

Optimal Multi-Dimensional Mechanisms are not Locally-Implementable

S. MATTHEW WEINBERG*, Princeton University, USA

ZIXIN ZHOU†, Stanford University, USA

We introduce *locality*: a new property of *multi-bidder auctions* that formally separates the simplicity of optimal single-dimensional multi-bidder auctions from the complexity of optimal multi-dimensional multi-bidder auctions. Specifically, consider the revenue-optimal, Bayesian Incentive Compatible auction for buyers with valuations drawn from $\tilde{D} := \times_i D_i$, where each distribution has support-size n . This auction takes as input a valuation profile \tilde{v} and produces as output an allocation of the items and prices to charge, $\text{Opt}_{\tilde{D}}(\tilde{v})$. When each D_i is single-dimensional, this mapping is *locally-implementable*: defining each input v_i requires $\Theta(\log n)$ bits, and $\text{Opt}_{\tilde{D}}(\tilde{v})$ can be fully determined using just $\Theta(\log n)$ bits from each D_i . This follows immediately from Myerson's virtual value theory [36].

Our main result establishes that optimal multi-dimensional mechanisms are *not locally-implementable*: in order to determine the output $\text{Opt}_{\tilde{D}}(\tilde{v})$ on one particular input \tilde{v} , one still needs to know (essentially) the entire distribution \tilde{D} . Formally, $\Omega(n)$ bits from each D_i is necessary: (essentially) enough to fully describe D_i , and exponentially more than the $\Theta(\log n)$ needed to define the input v_i . We show that this phenomenon already occurs with just two bidders, even when one bidder is single-dimensional, and even when the other bidder is barely multi-dimensional. More specifically, the multi-dimensional bidder is "inter-dimensional" from the FedEx setting with just two days [28].

Our techniques are fairly robust: we additionally establish that optimal mechanisms for single-dimensional buyers with budget constraints are not locally-implementable. This again occurs even with just two bidders, even when one has no budget constraint, and even when the other's budget is public.

CCS Concepts: • Theory of computation → Algorithmic mechanism design; Communication complexity.

Additional Key Words and Phrases: mechanism design, revenue maximization, communication complexity, Lagrangian duality

ACM Reference Format:

S. Matthew Weinberg and Zixin Zhou. 2022. Optimal Multi-Dimensional Mechanisms are not Locally-Implementable. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*, July 11–15, 2022, Boulder, CO, USA. ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/3490486.3538334>

1 INTRODUCTION

Consider the problem of selling multiple items to multiple bidders, where each bidder's valuation function (for the items) is drawn independently from a distribution known to the seller. The seller desires a truthful auction (formally, Bayesian Incentive Compatible. See Section 2) maximizing her expected revenue.

*smweinberg@princeton.edu. Supported by NSF CCF-1955205.

†jackzhou@stanford.edu. Supported by a Stanford SoE Fellowship.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Since Myerson's seminal work, it is well-established that revenue-optimal *single-dimensional* auctions are exceptionally simple, and satisfy many desirable properties. For example, revenue-optimal *single-bidder* single-dimensional auctions offer the bidder a take-it-or-leave-it price (the bidder can pay the price and get the item, or not pay and get nothing). Optimal single-bidder single-dimensional auctions are therefore deterministic, computable in poly-time, revenue-monotone,¹ and have menu-complexity one.² In contrast, optimal *single-bidder* multi-dimensional mechanisms (where, e.g., two distinct items are for sale) require randomization [39, 46], are computationally hard to find [15, 16, 20], are non-monotone [31, 44],³ and have unbounded menu complexity [6, 7, 21, 30, 35]. This vast (and still growing) line of works clearly establishes that optimal *single-bidder* multi-dimensional mechanisms are extremely complex when compared to their single-dimensional counterparts.

In the *multi-bidder* setting, however, the story is less written. Of course, optimal multi-bidder multi-dimensional mechanisms inherit all the complexities of optimal single-bidder multi-dimensional mechanisms. Still, it remains largely unknown to what extent *all* the complexities of optimal multi-bidder multi-dimensional auctions already manifest in the single-bidder setting. Indeed, in some special cases where the optimal single-bidder multi-dimensional auction is tractable, the optimal multi-bidder multi-dimensional auction is tractable as well [9]. Some further multi-dimensional special cases even admit formal multi-to-single-bidder reductions [1–3]. In this direction, our work identifies a novel complexity of optimal multi-bidder multi-dimensional mechanisms *driven by the multi-bidder aspect*. For example, our quantitative measure identifies complexity in broad classes of two-bidder multi-dimensional settings, even though the single-bidder problem for every instance in these classes is quite simple.

Locally-Implementable Mechanisms. Consider the following thought experiment: you run k -bidder auctions, and your Bayesian prior is that Bidder i 's valuation function is drawn from D_i . When each D_i remains permanently fixed, it makes sense to hard-code the revenue-optimal auction for $\times_i D_i$, and plug in each new valuation profile \vec{v} as input. But you are continuously gathering data on bidders' values to refine your beliefs (in fact, every additional auction executed itself refines your beliefs for future auctions). So while just a single \vec{v} is given as input, the problem you aim to solve is parameterized by the prior \vec{D} .

DEFINITION 1 (IMPLEMENTING A REVENUE OPTIMAL AUCTION). *Given as input k valuation functions v_1, \dots, v_k , and parameterized by k distributions D_1, \dots, D_k , determine $\text{Opt}_{\vec{D}}(\vec{v})$: an allocation of items and payments charged on valuation profile (v_1, \dots, v_k) that is consistent with some revenue-optimal mechanism for $\times_i D_i$.*

At the heart of our paper is the following (for now, informally-posed) question: *How much do you really need to know about each D_i in order to compute $\text{Opt}_{\vec{D}}(\vec{v})$ for just one particular \vec{v} ?*

When each D_i is single-dimensional, not much is needed, and this follows immediately from Myerson's theory of (ironed) virtual values [36]. Indeed, if each D_i is supported on n valuations, and each valuation has an integer value between 0 and $\text{poly}(n)$ for each outcome, and the probability of each valuation is an integer multiple of $1/\text{poly}(n)$, then $O(\log n)$ bits from each D_i suffice to compute $\text{Opt}_{\vec{D}}(\vec{v})$. Inspired by the concept of locally-decodable codes (see survey [48]), we term

¹Specifically, if D stochastically dominates D' , then the optimal revenue for D exceeds that of D' .

²The menu-complexity of a single-bidder auction is the number of distinct non-trivial allocations they might receive.

³Specifically, there exist distributions over additive valuations for two items D and D' which can be coupled so that $v \sim D$ values all sets of items more than $v' \sim D'$, yet the optimal revenue for D' is infinite and the optimal revenue for D is 1!

this property *locally-implementable*: to compute $\text{Opt}_{\vec{D}}(\vec{v})$, barely more bits are needed from each D_i than the bits needed to state v_i itself.⁴

To quickly see this (see Appendix A for a more detailed sketch), recall that Myerson's seminal work defines a (ironed) virtual valuation function $\bar{\varphi}_i^{D_i}(\cdot)$ (which depends only on D_i and not D_{-i}) such that the revenue-optimal auction gives the item to Bidder $i^* := \arg \max_i \{\bar{\varphi}_i^{D_i}(v_i)\}$ and charges them price $(\bar{\varphi}_{i^*}^{D_{i^*}})^{-1}(\max\{0, \max_{i \neq i^*} \{\bar{\varphi}_i^{D_i}(v_i)\}\})$ (if $\bar{\varphi}_{i^*}^{D_{i^*}}(v_{i^*}) \geq 0$, otherwise no one wins the item). In particular, *the winner can be determined just by knowing $\bar{\varphi}_i^{D_i}(v_i)$ for all i* , and the payment charged can be further determined *with one additional query to $\bar{\varphi}_{i^*}^{D_{i^*}}(\cdot)$* . Moreover, each $\bar{\varphi}_i^{D_i}(v)$ is the ratio of two integers of size at most $\text{poly}(n)$ (subject to the conditions at the start of this paragraph), so *just $O(\log n)$ bits from each D_i suffice to compute $\text{Opt}_{\vec{D}}(\vec{v})$, while $\Omega(n)$ bits are necessary to fully specify each D_i* .

Our main result shows that optimal multi-dimensional mechanisms are *not locally-implementable*, and in fact are as far from locally-implementable as possible: $\Omega(n)$ bits from each D_i are needed to determine $\text{Opt}_{\vec{D}}(\vec{v})$ — nearly as many bits as needed to fully specify D_i , and exponentially more than the $\Theta(\log n)$ bits needed to specify v_i . We further show that this already holds in essentially the simplest multi-dimensional setting: there are just two bidders, one of whom is single-dimensional, and another who is “inter-dimensional” according to the FedEx problem [28]. Specifically, there are two options for the item (call them one-day and two-day shipping). The single-dimensional bidder always has value 0 for two-day shipping. The multi-dimensional bidder either has the same value for both options, or has value 0 for two-day shipping. Such bidders are a (very) special case of unit-demand bidders.⁵ They are also a special case of buyers with a private budget constraint [24], and single-minded buyers [22]. Formally, we use the lens of communication complexity for the following problem to establish our main result.

DEFINITION 2. *Select-Outcome Problem* is a communication problem between Alice and Bob. Alice is given as input D_1 , and a valuation v_1 in its support. Bob is given D_2 , and a valuation v_2 in its support. When the input is of size n , each distribution has support-size at most n , all valuations in the support have integer values $\leq n^3$ for all outcomes, and all probabilities are an integer multiple of $\frac{1}{b}$, for some integer $b \leq n^8$.⁶

A solution to Select-Outcome Problem outputs an outcome x such that some revenue-optimal auction for $D_1 \times D_2$, on valuation profile (v_1, v_2) , selects outcome x with non-zero probability.⁷

In this language, the previous paragraphs state that Select-Outcome Problem can be solved in deterministic communication complexity $O(\log n)$ when both bidders are single-dimensional (and moreover, correct prices can be found in communication $O(\log n)$ as well). Formally, our main result is that the communication complexity is exponentially higher in multi-dimensional settings.

THEOREM 1.1 (MAIN RESULT). *Even when D_2 is single-dimensional, and D_1 is a FedEx bidder with two options, the communication complexity of Select-Outcome Problem is $\Omega(n)$. This holds for deterministic protocols, as well as randomized protocols which succeed with probability $\geq 2/3$.*

⁴To draw the (very high-level) conceptual connection to locally-decodable codes, think of the n -bit codeword C as parameterizing the decoding algorithm, which is given an index i as input (and the desired output is m_i , the i^{th} bit of the original message). Locality refers to the fact that m_i can be determined by querying just $o(n)$ bits of C . Similarly, locality in our context refers to the fact that $\text{Opt}_{\vec{D}}(\vec{v})$ can be determined using just $O(\log n)$ bits from each D_i .

⁵A valuation is unit-demand if its valuation for a set of items S is $v(S) := \max_{i \in S} v(\{i\})$.

⁶The particular choice of n^3 and n^8 are immaterial, and can be any sufficiently large (fixed) polynomials in n .

⁷For example, in a single-item auction with two bidders there are three outcomes: give the item to bidder one, bidder two, or no one. Note that many outcomes may be correct, both due to multiplicity of optimal auctions, and due to optimal auctions being randomized. Note also that the outcome selected may be the “null” outcome to keep all items with the seller.

We quickly motivate our precise choices in defining Select-Outcome Problem. Because all input numbers are integers $\leq \text{poly}(n)$, Theorem 1.1 must follow because any solution to Select-Outcome Problem requires many bits from each distribution (and not because $\Omega(n)$ bits are required just to do arithmetic on the input). Allowing the solution to be consistent with any allocation output with non-zero probability in any optimal auction ensures that hardness follows for any reasonable alternative definition as well (and not because of some technicality associated with multiplicity or randomization of optimal auctions).

We further note that the complexity uncovered by Theorem 1.1 arises *only* in the multi-bidder setting, as the single-bidder problems for both D_1 and D_2 are quite simple. Indeed, D_2 is a single-dimensional bidder, so the optimal single-bidder auction for D_2 is just a take-it-or-leave-it price. The optimal single-bidder auction for a FedEx bidder with two options is only slightly more complex: it has menu complexity at most two,⁸ is computationally tractable, etc. [28]. Additionally, for every possible instantiation of D_1 or D_2 used in our construction, the revenue-optimal single-bidder auction simply sets a take-it-or-leave-it price of $n^2 + 1$ (see Proposition 4.7). Put another way, every instantiation of D_1 and D_2 in our construction admits the simplest possible single-bidder solution. Yet, the optimal multi-bidder auction for both together is not locally-implementable (and recall that this phenomenon cannot occur with two single-dimensional bidders).

Finally, we emphasize that locality is an intrinsic measure of complexity for multi-bidder auctions. In this sense, the communication complexity of Select-Outcome Problem serves not as a problem to be solved in practice, but rather as a quantitative lens to view the extent to which the output of an optimal multi-bidder auction for one particular input \vec{v} depends on the underlying prior \vec{D} .

1.1 Extension: Budget-Constrained Bidders

In addition to our main result, we also consider bidders who are barely beyond the classic single-dimensional setting in a different direction: they have a budget constraint. Specifically, bidders have a value v and a budget B . If they receive the item and are charged $p \leq B$, their utility is $v - p$ as usual. If they are charged $p > B$, their utility is $-\infty$. That is, the bidder's utility is *not* quasi-linear.⁹

If v and B are both private information to the bidder, this is an inter-dimensional setting [24], and Theorem 1.1 already establishes that optimal mechanisms are not locally-implementable. If instead the budget is *public* (known to the designer), then this is still a single-dimensional setting (because the bidder's private information is just a single value), but it is *non-linear* (because the buyer is not quasi-linear). Again, this is essentially the simplest non-linear setting (and perhaps the most well-studied within the TCS literature, arguably by a significant margin): the buyer is still single-dimensional, and her utility with respect to price is piecewise-linear with two segments. We also show that optimal mechanisms for single-dimensional budget-constrained buyers are not locally-implementable.

THEOREM 1.2. *Even when D_2 is single-dimensional and quasi-linear, and D_1 is single-dimensional with a public budget constraint, the communication complexity of Select-Outcome Problem is $\Omega(n)$. This holds for deterministic protocols, as well as randomized protocols which succeed with probability $\geq 2/3$.*

We again note that this complexity arises in the multi-bidder setting, despite the fact that each single-bidder problem is quite simple. Again, D_2 is single-dimensional and quasi-linear, so the optimal single-bidder mechanism is just a take-it-or-leave-it price. The optimal single-bidder

⁸Specifically, it offers one-day shipping at a take-it-or-leave-it price. It may additionally offer one (perhaps randomized) option to receive two-day shipping at a discount.

⁹A bidder is quasi-linear if their utility for receiving the item is $v - p$ always. Throughout the paper, bidders will always be assumed to be quasi-linear unless otherwise specified.

mechanism for a single-dimensional buyer with a public budget has menu complexity at most two,¹⁰ is computationally tractable, etc. [13]. Additionally, for every possible instantiation of D_1 or D_2 used in our construction, the revenue-optimal single-bidder auction again sets a take-it-or-leave-it price of $n^2 + 1$ (see Proposition E.6 in the full version). We additionally emphasize that our constructions for Theorems 1.1 and Theorem 1.2 overlap significantly (see Section 3), highlighting the robustness of our technical contributions.

1.2 Additional Implications for Multi-Dimensional Virtual Values

We now provide an additional lens through which to view the implications of our main result. Specifically, several works provide some form of “multi-dimensional virtual values” [2, 9, 11, 12, 29, 42]. Their precise uses and derivations differ (see Section 1.4 for further detail), but they all share a theme of connecting truthful revenue maximization to algorithmic virtual welfare maximization. For example, [11] derives multi-dimensional virtual values through Lagrangian duality, and proves that for all instances D , there exists a virtual valuation function $\Phi_i^D(\cdot)$ for each bidder i such that for all valuation profiles \vec{v} , every revenue-optimal auction must select an outcome x maximizing $\sum_i(\Phi_i^D(v_i))(x)$.

In particular, this view shows a correspondence between Lagrangian multipliers/dual variables in a natural Linear Programming formulation (overviewed in Section 2) and these virtual valuation functions. This Linear Program has variables, constraints, and dual variables for each bidder.

One additional beautiful aspect of Myerson’s virtual value theory is the following: $\bar{\varphi}_i^{D_i}(\cdot)$ depends *only on D_i and not at all on D_{-i}* . In the language of LP duality, this implies the following remarkable property: in the LP formulation, *the optimal dual variables for Bidder i also depend only on D_i and not at all on D_{-i} !* The fact that the optimal dual variables can be computed separately for each bidder is remarkable because the optimal primal solution certainly cannot (Bidder i ’s allocation/price variables in the optimal auction certainly depend on D_{-i}). To emphasize the implications of this remarkable property: consider writing, separately for each i , the LP formulation to optimally sell a single item to a single bidder whose value is drawn from D_i . To solve the multi-bidder LP formulation for k bidders whose values are drawn from $\times_i D_i$, one could try naïvely stapling the k optimal primals to these single-bidder LPs together. There is no reason to expect this to succeed, and indeed it fails (in fact, it will generally fail to even produce a feasible primal). On the other hand, one could alternatively try naïvely stapling together the k optimal dual variables together, and hope that this produces the optimal dual variables for the k -bidder LP formulation. Somewhat miraculously, *this latter process succeeds in any single-dimensional setting* (with quasi-linear bidders).

For multi-dimensional bidders, the [11] framework still establishes the existence of $\Phi_i^D(\cdot)$, but not necessarily that $\Phi_i^D(\cdot)$ is agnostic to D_{-i} as in the single-dimensional case. Or in the language of LPs, the optimal dual variables for Bidder i may a priori depend on the entire prior, rather than just D_i . Viewed through this lens, [2, 29] discover restricted multi-dimensional settings where optimal duals retain this remarkable “bidder-separable” property. Through the same lens, Theorem 1.1 establishes that this bidder-separable property does not generally hold in multi-dimensional settings (even with just two bidders from the two-day FedEx problem). Theorem 1.2 rules out the bidder-separable property for single-dimensional non-linear settings as well. In addition, Theorems 5.2 and 5.5 further give concrete examples of how optimal dual variables for Bidder 1 can be quite sensitive to tiny changes in D_2 .

¹⁰Specifically, it offers the option to receive the item at some price p . If $p = B$, it may additionally offer the option to receive the item with probability $q < 1$ at price $r < B$.

1.3 Very Brief Technical Overview

The proof of Theorem 1.1 follows by a reduction from Disjointness (formally defined in Section 2). Our reduction makes heavy use of notation and concepts from prior work, so we defer an outline of the approach to Section 3 once appropriate language is built up. We provide here a brief highlight of the main challenge: we have just spent several paragraphs in Section 1 describing all the intractable properties that revenue-optimal auctions possess. To complete a reduction, we not only need to derive the optimal auction for a single instance, but for an entire class of instances. Moreover, this class must contain sufficiently many “intractable instances” in order to embed Disjointness. Indeed, reductions to Bayesian mechanism design are scarce, technically involved, and to-date exist only for single-bidder settings [10, 15, 16, 18, 20, 25]. In multi-bidder settings, the state-of-the-art only recently characterized optimal auctions for all instances with two additive bidders and two items where item values are drawn i.i.d. from distributions supported on $\{1, 2\}$ [47]!

Fortunately, the FedEx setting is a sweet spot which is both rich enough for optimal mechanisms to be non-locally-implementable, yet also structured enough for a tractable reduction to optimal mechanism design. We hope that the proof outline in Section 3 may serve as a roadmap for potential future reductions, recalling that any complexities established for the FedEx setting extend to the (significantly more general) multi-dimensional unit-demand setting as well.

1.4 Related Work

Complexity of Multi-Dimensional Mechanism Design. We have already discussed the thematically most-related work, which identifies formal complexity measures separating revenue-optimal single- and multi-dimensional mechanisms [6, 7, 15, 16, 20, 21, 30, 31, 35, 39, 44, 46, 47]. Among these, only [47] explicitly studies the multi-bidder setting, and establishes that while optimal single-dimensional auctions are dominant-strategy truthful,¹¹ optimal multi-dimensional auctions are not [47]. In comparison to this line of works, our paper provides a novel complexity *unique to multi-bidder settings*. Specifically, our work identifies complexity in broad classes of two-bidder settings, even though the single-bidder problem for every instance in these classes is quite simple.

Multi-Dimensional Virtual Values. Our work uses multi-dimensional virtual values in order to prove optimality of mechanisms. As previously referenced, several prior works introduce various notions of multi-dimensional virtual values [2, 9, 11, 12, 29, 42]. Some of these works consider continuous distributions, and derive multi-dimensional virtual values by explicitly choosing paths along which the incentive constraints might bind and then doing integration by parts. Others consider discrete settings, and derive multi-dimensional virtual values by drawing a connection to LP duality. Because our setting is discrete (and necessarily so, in order for communication complexity to be a meaningful measure), we adopt the language used in [11], which uses the lens of LP duality.

Within this line of works, [2, 29] also prove optimality in some multi-bidder settings, and in particular discover restricted settings where multi-dimensional virtual values are bidder-separable (termed “revenue-linear” and “MR-log-supermodular”, respectively). In their language, our main results rule out any extension to two-bidder settings generally (even when one bidder is single-dimensional, and the other is a two-day FedEx bidder or single-dimensional with a public budget). In our language, [2, 29] discover restricted multi-dimensional settings where optimal mechanisms are locally-implementable.

¹¹An auction is dominant-strategy truthful if it is in each bidder’s interest to report their true valuation no matter the other bidders’ reports. Contrast this with Bayesian Incentive Compatible (defined in Section 2).

Interdimensional Mechanism Design. [28] introduce the FedEx problem, and note that optimal single-bidder mechanisms inherit some-but-not-all of the nice properties of single-dimensional settings, along with some-but-not-all of the complexities associated with multi-dimensional settings. Our main result considers a bidder from the FedEx problem, and therefore our technical setup is similar to works such as [22–24, 28], but there is not much overlap with these works beyond Section 2.

Budget-Constrained Bidders. There is a substantial body of work involving mechanism design for budget-constrained buyers. The most related works to ours design revenue-optimal single-item auctions in Bayesian settings for buyers with a public or private budget constraint. Here, [13, 14, 24, 34] characterize the optimal single-buyer auction.¹² Several works also identify tractable structure for the optimal auction in restricted cases. For example, [34] consider the case of multiple buyers with values drawn i.i.d. from the same regular distribution and an identical public budget, and [38] consider the case that each buyer has a private budget and their value is drawn independently of their budget from an MHR distribution with decreasing density. In this context, our results establish that while optimal single-bidder mechanisms for budget-constrained buyers are quite tractable, optimal multi-bidder auctions remain intractable (without the restrictions imposed in works such as [34, 38]).

Reductions in Bayesian Mechanism Design. We have also briefly discussed reductions to optimal mechanism design, which previously exist only in single-bidder settings [10, 15, 16, 18, 20, 25]. Other styles of single-bidder reductions have been used to special cases of optimal single-bidder mechanism design (such as finding the optimal deterministic auction) [8, 17, 19]. In comparison to this line of works, our paper provides a technical contribution via the first reduction to multi-bidder Bayesian mechanism design.

Communication Complexity in Multi-Dimensional Mechanism Design. Recent work of [6] identifies a connection between the so-called *menu complexity* of single-bidder auctions and the deterministic communication required to implement it. More recent work of [45] further considers the randomized communication complexity required to implement single-bidder auctions, and in particular establishes that randomized implementations of auctions may sometimes communicate exponentially fewer bits than deterministic implementations. While this model is incomparable to ours,¹³ this context makes it significant that Theorems 1.1 and 1.2 hold for randomized communication protocols.

There is also a substantial body of work at the intersection of communication complexity and mechanism design generally, following seminal work of [37]. A parallel line of works following [27] considers the communication overhead specifically to compute payments (on top of any communication necessary to determine an outcome/allocation). Follow-up works of [5, 26, 43] show that this overhead can be quite significant, even with just two players. While their model is also incomparable to ours,¹⁴ this context makes it significant that our main results provide communication lower

¹²Specifically, it sets a single price if the budget is public and the valuation distribution is regular [34]. It has menu complexity at most two if the budget is public, no matter the valuation distribution [13]. It has menu complexity at most k if the budget is private and drawn from a distribution of support at most k , and the valuation conditioned on each possible budget satisfies a condition called “decreasing marginal revenues.” It has menu complexity at most $3 \cdot 2^{k-1} - 1$ if the budget is private and drawn from a distribution of support at most k , no matter the joint distribution of (value, budget) pairs.

¹³Specifically, these works study a single-bidder problem where the prior (and therefore the auction) is fully-known. Their goal is to implement the auction for a particular valuation v without necessarily learning v .

¹⁴For example, none of these works consider revenue-optimization at all.

bounds *just to solve Select-Outcome Problem* (rather than to solve Select-Outcome Problem and to also determine the payments).

1.5 Summary and Roadmap

We establish that optimal multi-dimensional mechanisms are not locally-implementable: *executing the auction on just a single valuation profile requires knowing (essentially) the entire distribution*. Formally, we study the communication complexity of Select-Outcome Problem. In single-dimensional settings, Select-Outcome Problem can be solved with $O(\log n)$ bits of communication, while simply stating the valuation profile also requires $\Theta(\log n)$ bits. In multi-dimensional settings, Theorem 1.1 gives a communication lower bound of $\Omega(n)$ on Select-Outcome Problem: exponentially more than the $\Theta(\log n)$ bits needed to state the input valuation profile, and nearly the $\Theta(n \log n)$ bits sufficient to fully define each D_i . In particular, recall that all D_i in our construction are especially simple from the single-bidder perspective: the optimal single-bidder auction sets a take-it-or-leave-it price of $n^2 + 1$ (Proposition 4.7). This makes clear that non-locality truly arises due to complexity of *multi-bidder* multi-dimensional auctions, and not simply due to complexity of the corresponding single-bidder problem.

Section 2 provides preliminaries. Section 3 provides a high-level overview of our approach. Section 4 provides our reduction, and Section 5 our analysis. Section 6 provides concluding thoughts. The appendix contains omitted proofs. In particular, Appendix E in the full version contains our complete analysis for the case of single-dimensional buyers with a public budget.

2 PRELIMINARIES

Below, we provide detailed preliminaries for our main result (Section 2.1) and detailed background on Lagrangian duality for Bayesian mechanism design (Section 2.2), so that we can present the key ideas behind our construction. We also formally define the setting we consider for our extension to budget-constrained bidders in Section 2.3 but defer to Appendix E in the full version full preliminaries necessary for the proofs. Section 2.4 quickly states the communication problem of Disjointness, which we use in our reductions.

2.1 The FedEx Problem

Setup and Notation. Our main result holds already when there are just two bidders and two options, which we refer to as day1 and day2. Bidders have a value and an interest. A bidder with (value, interest) pair $(v, 1)$ receives value v if they receive one-day shipping, and 0 if they receive two-day shipping. A bidder with (value, interest) pair $(v, 2)$ receives value v if they receive either one-day or two-day shipping. A bidder's *type* stores her full (value, interest) pair.

Each of the two bidders i have $2n_i + 1$ different types, we label them from t_i^0 to $t_i^{2n_i}$. There are n_i possible values among all types, which we label as v_i^j , for $j \in [n_i]$. In this labeling, t_i^{2k-1} represents the (value, interest) pair $(v_i^k, 1)$ and t_i^{2k} represents the (value, interest) pair $(v_i^k, 2)$. Finally, t_i^0 represents not participating in the auction, and has value $v_i^0 := 0$. We will alternate between referring to types as t^{2k+j-2} and (v_i^k, j) , depending on which notation is cleaner. We denote by $f_i(t_i)$ the probability that bidder i has type t_i , and we use D_i to represent the distribution of bidder i . Finally, we will also use the notation $R_i((v_i^k, j)) := \sum_{k' \geq k} f_i((v_i^{k'}, j))$.¹⁵

¹⁵Observe that $R_i(\cdot)$ is essentially a reverse CDF. Indeed, if there were only one possible interest, the definition would imply that $R_i(\cdot) := 1 - F_i(\cdot)$, where $F_i(\cdot)$ is the CDF for Bidder i .

Optimal Auctions. We have one item to ship, and can ship it to either bidder using either one- or two-day shipping.¹⁶ Note that this is a “service-constrained” environment, as defined in [2]. We seek the revenue-optimal Bayesian Incentive Compatible (BIC) auction, which asks each bidder to report their (value, interest) pair, and then decides to whom to ship the item (or to no one) and in how many days. It is observed in [28] that the revenue-optimal auction w.l.o.g. always ships the item on the reported interest (that is, it will never ship the item in two days to a bidder whose interest is day1, or vice versa).¹⁷ With this in mind, a mechanism is defined by its ex-post allocation rule X and ex-post payment rule P . Here, $X_i(t_1, t_2)$ denotes the probability that bidder i is shipped the item (matching their interest) when the reported types are t_1, t_2 , and $P_i(t_1, t_2)$ denotes the payment made by bidder i when the reported types are t_1, t_2 . Because we have one copy of the item to ship, an allocation rule is feasible iff for all t_1, t_2 , $X_1(t_1, t_2) + X_2(t_1, t_2) \leq 1$.

An auction is Bayesian Incentive Compatible (BIC) if it is in each bidder’s interest to report their true type in expectation over the types of the other bidder. More specifically, the revenue-optimal BIC auction is the solution to the following linear program. In the LP, the variables are X, P, π, p . X and P refer to the ex-post allocation/price rules, as defined above. π, p refer to the interim allocation/price rules, which satisfy the equalities in Equations (1) and (2).

$$\begin{aligned} \max_{X, P, \pi, p} \quad & \sum_i \sum_{j=1}^{2n_i} f_i(t_i^j) \cdot p_i(t_i^j) \\ \text{subject to} \quad & X_i(t_1^j, t_2^\ell) \in [0, 1] \text{ for all bidders } i \text{ and all } j, \ell. \\ & X_1(t_1^0, t_2^\ell) = P_1(t_1^0, t_2^\ell) = 0 \text{ for all } \ell \in [0, 2n_2]. \\ & X_2(t_1^k, t_2^0) = P_2(t_1^k, t_2^0) = 0 \text{ for all } k \in [0, 2n_1]. \\ & X_1(t_1^j, t_2^\ell) + X_2(t_1^j, t_2^\ell) \leq 1 \text{ for all } j \in [2n_1], \ell \in [2n_2]. \\ & \pi_i(t_i^j) = \sum_{\ell=1}^{2n_{3-i}} f_{3-i}(t_{3-i}^\ell) \cdot X_i(t_i^j; t_{3-i}^\ell) \text{ for all bidders } i \text{ and } j \in [0, 2n_i]. \quad (1) \\ & p_i(t_i^j) = \sum_{\ell=1}^{2n_{3-i}} f_{3-i}(t_{3-i}^\ell) \cdot P_i(t_i^j; t_{3-i}^\ell) \text{ for all bidders } i \text{ and } j \in [0, 2n_i]. \quad (2) \\ & \pi_i((v_i^k, j)) \cdot v_i^k - p_i((v_i^k, j)) \geq \pi_i((v_i^{k'}, j')) \cdot v_i^{k'} - p_i((v_i^{k'}, j')), \quad (3) \\ & \text{for all bidders } i, \text{ all } k, k' \in [0, n_i], \text{ and } 2 \geq j \geq j' \geq 1. \end{aligned}$$

The objective is simply the expected revenue. Constraints (1) and (2) simply confirm that the interim rules are computed correctly. Equation (3) guarantees that the mechanism is BIC. In particular, Equation (3) observes that it is only necessary to ensure that bidders don’t wish to underrepresent their interest (because overrepresenting their interest guarantees them non-positive utility).

Payment Identity. Myerson’s payment identity provides a closed-form to compute revenue-maximizing payments for a fixed (monotone) allocation rule. Observe in particular that a payment rule satisfying the payment identity exists for any (monotone) allocation rule.

¹⁶If desired, our construction can be easily modified so that the auctioneer has a copy of the item shippable on each day, and the bidders are unit-demand (or to many other settings), but we only present one to establish the desired hardness.

¹⁷To quickly see this: observe that two-day shipping an item to a bidder with interest day1 gives them zero value, so the item may as well not be shipped. A bidder with day2 interest is indifferent between one-day and two-day shipping, so giving them two-day instead of one-day shipping does not affect their utility and makes other types of that bidder less interested in misreporting this type.

DEFINITION 3 (MONOTONE, PAYMENT IDENTITY). An interim allocation rule π is monotone if for both players i and days j : $\pi_i((\cdot, j))$ is monotone non-decreasing. p satisfies the payment identity for π if for both players i , days j , and all k , we have: $p_i((v_i^k, j)) = \sum_{\ell=1}^k v_i^\ell \cdot (\pi_i((v_i^\ell, j)) - \pi_i((v_i^{\ell-1}, j)))$.

2.2 Lagrangian Duality

The purpose of this section is to build up the necessary notation/concepts in order to state Definition 6 and Theorem 2.1 at the end. Theorem 2.1 provides an approach to claim that a mechanism is or isn't optimal for a given instance. This approach uses Lagrangian duality, and specifically the language adopted in [11]. More specifically, we will put Lagrangian multipliers on the BIC constraints in the following manner, which creates a Lagrangian relaxation:

- (i) For constraints of the form: $\pi_i((v_i^k, 2)) \cdot v_i^k - p_i((v_i^k, 2)) \geq \pi_i((v_i^k, 1)) \cdot v_i^k - p_i((v_i^k, 1))$, we use a Lagrangian multiplier of $\alpha_i(k)$ (for all bidders i and $k \in [1, n_i]$).
- (ii) For constraints of the form: $\pi_i((v_i^k, j)) \cdot v_i^k - p_i((v_i^k, j)) \geq \pi_i((v_i^{k-1}, j)) \cdot v_i^k - p_i((v_i^{k-1}, j))$, we use a Lagrangian multiplier of $\lambda_i^j(k)$ (for all bidders i , items j , and $k \in [1, n_i]$).
- (iii) For all remaining BIC constraints, we use a Lagrangian multiplier of 0.
- (iv) To emphasize: for all other constraints (i.e. all the constraints which are unrelated to BIC), we don't use Lagrangian multipliers, and keep them as constraints.

Constraints in (i) guarantee that the bidder will not misreport its interest, and constraints in (ii) guarantee that the bidder will not underreport their value by the minimal amount possible. Definitions 4 and 5, and Theorem 2.1 below specialize the [11] framework to our setting. We refer the reader to [11] for further details surrounding their framework, but give brief intuition for each definition throughout. Recall that every choice of Lagrangian multipliers (α, λ) induces a Lagrangian relaxation with objective function:

$$\begin{aligned} \mathcal{L}(\alpha, \lambda) := & \sum_i \sum_{j=1}^{2n_i} f_i(t_i^j) \cdot p_i(t_i^j) + \sum_i \sum_{k=1}^{n_i} \alpha(k) \cdot \left(\pi_i((v_i^k, 2)) \cdot v_i^k - p_i((v_i^k, 2)) - \pi_i((v_i^k, 1)) \cdot v_i^k + p_i((v_i^k, 1)) \right) \\ & + \sum_i \sum_{k=1}^{n_i} \sum_{j=1}^2 \lambda_i^j(k) \cdot \left(\pi_i((v_i^k, j)) \cdot v_i^k - p_i((v_i^k, j)) - \pi_i((v_i^{k-1}, j)) \cdot v_i^k + p_i((v_i^{k-1}, j)) \right). \end{aligned}$$

The constraints are the same as in the initial LP, except removing the BIC constraints. The first concept in the [11] framework is that of a flow:

DEFINITION 4 (FLOW). A set of Lagrangian multipliers form a flow if the following hold for all i :

- $f_i(t_i^{2k-1}) + \lambda_i^1(k+1) + \alpha_i(k) = \lambda_i^1(k)$, for all $k \in [1, n_i - 1]$.
- $f_i(t_i^{2n_i-1}) + \alpha_i(n_i) = \lambda_i^1(n_i)$.
- $f_i(t_i^{2k}) + \lambda_i^2(k+1) = \alpha_i(k) + \lambda_i^2(k)$, for all $k \in [1, n_i - 1]$.
- $f_i(t_i^{2n_i}) = \alpha_i(n_i) + \lambda_i^2(n_i)$.

Intuitively, Definition 4 captures the following. In the relaxation, there are no constraints on the payment variables *at all*, so the relaxation is unbounded if any payment variable has a non-zero coefficient in $\mathcal{L}(\alpha, \lambda)$. (α, λ) form a flow if and only if all payment variables have a coefficient of zero in $\mathcal{L}(\alpha, \lambda)$.

DEFINITION 5 (VIRTUAL VALUES). For a given set of Lagrangian multipliers α, λ , define:¹⁸

$$\Phi_i^{\alpha, \lambda}((v_i^k, j)) := v_i^k - \frac{(v_i^{k+1} - v_i^k) \cdot \lambda_i^j(k+1)}{f_i((v_i^k, j))}.$$

¹⁸For simplicity of notation, denote by $\lambda_i^j(n_i + 1) := 0$, $v_i^{n_i+1} := v_i^{n_i}$.

OBSERVATION 1 ([11]). *For any (α, λ) which form a flow: $\mathcal{L}(\alpha, \lambda) = \sum_i \sum_{j=1}^{2n_i} f_i(t_i^j) \cdot \pi_i(t_i^j) \cdot \Phi_i^{\alpha, \lambda}(t_i^j)$.*

Intuitively, Observation 1 follows from algebraic manipulation, and Definition 5 is made for the sole purpose of yielding Observation 1, as it suggests that any optimal allocation rule for a particular Lagrangian relaxation should award the item to the bidder with highest virtual value (according to Definition 5).

DEFINITION 6 (WITNESS OPTIMALITY). *Let (α, λ) be a flow and (X, P) be a BIC auction such that:*

- *p satisfies the payment identity for π .*
- *For all $k: \alpha_i(k) > 0 \Rightarrow \pi_i((v_i^k, 2)) \cdot v_i^k - p_i((v_i^k, 2)) = \pi_i((v_i^k, 1)) \cdot v_i^k - p_i((v_i^k, 1))$.*
- *On all $(t_1^k, t_2^{k'})$, X awards the item to a bidder with highest non-negative virtual value.*

*Then we say that (α, λ) witnesses optimality for (X, P) , and (X, P) witnesses optimality for (α, λ) .*¹⁹

THEOREM 2.1 ([11]). *Let (α, λ) witness optimality for (X, P) . Then (X, P) is a revenue-optimal BIC auction. Moreover, all revenue-optimal auctions witness optimality for (α, λ) .*

Intuitively, (α, λ) and (X, P) witness optimality if (X, P) is optimal for the Lagrangian relaxation induced by (α, λ) (bullet three), and also (α, λ) and (X, P) satisfy complementary slackness (bullets one/two).

2.3 Public Budget Constraints

Setup and Notation. Our main extension considers bidders with a (value, budget) pair. A bidder with value v and budget B enjoys utility $v - p$ if they receive the item and pay $p \leq B$, and utility $-\infty$ if they pay price $p > B$. Each of the two bidders i have $n_i + 1$ different types. Because the budget is public, their type is fully specified by a value, which we label $v_i^0, \dots, v_i^{n_i}$. Again, v_i^0 refers to non-participation in the auction.

Optimal Auctions. We have one item for sale, and can give it to either bidder. We again seek the revenue-optimal BIC auction. Because the bidders are not quasi-linear, we must also specify that we seek an *ex-post* individually rational auction.²⁰ That is, even after learning the bid of the other player, and learning the outcome of all random coins of the mechanism, each bidder has non-negative utility. Appendix E in the full version provides a linear program for this setting, and more detailed preliminaries similar to Section 2.1.

2.4 Disjointness

Our communication complexity lower bound provides a reduction from Disjointness. In Disjointness, Alice is given $\vec{x} \in \{0, 1\}^n$, Bob is given $\vec{y} \in \{0, 1\}^n$, and their goal is to determine whether there exists an i such that $x_i = y_i = 1$. It is known that any deterministic communication protocol resolving Disjointness requires communication at least n , and any randomized protocol resolving Disjointness correctly with probability at least $2/3$ requires communication $\Omega(n)$ [32, 33, 41].

3 PROOF OVERVIEW

Our proof of Theorems 1.1 and 1.2 both follow the same outline below. All steps below apply to both proofs, although the referenced technical sections are for Theorem 1.1 (where significantly more detail is provided).

¹⁹[29] note that some (α, λ) cannot be optimal for any instance, because they cannot witness optimality for any incentive compatible (X, P) . However, note that the optimal $((\alpha, \lambda), (X, P))$ witness optimality for each other, by strong Lagrangian duality and complementary slackness.

²⁰When bidders are quasi-linear, any interim individually rational auction can be made *ex-post* individually rational with a simple reduction. This reduction fails when bidders have budget constraints.

(i) Section 4 defines our reduction from Disjointness. Specifically, we define a mapping from Alice's input x to a distribution D_1 , and from Bob's input y to a distribution D_2 . Section 4 states several properties of our reduction that will be used in later proofs. Here is an informal overview of some key properties.

- For the FedEx setting, all values lie in $\{n^2 + 1, n^2 + 2, \dots, n^2 + n + 2\}$. In the public budget setting, Bidder One's values lie in $\{n^2 + 1, n^2 + 2, \dots, n^2 + n + 2\}$, and Bidder Two's values lie in $\{n^2 + 1, n^2 + 2, \dots, n^2 + n + 1, n^2 + n + 1.9\}$. In all settings, both bidders' distributions have support-size $n + 2$. In the FedEx setting, the interesting index is 1 (the lowest value in the support). In the public budget setting, the interesting index is $n + 2$ (the highest value in the support). For the rest of this outline, refer to the interesting index as j .
- All distributions used in our constructions are nearly-uniform. Therefore, the optimal single-bidder auction for any distribution in our constructions is quite simple, and sets a price of $n^2 + 1$.
- Depending on the input to Disjointness, the distribution is perturbed slightly at all values.

(ii) Section 5.1 analyzes the canonical "Myerson flow" (see Section 5.1 for definition) for our construction, which yields virtual values equal to Myersonian virtual values. For the rest of this outline, refer to this flow as (α_1, λ_1) for both settings.

- We consider the allocation rule that awards the item to the bidder with maximum $\Phi_i^{\alpha_1, \lambda_1}(t_i)$ (and charges prices according to the payment identity). Refer to this auction (X_1, P_1) .
- *If and only if* $\text{Disjointness}(x, y) = \text{Yes}$, (X_1, P_1) happens to be a second price auction, breaking ties for Bidder One (Definition 11). We then show that (X_1, P_1) witnesses optimality for (α_1, λ_1) *if and only if* $\text{Disjointness}(x, y) = \text{Yes}$.
- This means that when $\text{Disjointness}(x, y) = \text{Yes}$, we've now found the optimal dual $((\alpha_1, \lambda_1))$ and optimal auction (Definition 11).
- We also show that $\Phi_1^{\alpha_1, \lambda_1}(t_1^j) > \Phi_2^{\alpha_1, \lambda_1}(t_2^j) > 0$ in both settings.
- Now, by Theorem 2.1, this leads to our first key conclusion: *when $\text{Disjointness}(x, y) = \text{Yes}$, every optimal auction must have $X_1(t_1^j, t_2^j) = 1$.*

(iii) Section 5.2 modifies the canonical Myerson flow, for instances where $\text{Disjointness}(x, y) = \text{No}$.

- For the FedEx setting, (X_1, P_1) is not BIC: when Buyer One has (value, interest) pair $(n^2 + n + 2, 2)$, she would rather misreport $(n^2 + n + 2, 1)$. For the public budget setting, (X_1, P_1) is not budget-respecting: Buyer One with value $n^2 + n + 2$ would have to pay more than her budget.
- We increase the Lagrangian multiplier for the violated constraint from the previous bullet, and adjust others in order to preserve flow-conservation. This step is the most intricate, and requires a very precise setting of each multiplier. For the rest of this outline, call this flow (α_2, λ_2) .
- We next find an allocation rule that witnesses optimality for (α_2, λ_2) , (X_2, P_2) . (X_2, P_2) is also a second-price auction, but ties must be broken in a precise (randomized) manner (Definition 14). We show that (X_2, P_2) witnesses optimality for (α_2, λ_2) *if and only if* $\text{Disjointness}(x, y) = \text{No}$. This step is also delicate, as we must simultaneously satisfy several constraints related to Theorem 2.1.
- We also show that $0 < \Phi_1^{\alpha_2, \lambda_2}(t_1^j) < \Phi_2^{\alpha_2, \lambda_2}(t_2^j)$ in both settings. Notice that the inequalities are flipped compared to (α_1, λ_1) .
- Now, by Theorem 2.1, this leads to our second key conclusion: *when $\text{Disjointness}(x, y) = \text{No}$, every optimal auction must have $X_1(t_1^j, t_2^j) = 0$.*

(iv) To conclude, (ii) and (iii) together establish that when $\text{Disjointness}(x, y) = \text{Yes}$, *every optimal auction awards bidder 1 the item with probability 1 on input (t_1^j, t_2^j)* . On the other hand, when

$\text{Disjointness}(x, y) = \text{No}$, every optimal auction awards bidder 1 the item with probability 0 on input (t_1^j, t_2^j) . Therefore, if we know any outcome consistent with any optimal auction on (t_1^j, t_2^j) , we know $\text{Disjointness}(x, y)$.

4 OUR REDUCTION AND ITS PROPERTIES

In this section, we define our reduction and state some useful properties. First, we define the type space. Throughout this section, n denotes the size of the input to Disjointness . We state the concrete lemmas which are relevant to give a detailed technical proof overview, but all proofs of these lemmas are in Appendix B in the full version. Note that the purpose of this section is *only* to define our flow and state basic properties. We will give intuition for these decisions in the subsequent sections as it will only become clear once we define our flow.

The type space. In our reduction, the type space does not depend on x, y (only the distribution does). For every input, $n_1 = n_2 = n + 2$ (meaning that each bidder has a total of $n + 2$ non-zero types per day, and $2n + 4$ non-zero types in total). For all $k \in [n + 2]$ and both i , $v_i^k := n^2 + k$.

The distribution. The distribution in our construction depends on x, y , but in all cases is nearly uniform. Below, for simplicity of notation let $b := 10n^6$, $a := \frac{b-n^5}{n+1}$. All probabilities will be an integer multiple of $\frac{1}{2b}$. Below, Bidder One's day1 distribution is fixed, and does not depend on x .

DEFINITION 7 (BIDDER ONE'S DAY1 DISTRIBUTION). Define $f_1((v_1^k, 1))$ (for all x) as follows:

- (1) Define $f_1((v_1^1, 1)) := \frac{b}{10n} = \frac{1}{20n}$.
- (2) For $k = 1$ to n , first define helper $z_{k+1} := \frac{b - \sum_{j=1}^k f_1((v_1^j, 1)) \cdot 2b}{n-k+2}$, then define $f_1((v_1^{k+1}, 1)) := \frac{\lfloor z_{k+1} + \frac{n^3}{n-k+2} \rfloor}{2b}$.
- (3) For $k = n + 1$, define helper $z_{n+2} := b - \sum_{j=1}^{n+1} f_1((v_1^j, 1)) \cdot 2b$, then define $f_1((v_1^{n+2}, 1)) := \frac{z_{n+2}}{2b}$.

We quickly establish that the total mass of Bidder One on day1 is always $1/2$ in this construction.

LEMMA 4.1. $\sum_{k=1}^{n+2} f_1((v_1^k, 1)) = 1/2$.

Lemma 4.2 is one key property which will be useful in our later analysis. It states that Bidder One's day1 distribution is nearly-uniform over v_1^2, \dots, v_1^{n+2} (recall that $a > n^5$).

LEMMA 4.2. $f_1((v_1^k, 1)) \cdot 2b \in [a - 2n^3, a + 2n^3]$ for all $k \in [2, n + 2]$.

We now proceed to construct the day2 distribution for Bidder One. Bidder One's day2 distribution depends on x , and is constructed so that x_k has a significant impact on $f_1((v_1^{k+1}, 2))$.

DEFINITION 8 (BIDDER ONE'S DAY2 DISTRIBUTION). Define $f_1((v_1^k, 2))$ (as a function of x) as follows:

- (1) Set $f_1((v_1^1, 2)) := \frac{b}{10n} = \frac{1}{20n}$, for all x .
- (2) For $k = 1$ to n , define helper $z_{k+1} = \frac{b - \sum_{j=1}^k f_1((v_1^j, 2)) \cdot 2b}{n-k+2}$.
 - If $x_k = 0$, then set $f_1((v_1^{k+1}, 2)) := \frac{\lfloor z_{k+1} + \frac{n^3}{n-k+2} \rfloor}{2b}$.
 - Otherwise ($x_k = 1$), set $f_1((v_1^{k+1}, 2)) := \frac{\lceil z_{k+1} \rceil}{2b}$.
- (3) For $k = n + 1$, define helper $z_{n+2} := b - \sum_{j=1}^{n+1} f_1((v_1^j, 2)) \cdot 2b$. Set $f_1((v_1^{n+2}, 2)) := \frac{z_{n+2}}{2b}$

The two lemmas below similarly establish that the total mass of Bidder One on day2 is always $1/2$, and that Bidder One's day2 distribution is always nearly-uniform over v_1^2, \dots, v_1^{n+2} .

LEMMA 4.3. $\sum_{k=1}^{n+2} f_1((v_1^k, 2)) = 1/2$.

LEMMA 4.4. For all x , $f_1((v_1^k, 2)) \cdot 2b \in [a - n^3, a + n^3]$ for all $k \in [2, n+2]$.

Finally, we define the distribution for Bidder Two day1. Bidder Two's distribution will be truly single-parameter in that their interest is always day1. Bidder Two's distribution depends on y , and is constructed so that y_k has a significant impact on $f_2((v_2^{k+1}, 1))$.

DEFINITION 9 (BIDDER TWO'S DISTRIBUTION). Define $f_2((v_2^k, 1))$ (as a function of y) as follows:

(1) Set $f_2((v_2^1, 1)) := \frac{\frac{b}{10n} - 1}{b}$.

(2) For $k = 1$ to n , define helper $z_{k+1} := \frac{b - \sum_{j=1}^k f_2((v_2^j, 1)) \cdot b}{n - k + 2}$.

- If $y_k = 1$, then set $f_2((v_2^{k+1}, 1)) := \frac{\lfloor z_{k+1} + \frac{n^2}{n - k + 2} \rfloor}{b}$.

- Otherwise, set $f_2((v_2^{k+1}, 1)) := \frac{\lfloor z_{k+1} - 1 \rfloor}{b}$.

(3) For $k = n+1$, define helper $z_{n+2} := b - \sum_{j=1}^{n+1} f_2((v_2^j, 1)) \cdot b$, set $f_2((v_2^{n+2}, 1)) := \frac{z_{n+2}}{b}$.

Again, we confirm quickly that this is a valid distribution, and that it is nearly-uniform over v_2^2, \dots, v_2^{n+2} .

LEMMA 4.5. $\sum_{k=1}^{n+2} f_2((v_2^k, 1)) = 1$.

LEMMA 4.6. $f_2((v_2^k, 1)) \cdot b \in [a - 2n^3, a + 2n^3]$ for all $k \in [2, n+2]$.

Finally, we quickly state that the optimal *single-bidder* auction for any distribution considered in our reduction is especially simple: it sets the same take-it-or-leave-it price of $n^2 + 1$. The proof is in Appendix D in the full version.

PROPOSITION 4.7. For all x (resp., y), the revenue-optimal single-bidder auction for the resulting distribution D_1 (resp. D_2) simply sets a take-it-or-leave-it price of $n^2 + 1$ on one-day shipping.

5 CONSTRUCTING A FLOW

In this section, we construct a flow which is optimal for all instances of our construction. We proceed in two steps. First, we consider a canonical flow and establish that this flow is optimal if and only if $\text{Disjointness}(x, y) = \text{yes}$. Next, we show how to modify the flow to be optimal when $\text{Disjointness}(x, y) = \text{no}$.

5.1 A Canonical Flow

We first define a canonical flow, and then argue it is optimal when $\text{Disjointness}(x, y) = \text{yes}$.

DEFINITION 10 (CANONICAL FLOW). (α^M, λ^M) is the canonical Myerson flow, where:

- $\alpha_i(k) = 0$ for both bidders i and all $k \in [n+2]$.
- $\lambda_i^j(k) = R_i((v_i^k, j))$ for both bidders i , days j , and all $k \in [n+2]$.

It is easy to confirm that (α^M, λ^M) is a flow. We can also quickly execute Definition 5 (recalling that $v_i^k = n^2 + k$) to compute $\Phi^{\alpha^M, \lambda^M}$:

OBSERVATION 2. For both bidders i , days j , and all $k \in [1, n+2]$, $\Phi_i^{\alpha^M, \lambda^M}((v_i^k, j)) = v_i^k - \frac{R_i((v_i^{k+1}, j))}{f_i((v_i^k, j))}$.

Proposition 5.1 below captures the key properties of our construction and this flow. These properties are motivated by bullet three of Definition 6: we need to compare virtual values of types of Bidder One with those for types of Bidder Two to determine if a certain allocation is optimal. The proof is in Appendix C in the full version.

PROPOSITION 5.1. *For all x, y , the flow (α^M, λ^M) satisfies the following:*

- *For both $j: k > k' \Rightarrow \Phi_1^{\alpha^M, \lambda^M}((v_1^k, j)) > \Phi_2^{\alpha^M, \lambda^M}((v_2^{k'}, 1))$.*
- *For both $j: k < k' \Rightarrow \Phi_1^{\alpha^M, \lambda^M}((v_1^k, j)) < \Phi_2^{\alpha^M, \lambda^M}((v_2^{k'}, 1))$.*
- *For all $k \in [n]:$ if $x_k = 0$ OR $y_k = 0$, then for both $j: \Phi_1^{\alpha^M, \lambda^M}((v_1^{k+1}, j)) > \Phi_2^{\alpha^M, \lambda^M}((v_2^{k+1}, 1))$.*
- *For all $k \in [n]$, if $x_k = 1$ AND $y_k = 1$, then: $\Phi_1^{\alpha^M, \lambda^M}((v_1^{k+1}, 1)) > \Phi_2^{\alpha^M, \lambda^M}((v_2^{k+1}, 1)) > \Phi_1^{\alpha^M, \lambda^M}((v_1^{k+1}, 2))$.*
- *For both $j: \Phi_1^{\alpha^M, \lambda^M}((v_1^1, j)) > \Phi_2^{\alpha^M, \lambda^M}((v_2^1, 1))$.*
- *For both $j: \Phi_1^{\alpha^M, \lambda^M}((v_1^{n+2}, j)) = \Phi_2^{\alpha^M, \lambda^M}((v_2^{n+2}, 1))$.*
- *For both i , both j , and all $k \geq 1: \Phi_i^{\alpha^M, \lambda^M}((v_i^k, j)) > 0$.*

The first two bullets assert that a bidder with strictly higher value has strictly higher virtual value as well. The next two bullets concern virtual values when both bidders' values are the same. Importantly, they assert that the relative comparison of virtual values when both bidders have value $n^2 + k$ depends *only on x_k and y_k , and not on x_{-k} or y_{-k}* . Just as importantly, they assert that Bidder One's interest is *only relevant if $x_k = y_k = 1$* . Bullet seven implies that any allocation rule that witnesses optimality for (α^M, λ^M) must learn which bidder has higher virtual value. We analyze a potential such auction next.

DEFINITION 11 (SECOND-PRICE AUCTION, TIE-BREAKING FOR BIDDER ONE). *The second-price auction, tie-breaking for Bidder One, gives the item to the bidder with highest value and breaks ties in favor of Bidder One. Payments are charged to satisfy the payment identity.*

THEOREM 5.2. *The second-price auction, tie-breaking for Bidder One, witnesses optimality for (α^M, λ^M) if and only if $\text{Disjointness}(x, y) = \text{yes}$.*

PROOF. First, it is well-known (and easy to see) that the second-price auction, tie-breaking for Bidder One, with the payment identity is BIC. After this, there are three bullets to check in Definition 6. We claim that the first two hold for all x, y , and the third holds if and only if $\text{Disjointness}(x, y) = \text{yes}$.

The first bullet holds trivially, as payments are specifically defined to satisfy the payment identity.

The second bullet holds vacuously, as $\alpha_i^M(k) = 0$ for both i and all k (in fact, the implied condition holds anyway, as the bidders' allocation/payment doesn't depend on their interest).

To see that final bullet holds if and only if $\text{Disjointness}(x, y) = \text{yes}$, observe that the second-price auction awards the item to the bidder with highest value, tie-breaking for Bidder One. So the final bullet holds if and only if: (a) a higher value implies a higher virtual value (which immediately follows from the first two bullets of Proposition 5.1), (b) all virtual values are non-negative (which immediately follows from bullet seven of Proposition 5.1), and (c) Bidder One has a higher virtual value whenever both bidders have the same value (which holds *if and only if* $\text{Disjointness}(x, y) = \text{yes}$, by bullets three and four of Proposition 5.1).

This completes the proof: Definition 6 is satisfied if and only if $\text{Disjointness}(x, y) = \text{yes}$. \square

Observe that Theorem 5.2 implies that the Second-Price Auction, tie-breaking for Bidder One is one optimal auction for (D_1, D_2) when $\text{Disjointness}(x, y) = \text{yes}$. We now conclude the following simple corollary:

COROLLARY 5.2.1. *If $\text{Disjointness}(x, y) = \text{yes}$, every optimal BIC auction (X, P) for D_1, D_2 has $X_1(t_1^1, t_2^1) = 1$.*

PROOF. Because a BIC auction witnesses optimality for (α^M, λ^M) (by Theorem 5.2), every optimal BIC auction for D_1, D_2 witnesses optimality for (α^M, λ^M) . Because $\Phi_1^{\alpha^M, \lambda^M}((v_1^1, 1)) >$

$\Phi_2^{\alpha^M, \lambda^M}((v_2^1, 1))$ by Proposition 5.1, bullet three of Definition 6 asserts that every optimal BIC auction satisfies $X_1(t_1^1, t_2^1) = 1$. \square

Corollary 5.2.1 proves half of Theorem 1.1: that Bidder One wins the item in all optimal auctions on (t_1^1, t_2^1) when $\text{Disjointness}(x, y) = \text{yes}$. Theorem 5.2 makes clear the key distinction when $\text{Disjointness}(x, y) = \text{no}$: (α^M, λ^M) does not witness optimality, so we need a new flow.

5.2 Modifying the Canonical Flow

We now modify the canonical flow to find an optimal (α', λ') in the case when $\text{Disjointness}(x, y) = \text{no}$. Fortunately, the necessary modification is simple to describe (although verifying the desired properties is complex). We will only make one modification, defined below, and first used in [24].

DEFINITION 12 (BOOSTING, [24]). *Beginning with a flow (α, λ) , boosting (α, λ) at k for Bidder i by ε , for any $\varepsilon \leq \lambda_i^2(k')$ for all $k' \leq k$, produces a new flow (α', λ') with:*

- $\alpha'_i(k) := \alpha_i(k) + \varepsilon$.
- $(\lambda')_i^2(k') := \lambda_i^2(k') - \varepsilon$, for all $k' \leq k$.
- $(\lambda')_i^1(k') := \lambda_i^1(k') + \varepsilon$, for all $k' \leq k$.
- If not already specified, then $\alpha' = \alpha$ and $\lambda' = \lambda$.

It is not hard to see that Boosting at k preserves the flow conditions (provided that $\varepsilon \leq \lambda_i^2(k')$ for all $k' \leq k$). It is also not hard to see that Boosting at k for Bidder i *increases* the virtual value for all $(v_i^{k'}, 2)$ for all $k' < k$, *decreases* the virtual value for all $(v_i^{k'}, 1)$ for all $k' < k$, and leaves all other virtual values unchanged (see [24, Observation 3] — although we will prove this ourselves whenever this is used in calculations).

DEFINITION 13 (MODIFIED FLOW). *The modified flow (α^*, λ^*) proceeds as follows:*

- (1) *Begin with $(\alpha, \lambda) = (\alpha^M, \lambda^M)$.*
- (2) *Boost (α, λ) at $n + 2$ for Bidder One by ε . Here, ε is the minimum boost which results in $\Phi_1^{\alpha', \lambda'}((v_1^k, 2)) \geq \Phi_2^{\alpha', \lambda'}((v_2^k, 1))$ for all k .*

The rest of our analysis proceeds as follows. First, we need to establish that this modified flow indeed exists, because the required boost for Bullet 2 is small enough to be valid. Proposition 5.3 states this, and also several useful properties of this flow. The proof of Proposition 5.3 is in Appendix C in the full version, and this relies on many of the precise choices in defining our instance.

PROPOSITION 5.3. *For all $x, y, (\alpha^*, \lambda^*)$ is a valid flow. Moreover, it satisfies the following properties:*

- For both $j: k > k' \Rightarrow \Phi_1^{\alpha^*, \lambda^*}((v_1^k, j)) > \Phi_2^{\alpha^*, \lambda^*}((v_2^{k'}, 1))$.
- For both $j: k < k' \Rightarrow \Phi_1^{\alpha^*, \lambda^*}((v_1^k, j)) < \Phi_2^{\alpha^*, \lambda^*}((v_2^{k'}, 1))$.
- For both j , and all $k \geq 2$, $\Phi_1^{\alpha^*, \lambda^*}((v_1^k, j)) \geq \Phi_2^{\alpha^*, \lambda^*}((v_2^k, 1))$.
- When $\text{Disjointness}(x, y) = \text{no}$, there exists a $k^* \in [2, n + 1]$ such that: $\Phi_1^{\alpha^*, \lambda^*}((v_1^{k^*}, 2)) = \Phi_2^{\alpha^*, \lambda^*}((v_2^{k^*}, 1))$.
- When $\text{Disjointness}(x, y) = \text{no}$, then: $\Phi_1^{\alpha^*, \lambda^*}((v_1^1, 2)) > \Phi_2^{\alpha^*, \lambda^*}((v_2^1, 1)) > \Phi_1^{\alpha^*, \lambda^*}((v_1^1, 1))$.
- For both $j: \Phi_1^{\alpha^*, \lambda^*}((v_1^{n+2}, j)) = \Phi_2^{\alpha^*, \lambda^*}((v_2^{n+2}, 1))$.
- For both i , both j , and all $k \geq 1: \Phi_i^{\alpha^*, \lambda^*}((v_i^k, j)) > 0$.

We now define an auction that witnesses optimality for (α^*, λ^*) , and conclude implications for $X_1(t_1^1, t_2^1)$.

DEFINITION 14 (SECOND-PRICE AUCTION, CAREFUL TIE-BREAKING AT k^*). *The second-price auction with careful tie-breaking at $k^* \in [2, n + 1]$ gives the item to the bidder with highest value. If both*

bidders have the same value $n^2 + k$, break ties in the following manner (in all cases, charge payments satisfying the payment identity):

- If $k \neq 1$, and Bidder One's interest is day1, give the item to Bidder One.
- If $k = 1$, and Bidder One's interest is day1, give the item to Bidder Two.
- If $k \neq k^*$, and Bidder One's interest is day2, give the item to Bidder One.
- If $k = k^*$, and Bidder One's interest is day2, give Bidder One the item with probability $1 - \frac{f_2((v_2^1, 1))}{f_2((v_2^{k^*}, 1))}$, and to Bidder Two with probability $\frac{f_2((v_2^1, 1))}{f_2((v_2^{k^*}, 1))}$.²¹

Let us quickly get some intuition for the Second-Price Auction with careful tie-breaking at k^* . First, observe that when Bidder One's value is neither $n^2 + 1$ nor $n^2 + k^*$, the allocation rule is agnostic to Bidder One's interest. However, when Bidder One's value is $n^2 + 1$, ties are more often broken in favor of Bidder One when their interest is day2 versus day1. Similarly, when their value is $n^2 + k^*$, ties are more often broken in favor of Bidder One when their interest is day1 versus day2. When calculating the payment identity, this implies that *no matter Bidder One's value, their payment depends on their interest, even when their allocation probability does not*. In particular, the Second-Price Auction with careful tie-breaking at k^* is not DSIC, and the precise probability chosen in bullet four is chosen exactly so that Lemma 5.4 (below) holds.²²

LEMMA 5.4. *For all x, y such that $\text{Disjointness}(x, y) = \text{no}$, the Second-Price Auction with careful tie-breaking at k^* is BIC. Moreover, $\pi_i((v_i^k, 2)) \cdot v_i^k - p_i((v_i^k, 2)) = \pi_i((v_i^k, 1)) \cdot v_i^k - p_i((v_i^k, 1))$ for all $k > k^*$.*

With Lemma 5.4 in hand, the proof of Theorem 5.5 follows similarly to that of Theorem 5.2.

THEOREM 5.5. *When $\text{Disjointness}(x, y) = \text{no}$, let k^* be the index promised by bullet four of Proposition 5.3. Then the second-price auction with careful tie-breaking at k^* witnesses optimality for (α^*, λ^*) .*

PROOF. We have already established in Lemma 5.4 that the second-price auction with careful tie-breaking at k^* is BIC, so we just need to check the three bullets. We have also explicitly defined payments to satisfy the payment identity, so bullet one is satisfied. For bullet two, the condition is vacuously satisfied for all $k < n + 2$ because $\alpha_1(k) = 0$. At $k = n + 2$, we are guaranteed that $\pi_i((v_i^{n+2}, 2)) \cdot v_i^{n+2} - p_i((v_i^{n+2}, 2)) = \pi_i((v_i^{n+2}, 1)) \cdot v_i^{n+2} - p_i((v_i^{n+2}, 1))$ by Lemma 5.4, as $k^* < n + 2$. Therefore, bullet two is satisfied for all k .

Finally, we just need to confirm bullet three: that the auction always awards the item to a bidder with highest non-negative virtual value. Indeed, bullets one, two, and seven of Proposition 5.3 imply that the bidder with highest value also has the highest non-negative virtual value, so the second-price auction with careful tie-breaking at k^* is correct whenever the two bidders have different values. Bullet three confirms that Bidder One's virtual value is always weakly higher than Bidder Two's in case they have the same value $> n^2 + 1$. This implies that the second-price auction with careful tie-breaking at k^* breaks ties correctly in all cases when it gives the item to Bidder One. When both bidders have value $n^2 + 1$ and Bidder One's interest is day1, bullet five confirms that Bidder Two has higher virtual value (and the auction gives the item to Bidder Two). When

²¹Observe that this is feasible, as we've guaranteed in our construction that $f_2((v_2^1, 1)) < f_2((v_2^k, 1))$ for all $k > 1$.

²²For example, if Bidder One wins ties in bullet four with any probability $> 1 - \frac{f_2((v_2^1, 1))}{f_2((v_2^{k^*}, 1))}$, the auction would remain BIC

(but the 'Moreover,...' property in Lemma 5.4 would not hold). If Bidder One wins ties with probability $< 1 - \frac{f_2((v_2^1, 1))}{f_2((v_2^{k^*}, 1))}$, then Bidder One would have strict incentive to misreport their day2 interest as day1 whenever their value is $> n^2 + k^*$.

both bidders have value $n^2 + k^*$ and Bidder One's interest is k^* , bullet four confirms that the bidders have the same virtual value, so ties can be broken arbitrarily (and in particular, the randomization proposed is guaranteed to award the item to a bidder of highest virtual value).

This confirms all three bullets, and the proof. \square

Again observe that Theorem 5.5 implies that the Second-Price Auction with careful tie-breaking at k^* is one optimal auction for (D_1, D_2) when $\text{Disjointness}(x, y) = \text{no}$. We again conclude the following corollary:

COROLLARY 5.5.1. *If $\text{Disjointness}(x, y) = \text{no}$, every optimal BIC auction (X, P) for D_1, D_2 has $X_1(t_1^1, t_2^1) = 0$.*

PROOF. Because a BIC auction witnesses optimality for (α^*, λ^*) (by Theorem 5.5), every optimal BIC auction for D_1, D_2 witnesses optimality for (α^*, λ^*) . Because $\Phi_1^{\alpha^*, \lambda^*}((v_1^1, 1)) < \Phi_2^{\alpha^*, \lambda^*}((v_2^1, 1))$ by Proposition 5.3, bullet three of Definition 6 asserts that every optimal BIC auction satisfies $X_1(t_1^1, t_2^1) = 0$. \square

This wraps up the proof of Theorem 1.1.

PROOF OF THEOREM 1.1. Corollary 5.2.1 establishes that when $\text{Disjointness}(x, y) = \text{yes}$, any optimal BIC auction must allocate the item to Bidder One on (t_1^1, t_2^1) with probability one. Corollary 5.5.1 establishes that when $\text{Disjointness}(x, y) = \text{no}$, any optimal BIC auction must allocate the item to Bidder One on (t_1^1, t_2^1) with probability zero. Because D_1 can be constructed only as a function of x , and D_2 can be constructed only as a function of y , any communication protocol which correctly allocates the item on (t_1^1, t_2^1) in accordance with *any* optimal BIC mechanism (even with probability 2/3) can also solve Disjointness (with probability 2/3). Because any deterministic (resp. randomized, succeeding with probability 2/3) protocol for disjointness requires communication n (resp. $\Omega(n)$), this means that any deterministic (resp. randomized, succeeding with probability 2/3) protocol which can correctly allocate the item on (t_1^1, t_2^1) in accordance with *any* optimal BIC mechanism (resp. with probability 2/3) requires communication at least n (resp. $\Omega(n)$). \square

6 CONCLUSION

We establish that optimal multi-dimensional mechanisms are not locally-implementable: in order to evaluate the auction on just a single valuation profile, one must know (essentially) the entire distribution. In contrast, optimal single-dimensional mechanisms are locally-implementable: evaluating the auction on a single valuation profile requires barely more bits from each D_i than simply stating v_i itself. Our construction establishes that this separation holds already in (essentially) the simplest possible multi-dimensional setting: one single-dimensional bidder and one two-day FedEx bidder. We also show that optimal auctions for single-dimensional buyers with public budget constraints are not locally-implementable. Moreover, both results follow the same outline, highlighting the robustness of our techniques.

Our work establishes a novel complexity of optimal multi-dimensional mechanisms distinct from optimal single-dimensional mechanisms. In particular, unlike prior work, this complexity is inherently a multi-bidder phenomenon, rather than inherited from the single-bidder setting. Indeed, every optimal single-bidder auction for any instance considered by our reductions simply sets a price of $n^2 + 1$. Locality can serve as a quantitative lens for future work to study the complexity of multi-bidder auctions in multi-dimensional settings where single-bidder auctions are tractable. For example:

- Do there exist approximately-optimal multi-dimensional auctions that are locally-implementable? One significant technical barrier to this direction is an alternative line of attack beyond complementary slackness (as complementary slackness holds only for optimal primal/dual pairs). Note that the Marginal Revenue Mechanism of [2] is locally-implementable and approximately optimal in restricted “approximately revenue-linear” settings. But, it remains unknown whether approximately-optimal locally-implementable mechanisms exist generally.
- What are the implications of (non)-locality for streaming or online-learning variants of optimal auction design? In this direction, it is important that we study locality via communication complexity, due to the strong connection between communication complexity and streaming lower bounds [4, 40].

REFERENCES

- [1] Saeed Alaei. 2014. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. *SIAM J. Comput.* 43, 2 (2014), 930–972. <https://doi.org/10.1137/120878422>
- [2] Saeed Alaei, Hu Fu, Nima Haghpanah, and Jason Hartline. 2013. The Simple Economics of Approximately Optimal Auctions. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- [3] Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. 2012. Bayesian Optimal Auctions via Multi- to Single-agent Reduction. In *the 13th ACM Conference on Electronic Commerce (EC)*.
- [4] Noga Alon, Yossi Matias, and Mario Szegedy. 1999. The Space Complexity of Approximating the Frequency Moments. *J. Comput. Syst. Sci.* 58, 1 (1999), 137–147. <https://doi.org/10.1006/jcss.1997.1545>
- [5] Moshe Babaioff, Liad Blumrosen, and Michael Schapira. 2013. The communication burden of payment determination. *Games and Economic Behavior* 77, 1 (2013), 153 – 167. <https://doi.org/10.1016/j.geb.2012.08.007>
- [6] 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>
- [7] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2015. Pricing lotteries. *J. Economic Theory* 156 (2015), 144–174. <https://doi.org/10.1016/j.jet.2014.04.011>
- [8] Patrick Briest and Piotr Krysta. 2007. Buying Cheap is Expensive: Hardness of Non-Parametric Multi-Product Pricing. In *the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [9] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2012. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*. 459–478. <https://doi.org/10.1145/2213977.2214021>
- [10] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. 2013. Understanding Incentives: Mechanism Design Becomes Algorithm Design. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*. 618–627. <https://doi.org/10.1109/FOCS.2013.72>
- [11] 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)*.
- [12] Gabriel Carroll. 2017. Robustness and Separation in Multidimensional Screening. *Econometrica* 85, 2 (2017), 453–488.
- [13] Shuchi Chawla, David L. Malec, and Azarakhsh Malekian. 2011. Bayesian mechanism design for budget-constrained agents. In *Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011*. 253–262. <https://doi.org/10.1145/1993574.1993613>
- [14] Yeon-Koo Che and Ian Gale. 2000. The Optimal Mechanism for Selling to a Budget-Constrained Buyer. *Journal of Economic theory* 92, 2 (2000), 198–233.
- [15] 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>
- [16] 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>
- [17] Xi Chen, George Matikas, Dimitris Paparas, and Mihalis Yannakakis. 2018. On the Complexity of Simple and Optimal Deterministic Mechanisms for an Additive Buyer. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018, Artur Czumaj (Ed.)*. SIAM, 2036–2049. <https://doi.org/10.1137/1.9781611975031.133>
- [18] Natalie Collina and S. Matthew Weinberg. 2020. On the (in-)approximability of Bayesian Revenue Maximization for a Combinatorial Buyer. In *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary*,

July 13–17, 2020, Péter Biró, Jason Hartline, Michael Ostrovsky, and Ariel D. Procaccia (Eds.). ACM, 477–497. <https://doi.org/10.1145/3391403.3399496>

- [19] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2012. Optimal Pricing is Hard. In *the 8th Workshop on Internet & Network Economics (WINE)*.
- [20] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2014. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [21] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2017. Strong Duality for a Multiple-Good Monopolist. *Econometrica* 85, 3 (2017), 735–767.
- [22] Nikhil R. Devanur, Kira Goldner, Raghuvansh R. Saxena, Ariel Schwartzman, and S. Matthew Weinberg. 2020. Optimal Mechanism Design for Single-Minded Agents. In *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13–17, 2020*, Péter Biró, Jason Hartline, Michael Ostrovsky, and Ariel D. Procaccia (Eds.). ACM, 193–256. <https://doi.org/10.1145/3391403.3399454>
- [23] Nikhil R. Devanur, Nima Haghpanah, and Christos-Alexandros Psomas. 2017. Optimal Multi-Unit Mechanisms with Private Demands. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26–30, 2017*, Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin (Eds.). ACM, 41–42. <https://doi.org/10.1145/3033274.3085122>
- [24] Nikhil R. Devanur and S. Matthew Weinberg. 2017. The Optimal Mechanism for Selling to a Budget Constrained Buyer: The General Case. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26–30, 2017*, 39–40. <https://doi.org/10.1145/3033274.3085132>
- [25] Shahar Dobzinski, Hu Fu, and Robert D. Kleinberg. 2011. Optimal Auctions with Correlated Bidders are Easy. In *the 43rd ACM Symposium on Theory of Computing (STOC)*.
- [26] Shahar Dobzinski and Shiri Ron. 2021. The communication complexity of payment computation. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021*, Samir Khuller and Virginia Vassilevska Williams (Eds.). ACM, 933–946. <https://doi.org/10.1145/3406325.3451083>
- [27] Ronald Fadel and Ilya Segal. 2009. The communication cost of selfishness. *Journal of Economic Theory* 144, 5 (2009), 1895–1920.
- [28] Amos Fiat, Kira Goldner, Anna R. Karlin, and Elias Koutsoupias. 2016. The FedEx Problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24–28, 2016*, 21–22. <https://doi.org/10.1145/2940716.2940752>
- [29] Nima Haghpanah and Jason Hartline. 2015. Reverse Mechanism Design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15–19, 2015*.
- [30] Sergiu Hart and Noam Nisan. 2019. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *J. Econ. Theory* 183 (2019), 991–1029. <https://doi.org/10.1016/j.jet.2019.07.006>
- [31] Sergiu Hart and Philip J. Reny. 2015. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics* 10, 3 (2015), 893–922.
- [32] Bala Kalyanasundaram and Georg Schnitger. 1992. The Probabilistic Communication Complexity of Set Intersection. *SIAM J. Discret. Math.* 5, 4 (1992), 545–557. <https://doi.org/10.1137/0405044>
- [33] Eyal Kushilevitz and Noam Nisan. 1997. *Communication complexity*. Cambridge University Press.
- [34] Jean-Jacques Laffont and Jacques Robert. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2 (1996), 181–186.
- [35] A. M. Manelli and D. R. Vincent. 2007. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory* 137, 1 (2007), 153–185.
- [36] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [37] 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>
- [38] Malleesh M. Pai and Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150, C (2014), 383–425.
- [39] Gregory Pavlov. 2011. Optimal Mechanism for Selling Two Goods. *The B.E. Journal of Theoretical Economics* 11, 3 (2011).
- [40] Anup Rao and Amir Yehudayoff. 2020. *Communication Complexity and Applications*. Cambridge University Press.
- [41] Alexander A. Razborov. 1992. On the Distributional Complexity of Disjointness. *Theor. Comput. Sci.* 106, 2 (1992), 385–390. [https://doi.org/10.1016/0304-3975\(92\)90260-M](https://doi.org/10.1016/0304-3975(92)90260-M)
- [42] Jean-Charles Rochet and Lars A. Stole. 2003. The Economics of Multidimensional Screening. *Advances in Economic Theory, 7th World Congress* (2003).
- [43] Aviad Rubinstein, Raghuvansh R. Saxena, Clayton Thomas, S. Matthew Weinberg, and Junyao Zhao. 2021. Exponential communication separations between notions of selfishness. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021*, Samir Khuller and Virginia Vassilevska Williams (Eds.).

ACM, 947–960. <https://doi.org/10.1145/3406325.3451127>

[44] 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>

[45] Aviad Rubinstein and Junyao Zhao. 2021. The randomized communication complexity of optimal randomized auctions. In *Symposium on Theory of Computing, STOC 2021*. ACM.

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

[47] Andrew Chi-Chih Yao. 2017. Dominant-Strategy versus Bayesian Multi-item Auctions: Maximum Revenue Determination and Comparison. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26–30, 2017*, Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin (Eds.). ACM, 3–20. <https://doi.org/10.1145/3033274.3085120>

[48] Sergey Yekhanin. 2012. Locally Decodable Codes. *Found. Trends Theor. Comput. Sci.* 6, 3 (2012), 139–255. <https://doi.org/10.1561/0400000030>

A OPTIMAL SINGLE-DIMENSIONAL MECHANISMS ARE LOCAL

In this section we are going to show that optimal single-item auctions are local. Our analysis here adopts the characterization by [11].

We assume each D_i is a distribution supported on n valuations $(v_i^1, v_i^2, \dots, v_i^n)$, and each valuation has an integer value between 0 and $p(n)$ (p is a polynomial in n) for each outcome, and the probability of each valuation is an integer multiple of $1/g(n)$ (g is another polynomial in n). We use the notation $R_i(v_i^k) := \sum_{k' \geq k} f_i(v_i^{k'})$.²³

First we define the discrete Myerson virtual value.

DEFINITION 15 (SINGLE-DIMENSIONAL VIRTUAL VALUE). *For a discrete distribution D_i , define*

$$\varphi^{D_i}(v_i^k) := v_i^k - \frac{(v_i^{k+1} - v_i^k) \cdot R_i(v_i^{k+1})}{f_i(v_i^k)}.$$

We now show that a single virtual value $\varphi^{D_i}(v_i^k)$ contains barely more information from D_i than the value v_i^k itself (which has zero information from D_i).

LEMMA A.1. *For fixed polynomials p and g , given distribution D_i and k , $\varphi^{D_i}(v_i^k)$ can be represented in $O(\log n)$ bits.*

PROOF. We can use v_i^k , v_i^{k+1} , $R_i(v_i^{k+1})$, and $f_i(v_i^k)$ to compute $\varphi^{D_i}(v_i^k)$, thus we can trivially encode each of them into $O(\log n)$ bits separately. \square

We will then prove a bit stronger result to show the weighted average virtual value over an interval can also be represented in $O(\log n)$ bits.

LEMMA A.2. *For fixed polynomials p and g , given distribution D_i and k, l , the weighted average virtual value over $[k, l]$*

$$\frac{\sum_{j=k}^l f_i(v_i^j) \cdot \varphi^{D_i}(v_i^j)}{\sum_{j=k}^l f_i(v_i^j)}$$

can be represented in $O(\log n)$ bits.

PROOF. Since the denominator $\sum_{j=k}^l f_i(v_i^j)$ is an integer multiple of $1/g(n)$ and less than or equal to 1, it can be easily encoded in $O(\log n)$ bits. So it is sufficient to show that the numerator $\sum_{j=k}^l f_i(v_i^j) \cdot \varphi^{D_i}(v_i^j)$ can be represented in $O(\log n)$ bits. To achieve this, observe that:

$$\sum_{j=k}^l f_i(v_i^j) \cdot \varphi^{D_i}(v_i^j) = \sum_{j=k}^l \left(f_i(v_i^j) \cdot v_i^j - (v_i^{k+1} - v_i^k) \cdot R_i(v_i^{k+1}) \right)$$

²³For simplicity of notation, denote by $R_i(v_i^{n+1}) := 0$.

$$\begin{aligned}
&= \sum_{j=k}^l f_i(v_i^j) \cdot v_i^j - \left(v_i^{l+1} \cdot R_i(v_i^{l+1}) - v_i^k \cdot R_i(v_i^{k+1}) + \sum_{j=k+1}^l f_i(v_i^j) \cdot v_i^j \right) \\
&= f_i(v_i^k) \cdot v_i^k - v_i^{l+1} \cdot R_i(v_i^{l+1}) + v_i^k \cdot R_i(v_i^{k+1}) \\
&= v_i^k \cdot R_i(v_i^k) - v_i^{l+1} \cdot R_i(v_i^{l+1}).
\end{aligned}$$

Note that $v_i^k, R_i(v_i^k), v_i^{l+1}$, and $R_i(v_i^{l+1})$ can all be encoded in $O(\log n)$ bits. \square

The following theorem characterizes the optimal mechanism in this single-dimensional setting.

THEOREM A.3 ([11, 36]). *The revenue-optimal BIC mechanism awards the item to the bidder with the highest non-negative ironed virtual value (if one exists), breaking ties arbitrarily but consistently across inputs, where ironed virtual value for type $v_i^k: \bar{\varphi}_i^{D_i}(v_i^k)$ is the weighted average virtual value over an interval containing v_i^k .²⁴ If no such bidder exists, the item remains unallocated.*

Finally with Lemma A.2 and Theorem A.3 we conclude the following theorem, which implies that the optimal single-dimensional auction is local.

THEOREM A.4. *When each D_i is single-dimensional, supported on n valuations, and each valuation has an integer value between 0 and $\text{poly}(n)$ for each outcome, and the probability of each valuation is an integer multiple of $1/\text{poly}(n)$, then $O(\log n)$ bits from each D_i suffice to compute $\text{Opt}(\vec{v}, \vec{D})$.*

²⁴See [11] for explicit expressions of ironed virtual values. Here we only need to know that an ironed virtual value is the weighted average over an interval of virtual values.