randomized [Thanassoulis 2004; Manelli and Vincent 2010; Briest et al. 2010; Pavlov 2011; Hart and Reny 2015] (as well as by an open problem of Babaioff, Gonczarowski, and Nisan [Babaioff et al. 2017]), 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 [Fadel and Segal 2009].

Categories and Subject Descriptors: J.4 [**Social and Behaviorial Science**]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: Revenue Maximization, Randomized Protocols, Communication Complexity, Approximation

1. INTRODUCTION

The central goal of Algorithmic Mechanism Design is to design mechanisms that guarantee good outcomes while taking into account both (i) the selfish agents' incentives and (ii) 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, optimal mechanisms are inherently complex, e.g. they involve randomized lotteries [Thanassoulis 2004; Manelli and Vincent 2010; Briest

Authors' addresses: `aviad@cs.stanford.edu`, `junyao@cs.stanford.edu`

et al. 2010; Pavlov 2011; Hart and Reny 2015] and are often computationally intractable [Daskalakis et al. 2014; Chen et al. 2014; Chen et al. 2015].

One particularly influential measure of complexity of mechanisms is the *menu-size complexity* of [Hart and Nisan 2019]: 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 [Daskalakis et al. 2017], 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.

[Daskalakis et al. 2017] 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 of Related work in [Rubinstein and Zhao 2021]).

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[1] 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 [Daskalakis et al. 2017]. However, [Babaioff et al. 2017] recently observed that, once the mechanism is announced, the deterministic communication complexity to implement it is equal (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. [Babaioff et al. 2017] 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[2].

In our paper [Rubinstein and Zhao 2021], inspired by [Babaioff et al. 2017]'s open question, we formulate a notion 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.

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.

[1] Throughout this note, we use feminine pronouns for the seller and masculine for the buyer.

[2] 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.

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[3] over the protocol's randomness. 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.

Consistent with essentially all prior work on this problem, 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, or a publicly observable, renewable 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 (IR). 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.

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 does not depend on U, e.g. U can be all the money in the universe (typically much smaller).

2. OUR RESULTS

2.1 Communication-Efficient IC Auction Protocols

We design IC auction protocols that are simple, surprisingly efficient, and are *exactly* revenue-optimal. For instance, 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[4].

THEOREM 2.1. *For any prior $\mathcal{D}$ of Buyer's (additive/unit-demand/combinatorial) valuations over n items bounded by maximum value U, and any IC mechanism $\mathcal{M}$,*

[3] 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.

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

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*[5]. Specifically, at the cost of an ε-fraction loss in revenue, we obtain an IC auction protocol that uses only $\text{polylog}(n)$ communication.

THEOREM 2.2. *Let* $\mathcal{D}$ *be a distribution of independent unit-demand valuations over* n *items bounded by maximum value* U*. Then, for any constant* $\varepsilon > 0$*, there is a* $(1-\varepsilon)$*-approximately revenue-optimal IC auction protocol using* $\text{polylog}(n)$ *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 [Kothari et al. 2019].

Remark 2.3. For simplicity of presentation we focus on the expected communication complexity. Here we briefly remark that 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). Using trivial batching, one can further compress the number of rounds: at the cost of a constant factor increase in the communication complexity, our 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).

2.2 Communication Complexity Lower Bounds

We show that beyond the (important) special case covered by Theorem 2.2, the communication complexity of our protocols is almost the best possible, in the following strong sense:

THEOREM 2.4. *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*[6]*:*

—$\Omega(n)$ *communication for unit-demand valuations*[7]*;*

—$2^{\Omega(n^{1/3})}$ *communication for gross substitutes valuations;*

—$2^{\Omega(n)}$ *for XOS valuations.*

Furthermore, any incentive compatible auction protocol that obtains more than (80%/91%) *of the optimal revenue must use at least:*

[5] A valuation $v : \{0,1\}^n \to \mathbb{R}_{\geq 0}$ is unit-demand if for all $S \subseteq [n]$, $v(S) = \max_{i \in S} v(\{i\})$. We say a prior distribution of unit-demand valuations has independent items if for each $i \in [n]$, $v(\{i\})$ is sampled independently from an arbitrary distribution supported on $\mathbb{R}_{\geq 0}$.

[6] To be more precise, in general, an auction protocol takes a prior distribution as input and then specifies how the buyer and the seller communicate with each other given their valuations, and our theorem states that for each specified valuation class, there exists a prior distribution such that any constant-approximate auction protocol requires certain amount of randomized communication.

[7] *additive, unit-demand* $\subset$ *gross-substitutes* $\subset$ *submodular* $\subset$ *XOS* $\subset$ *subadditive..*

—$2^{\Omega(n)}$ *communication for (XOS/submodular) 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: [Rubinstein and Weinberg 2018] showed that a constant fraction of revenue can be guaranteed by simple mechanisms; this constant has been improved in followup works [Cai et al. 2016; Chawla and Miller 2016; Cai and Zhao 2017], but no non-trivial upper bound on the best approximation factor were known. 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 submodular 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 2.2.

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

Open Question 2.5. *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 2.2) does not guarantee exact revenue optimality.

2.3 Separating the Complexity of Implementing and Incentivizing

Our results also have implications for a question of Fadel and Segal [Fadel and Segal 2009]. 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 [Fadel and Segal 2009, Open Question 3]. The source of hardness of our lower bounds is inherently different from the instances in [Fadel and Segal 2009]: 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[8] 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 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 [Fadel and Segal 2009]'s open question on the negative[9]. We remark that by [Fadel and Segal 2009, Corollary 3], this exponential separation is tight.

Corollary 2.6. *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 2.7 *Separations for deterministic vs randomized protocols.*
Formally, [Fadel and Segal 2009] phrase their open question for deterministic protocols. To view Corollary 2.6 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).

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 appendix of [Rubinstein and Zhao 2021], we show a *non-IC* auction protocol 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.

3. TECHNICAL HIGHLIGHTS

In this section, we give an overview of some highlights of our techniques for establishing our efficient protocols and nearly tight lower bounds.

3.1 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 2.1 and 2.2). 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

[8]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.
[9]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.

using unbiased coin tosses[10] 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 $p > \tau$.

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[11] $\tau \in [0, 1]$ 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.

3.2 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. [Papadimitriou et al. 2008; Dobzinski 2016; Immorlica et al. 2018; Göös and Rubinstein 2018]). However, our non-IC auction protocol in appendix of [Rubinstein and Zhao 2021] 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 alloca-

[10] 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 [Babaioff et al. 2017].
[11] τ can be drawn on the fly so the expected number of random bits is also bounded.

tions conditioned on sending "0" or "1"; this means that it can only partition the buyer's types into halfspaces in valuation space. In contrast, an non-IC protocol can make an arbitrary partition of the valuation space into two disjoint sets. 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.

REFERENCES

Babaioff, M., Gonczarowski, Y. A., and Nisan, N. 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.

Briest, P., Chawla, S., Kleinberg, R., and Weinberg, S. M. 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.

Cai, Y., Devanur, N., and Weinberg, S. M. 2016. A duality based unified approach to bayesian mechanism design. In *Proceedings of the 48th ACM Conference on Theory of Computation(STOC)*.

Cai, Y. and Zhao, M. 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.

Chawla, S. and Miller, J. B. 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.

Chen, X., Diakonikolas, I., Orfanou, A., Paparas, D., Sun, X., and Yannakakis, M. 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.

Chen, X., Diakonikolas, I., Paparas, D., Sun, X., and Yannakakis, M. 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.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2014. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2017. Strong duality for a multiple-good monopolist. *Econometrica 85,* 3, 735–767.

Dobzinski, S. 2016. Computational efficiency requires simple taxation. In *FOCS*.

Fadel, R. and Segal, I. 2009. The communication cost of selfishness. *J. Econ. Theory 144,* 5, 1895–1920.

Göös, M. and Rubinstein, A. 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.

Hart, S. and Nisan, N. 2019. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *J. Economic Theory 183*, 991–1029.

Hart, S. and Reny, P. J. 2015. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics 10,* 3, 893–922.

Immorlica, N., Lucier, B., Mao, J., Syrgkanis, V., and Tzamos, C. 2018. Combinatorial assortment optimization. In *Web and Internet Economics - 14th International Conference, WINE 2018, Oxford, UK, December 15-17, 2018, Proceedings*, G. Christodoulou and T. Harks, Eds. Lecture Notes in Computer Science, vol. 11316. Springer, 218–231.

Kothari, P., Mohan, D., Schvartzman, A., Singla, S., and Weinberg, S. M. 2019. Approximation schemes for a buyer with independent items via symmetries. In *the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS).*

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

Papadimitriou, C. H., Schapira, M., and Singer, Y. 2008. On the hardness of being truthful. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS).*

Pavlov, G. 2011. Optimal mechanism for selling two goods. *The B.E. Journal of Theoretical Economics 11,* 3.

Rubinstein, A. and Weinberg, S. M. 2018. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Trans. Economics and Comput. 6,* 3-4, 19:1–19:25.

Rubinstein, A. and Zhao, J. 2021. The randomized communication complexity of randomized auctions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 882–895.

Thanassoulis, J. 2004. Haggling over substitutes. *Journal of Economic Theory 117*, 217–245.