The menu complexity of "one-and-a-half-dimensional" mechanism design

Raghuvansh R. Saxena * Ariel Schvartzman † S. Matthew Weinberg ‡

Abstract

We study the menu complexity of optimal and approximately-optimal auctions in the context of the "FedEx" problem, a so-called "one-and-a-half-dimensional" setting where a single bidder has both a value and a deadline for receiving an item [FGKK16]. The menu complexity of an auction is equal to the number of distinct (allocation, price) pairs that a bidder might receive [HN13]. We show the following when the bidder has n possible deadlines:

- **Exponential menu complexity is necessary to be exactly optimal:** There exist instances where the optimal mechanism has menu complexity $\geq 2^n - 1$. This matches *exactly* the upper bound provided by Fiat et al.'s algorithm, and resolves one of their open questions [FGKK16].
- **Fully polynomial menu complexity is necessary and sufficient for approximation:** For all instances, there exists a mechanism guaranteeing a multiplicative $(1-\epsilon)$-approximation to the optimal revenue with menu complexity $O(n^{3/2}\sqrt{\frac{\min\{n/\epsilon, \ln(v_{\max})\}}{\epsilon}}) = O(n^2/\epsilon)$, where $v_{\max}$ denotes the largest value in the support of integral distributions.
- There exist instances where any mechanism guaranteeing a multiplicative $(1 - O(1/n^2))$-approximation to the optimal revenue requires menu complexity $\Omega(n^2)$.

Our main technique is the polygon approximation of concave functions [Rot92], and our results here should be of independent interest. We further show how our techniques can be used to resolve an open question of [DW17] on the menu complexity of optimal auctions for a budget-constrained buyer.

*Department of Computer Science, Princeton University, rrsaxena@cs.princeton.edu.

†Department of Computer Science, Princeton University, acohenca@cs.princeton.edu.

‡Department of Computer Science, Princeton University, smweinberg@princeton.edu.

1 Introduction

It is by now quite well understood that optimal mechanisms are far from simple: they may be randomized [Tha04, BCKW10, HN13], behave non-monotonically [HR15, RW15], and be computationally hard to find [CDW13, DDT14, CDP$^+$14, Rub16]. To cope with this, much recent attention has shifted to the design of simple, but approximately optimal mechanisms (e.g. [CHK07, CHMS10, HN12, BILW14]). However, the majority of these works take a binary view on simplicity, developing simple mechanisms that guarantee constant-factor approximations. Only recently have researchers started to explore the tradeoff space between simplicity and optimality through the lens of *menu complexity.*

Hart and Nisan first proposed the menu complexity as one quantitative measure of simplicity, which captures the *number of different outcomes* that a buyer might see when participating in a mechanism [HN13]. For example, the mechanism that offers only the grand bundle of all items at price p (or nothing at price 0) has menu complexity 1. The mechanism that offers any single item at price p (or nothing at price 0) has menu complexity n, and randomized mechanisms could have infinite menu complexity.

Still, all results to date regarding menu complexity have really been more qualitative than quantitative. For example, only just now is the state-of-the-art able to show that for a single additive bidder with independent values for multiple items and all $\varepsilon > 0$, the menu complexity required for a $(1 - \varepsilon)$ approximation is finite [BGN17] (and even reaching this point was quite non-trivial). On the quantitative side, the best known positive results for a single additive or unit-demand bidder with independent item values require menu complexity exp(n) for a $(1 - \varepsilon)$-approximation, but the best known lower bounds have yet to rule out that poly(n) menu complexity suffices for a $(1 - \varepsilon)$-approximation in either case. In this context, our work provides *the first nearly-tight quantitative bounds on menu complexity in any multi-dimensional setting.*

1.1 One-and-a-half dimensional mechanism design

The setting we consider is the so-called "FedEx Problem," first studied in [FGKK16]. Here, there is a single bidder with a value v for the item and a deadline

i for receiving it, and the pair (v, i) is drawn from an arbitrarily correlated distribution where the number of possible deadlines is finite (n). The buyer's value for receiving the item by her deadline is v, and her value for receiving the item after her deadline (or not at all) is 0. While technically a two-dimensional problem, optimal mechanisms for the FedEx problem don't suffer the same undesirable properties as "truly" two-dimensional problems. Still, the space of optimal mechanisms is considerably richer than single-dimensional problems (hence the colloquial term "one-and-a-half dimensional"). More specifically, while the optimal mechanism might be randomized, it has menu complexity at most $2^n - 1$, and there is an inductive closed-form solution describing it. Additionally, there is a natural condition on each F_i (the marginal distribution of v conditioned on i) guaranteeing that the optimal mechanism is deterministic (and therefore has menu complexity $\leq n$).[1]

A number of recent (and not-so-recent) works examine similar settings such as when the buyer has a value and budget [LR96, CG00, CMM11, DW17], or a value and a capacity [DHP17], and observe similar structure on the optimal mechanism. Such settings are quickly gaining interest within the algorithmic mechanism design community as they are rich enough for optimal mechanisms to be highly non-trivial, but not quite so chaotic as truly multi-dimensional settings.

1.2 Our results We study the menu complexity of optimal and approximately optimal mechanisms for the FedEx problem. Our first result proves that the $2^n - 1$ upper bound on the menu complexity of the optimal mechanism provided by Fiat et al.'s algorithm is *exactly* tight:

THEOREM 1.1. *For all n, there exist instances of the FedEx problem on n days where the menu complexity of the optimal mechanism is $2^n - 1$.*

From here, we turn to approximation and prove our main results. First, we show that fully polynomial menu complexity suffices for a $(1-\varepsilon)$-approximation. The guarantee below is always $O(n^2/\varepsilon)$, but is often improved for specific instances. Below, if the FedEx instance happens to have integral support and the largest value is $v_{\max}$, we can get an improved bound (but if the support is continuous or otherwise non-integral, we can just take the n/ε term instead).[2]

[1] This condition is called "decreasing marginal revenues," and is satisfied by distributions with CDF F and PDF f such that $x \cdot f(x) - 1 + F(x)$ is monotone non-decreasing.

[2] Actually our bounds can be be improved to replace $v_{\max}$ with many other quantities that are always $\leq v_{\max}$, and will still be well-defined for continuous distributions, more on this in Section 4.

THEOREM 1.2. *For all instances of the FedEx problem on n days, there exists a mechanism of menu complexity $O\left(n\sqrt{\frac{\min\{n/\varepsilon, \ln(v_{\max})\}}{\varepsilon/n}}\right)$ guaranteeing a $(1-\varepsilon)$ approximation to the optimal revenue.*

In Theorem 1.2, observe that for any fixed instance, as $\varepsilon \to 0$, our bound grows like $O(1/\sqrt{\varepsilon})$ (because eventually n/ε will exceed $\ln(v_{\max})$). Similarly, our bound is always $O(n^2/\varepsilon)$ for any $v_{\max}$. Both of these dependencies are provably tight for our approach (discussed shortly in Section 1.3), and in general tight up to a factor of $\sqrt{n \log n}$.[3]

THEOREM 1.3. *For all n, there exists an instance of the FedEx problem on n days with $v_{\max} = O(n)$, such that the menu complexity of every $(1 - O(1/n^2))$-optimal mechanism is $\Omega(n^2)$.*

We consider Theorems 1.2 and 1.3 to be our main results, with Theorem 1.1 motivating the study of approximation in the first place. Taken together, the picture provided by these results is the following:

- Exactly optimal mechanisms can require exponential menu complexity (Theorem 1.1), while $(1-\varepsilon)$-approximate mechanisms exist with fully polynomial menu complexity (Theorem 1.2).
- The menu complexity required to guarantee a $(1-\varepsilon)$-approximation is nailed down within a multiplicative $\sqrt{n \log n}$ gap, and lies in $\left[\Omega\left(\sqrt{n/\log n} \cdot \sqrt{\frac{\min\{n/\varepsilon, \ln(v_{\max})\}}{\varepsilon/n}}\right), O\left(n \cdot \sqrt{\frac{\min\{n/\varepsilon, \ln(v_{\max})\}}{\varepsilon/n}}\right)\right]$ (lower bound: Theorem 1.3, upper bound Theorem 1.2).

1.3 Our techniques We'll provide an intuitive proof overview for each result in the corresponding technical section, but we briefly want to highlight one aspect of our approach that should be of independent interest.

It turns out that the problem of revenue maximization with bounded menu complexity really boils down to a question of how well piece-wise linear functions with bounded number of segments can approximate concave functions (we won't get into details of *why* this is the case until Section 4). This is a quite well-studied problem called *polygon approximation* (e.g. [Rot92, YG97, BHR91]). Questions asked here are typically of the form "for a concave function f and

[3] The gap of $\sqrt{n \log n}$ comes as our upper bound approach requires that we lose at most $\varepsilon\mathsf{OPT}/n$ "per day," while our lower bound approach shows that any mechanism with lower menu complexity loses at least $\varepsilon\mathsf{OPT}$ on some day.

interval $[0, v_{\max}]$ such that $f'(0) = 1$, $f'(v_{\max}) = 0$, what is the minimum number of segments a piece-wise linear function g must have to guarantee $f(x) \geq g(x) \geq f(x) - \varepsilon$ for all $x \in [0, v_{\max}]$?"

The answer to the above question is $\Theta(\sqrt{v_{\max}/\varepsilon})$ [Rot92, YG97]. This bound certainly suffices for our purposes to get *some* bound on the menu complexity of $(1-\varepsilon)$-approximate auctions, but it would be much weaker than what Theorem 1.2 provides (we'd have linear instead of logarithmic dependence on $v_{\max}$, and no option to remove $v_{\max}$ from the picture completely). Interestingly though, for our application absolute additive error doesn't tightly characterize what we need (again, we won't get into why this is the case until Section 4). Instead, we are really looking for the following kind of guarantee, which is a bit of a hybrid between additive and multiplicative: for a concave function f and interval $[0, v_{\max}]$ such that $f'(0) = 1$, $f'(v_{\max}) = 0$, what is the minimum number of segments a piece-wise linear function g must have to guarantee $f(x) \geq g(x) \geq f(x) - \varepsilon - \varepsilon(f(v_{\max}) - f(0))$?

At first glance it seems like this really shouldn't change the problem at all: why don't we just redefine $\varepsilon' := \varepsilon(1 + f(v_{\max}) - f(0))$ and plug into upper bounds of Rote for ε'? This is indeed completely valid, and we could again chase through and obtain some weaker version of Theorem 1.2 that also references additional parameters in unintuitive ways. But it turns out that for all examples in which this $\Omega(\sqrt{v_{\max}/\varepsilon})$ dependence is tight, there is actually quite a large gap between $f(0)$ and $f(v_{\max})$, and a greatly improved bound is possible (which replaces the linear dependence on $v_{\max}$ with logarithmic dependence, and provides an option to remove $v_{\max}$ from the picture completely at the cost of worse dependence on ε).

THEOREM 1.4. *For any concave function f and any $\varepsilon > 0$ such that $f'(0) \leq 1$, $f'(v_{\max}) \geq 0$, there exists a piece-wise linear function g such that $f(x) \geq g(x) \geq f(x) - \varepsilon(1 + f(v_{\max}) - f(0))$ with $\Theta(\sqrt{\ln(v_{\max})/\varepsilon})$ segments, and this is tight.*

If one wishes to remove the dependence on $v_{\max}$, then one can replace the bound with $\Theta(1/\varepsilon)$, which is also tight (among bounds that don't depend on $v_{\max}$).

The proof of Theorem 1.4 is self-contained and appears in Section 4. Both the statement of Theorem 1.4 and our proof will be useful for future work on menu complexity, and possibly outside of mechanism design as well - to the best of our knowledge these kinds of hybrid guarantees haven't been previously considered.[4]

1.4 Related work

Menu complexity. Initial results on menu complexity prove that for a single additive or unit-demand bidder with arbitrarily correlated item values over just 2 items, there exist instances where the optimal (randomized, with infinite menu complexity) mechanism achieves infinite revenue, while any mechanism of menu complexity C achieves revenue $\leq C$ (so no finite approximation is possible with bounded menu complexity) [BCKW10, HN13]. This motivated follow-up work subject to assumptions on the distributions, such as a generalized hazard rate condition [WT14], or independence across item values [DDT13, BGN17]. Even for a single bidder with independent values for two items, the optimal mechanism could have uncountable menu complexity [DDT13], motivating the study of approximately optimal mechanisms subject to these assumptions. Only just recently did we learn that the menu complexity is indeed finite for this setting [BGN17].

It is also worth noting that other notions of simplicity have been previously considered as well, such as the sample complexity (how many samples from a distribution are required to learn an approximately optimal auction?). Here, quantitative bounds are known for the single-item setting (where the menu complexity question is trivial: optimal mechanisms have menu complexity 1) [CR14, HMR15, DHP16, GN17], but again only binary bounds are known for the multi-item setting: few samples suffice for a constant-factor approximation if values are independent [MR15, MR16], while exponentially many samples are required when values are arbitrarily correlated [DHN14]. In comparison to works of the previous paragraphs, we are the first to nail down "the right" quantitative menu complexity bounds in any multi-dimensional setting.

One-and-a-half dimensional mechanism design. One-and-a-half dimensional settings have been studied for decades by economists, the most notable example possibly being that of a single buyer with a value and a budget [LR96, CG00]. Recently, such problems have become popular within the AGT community as optimal auctions are more involved than single-dimensional settings, but not quite so chaotic as truly multidimensional settings [FGKK16, DW17, DHP17]. Each of these works focus exclusively on exactly optimal mechanisms (and exclusively on positive results). In comparison, our work is both the first to prove lower bounds on the complexity of (approximately) optimal mechanisms in these settings, and the first to provide nearly-optimal mechanisms that are considerably less complex.

[4] Interestingly (and completely unrelated to this work), hybrid additive-multiplicative approximations for core problems in online learning have also found use in other recent directions in AGT [DJF16, SBN17].

Polygon approximation. Prior work on polygon approximation is vast, and includes, for instance, core results on univariate concave functions [Rot92, BHR91, YG97], the study of multi-variate functions [Bro08, GG09, DDY16], and even applications in robotics [BGO07]. The more recent work has mostly been pushing toward better guarantees for higher dimensional functions. To the best of our knowledge, the kinds of guarantees we target via Theorem 1.4 haven't been previously considered, and could prove more useful than absolute additive guarantees for some applications.

1.5 Organization In Section 2, we formally describe the FedEx problem and recap the main result of [FGKK16]. In Section 3 we present an instance of the FedEx problem whose menu complexity for optimal auctions is exponential, the worst possible. In Section 4 we present a mechanism that guarantees a (1-ε) fraction of the optimal revenue with a menu complexity of $O(\frac{n^2}{\varepsilon})$. We also explain the connection between approximate auctions and polygon approximation. In Section 5 we present an instance of the FedEx problem that requires a menu complexity of $\Omega(n^2)$ in order to approximate the revenue within $1 - O(1/n^2)$. In Section 7 we use similar techniques to those of Section 3 to construct an example resolving an open question of [DW17].[5]

For the sake of brevity, most proofs have been omitted from this version. Please consult the full version of this paper for detailed proofs.

2 Preliminaries

We consider a single bidder who's type depends on two parameters: a value v and a deadline $i \in [n]$. Deterministic outcomes that the seller can award are just a day $\in [n]$ to ship the item, or to not ship the item at all (and the seller may also randomize over these outcomes). A buyer receives value v if the item is shipped by her deadline, and 0 if it is shipped after her deadline (or not at all).

The types (v, i) are drawn from a known (possibly correlated) distribution $\mathcal{F}$. Let q_i denote the probability that the bidder's deadline is i and $\mathcal{F}_i$ the marginal distribution of v conditioned on a deadline of i. For simplicity of exposition, in several parts of this paper we'll assume that $\mathcal{F}$ is supported on $\{0, 1, \dots, v_{\max}\} \times \{1, \dots, n\}$. This assumption is w.l.o.g., and all results extend to continuous distributions, or distributions with arbitrary discrete support if desired [CDW16].

We assume familiarity with the standard linear program whose solution yields the revenue-optimal auction for the FedEx problem. We only note here the relevant incentive compatibility constraints (observed in [FGKK16]). First, note that w.l.o.g. whenever the buyer has deadline i, the optimal mechanism can ship her the item (if at all) exactly on day i. Shipping the item earlier doesn't make her any happier, but might make the buyer interested in misreporting and claiming a deadline of i if her deadline is in fact earlier. Next note that, subject to this, the buyer never has an incentive to overreport her deadline, but she still might have incentive to underreport her deadline (or misreport her value).

We will be interested in understanding the *menu complexity* of auctions, which is the number of different outcomes that, depending on the buyer's type, are ever selected. If $\pi(v, i)$ denotes the probability that a buyer with value v and deadline i receives the item, then we define the i-deadline menu complexity to be the number of distinct options on deadline i ($|\{p | \exists v, \pi(v, i) = p\}|$). The menu complexity then just sums the i-deadline menu complexities, and we will sometimes refer also to the "deadline menu complexity" as the maximum of the i-deadline menu complexities.

2.1 Optimal auctions for the FedEx problem

Here, we recall some tools from [FGKK16] regarding optimal mechanisms for the FedEx problem. The first tool they use is the notion of a *revenue curve*.[6]

DEFINITION 2.1. (REVENUE CURVES) *For a given deadline i, define the i^{th} revenue curve R_i so that*

$$R_i(v) = q_i \cdot v \cdot \Pr_{x \leftarrow \mathcal{F}_i}[x \geq v].$$

Intuitively, $R_i(v)$ captures the achievable revenue by setting price v exclusively for consumers on deadline i. It is also necessary to consider the *ironed* revenue curve, defined below.

DEFINITION 2.2. (IRONED REVENUE CURVES) *For any revenue curve R_i, define $\tilde{R}_i$ to be its upper concave envelope.*[7] *We say $\tilde{R}_i$ is ironed at v if $\tilde{R}_i(v) \neq R_i(v)$, and we call $[x, y]$ an ironed interval of $\tilde{R}_i$ if $\tilde{R}_i$ is not ironed at x or y, but is ironed at v for all $v \in (x, y)$.*

Of course, it is not sufficient to consider each possible deadline of the buyer in isolation. In particular, offering

[5] Specifically, [DW17] ask whether the optimal mechanism for a single buyer with a private budget and a regular value distribution conditioned on each possible budget is deterministic. The answer is yes if we replace "regular" with "decreasing marginal revenues," or "private budget" with "public budget." We show that the answer is no in general: the optimal mechanism, even subject to regularity, could be randomized.

[6] For those familiar with revenue curves, note that this revenue curve is intentionally drawn in value space, and not quantile space.

[7] That is, $\tilde{R}_i$ is the smallest concave function such that $\tilde{R}_i(x) \geq R_i(x)$ for all x.

certain options on day i constrains what can be offered on days $\geq i$ subject to incentive compatibility. For instance, if some (v, i) pair receives the item with probability 1 on day 1 for price p, no bidder with a deadline ≥ 1 will ever choose to pay $> p$. So we would also like a revenue curve that captures the optimal revenue we can make from days $\geq i$ conditioned on selling the item deterministically at price p on day i. It's not obvious how to construct such a curve, but this is one of the main contributions of [FGKK16], stated below.

Definition 2.3. *Let $R_{\geq n}(v) := R_n(v)$, and $r_{\geq n} := \arg\max_v R_{\geq n}(v)$. Define for $i = n-1$ to 1:*

$$R_{\geq i}(v) = \begin{cases} R_i(v) + \tilde{R}_{\geq i+1}(v) & v < r_{\geq i+1} \\ R_i(v) + \tilde{R}_{\geq i+1}(r_{\geq i+1}) & v \geq r_{\geq i+1}. \end{cases}$$

Lemma 2.1. ([FGKK16]) *$R_{\geq i}(v)$ is the optimal revenue of any mechanism that satisfy the following:*

- *The buyer can either receive the item on day i and pay v, or receive nothing/pay nothing.*
- *The buyer cannot receive the item on any day $< i$.*

Moreover, for any $v_1 < \ldots < v_k$, and $a_i(1), \ldots, a_i(k) \geq 0$ such that $\sum_j a_i(j) \leq 1$, $\sum_j a_i(j) R_{\geq i}(v_j)$ is the optimal revenue of any mechanism that satisfy the following:

- *The buyer can receive the item on day i with probability $\sum_{j \leq \ell} a_i(j)$ and pay $\sum_{j \leq \ell} a_i(j) v_j$, for any $\ell \in [k]$ (or not receive the item on day i and pay nothing).*
- *The buyer cannot receive the item on any day $< i$.*

Finally, we describe the optimal mechanism provided by [FGKK16], which essentially places mass optimally upon each day's revenue curve, subject to constraints imposed by the decisions of previous days. First, simply set any price p maximizing $R_{\geq 1}(p)$ to receive the item on day 1 (as day 1 is unconstrained by previous days). Now inductively, assume that the options for day i have been set and we're deciding what to do for day $i+1$. If the menu options offered on day i are $(\pi_0, p_0), \ldots, (\pi_k, p_k)$ (interpret the option (π_j, p_j) as "charge p_j to ship the item on day i with probability π_j"), think of this instead as a distribution over prices, where price $\frac{p_j - p_{j-1}}{\pi_j - \pi_{j-1}}$ has mass $\pi_j - \pi_{j-1}$.[8] For each such price p, it will undergo one of the following three operations to become an option for day $i+1$.

- If $p \geq r_{\geq i+1}$, move all mass from p to $r_{\geq i+1}$.
- If $\hat{R}_{\geq i+1}$ is not ironed at p, and $p \leq r_{\geq i+1}$, keep all mass at p.
- If $\hat{R}_{\geq i+1}$ is ironed at p, and $p \leq r_{\geq i+1}$, let $[x, y]$ denote the ironed interval containing p, and let $qx + (1-q)y = p$. Move a q fraction of the mass at p to x, and a $(1-q)$ fraction of the mass at p to y.

Once the mass is resettled, if there is mass $a_i(j)$ on price p_j for $p_1 < \ldots < p_k$, the buyer will have the option to receive the item on day i with probability $\sum_{j \leq \ell} a_i(j)$ for price $\sum_{j \leq \ell} a_i(j) p_j$ for any $\ell \in [k]$ (or not at all). Note that due to case three in the transformation above, there could be up to twice as many menu options on day i as day $i-1$.

Theorem 2.1. ([FGKK16]) *The allocation rule described above is the revenue-optimal auction.*

3 Optimal Mechanisms Require Exponential Menu Complexity

In this section we overview our construction for an instance of the FedEx problem with $v_{\max}$ integral values for each day and $n \leq \log(v_{\max})$ days where the i-deadline menu complexity of the optimal mechanism is 2^{i-1} for all i (and this is the maximum possible [FGKK16]), implying that the menu complexity is $2^n - 1$. Note that the deadline menu complexity is always upper bounded by $v_{\max}$, so $v_{\max}$ must be at least 2^n.

At a high level, constructing the example appears straight-forward, once one understands Fiat et al.'s algorithm (end of Section 2). Every menu option from day i is either "shifted" to $r_{\geq i+1}$, "copied," or "split." If the option is shifted or copied, it spawns only a single menu option on day $i+1$, while if split it spawns two (hence the upper bound of $2^n - 1$). So the goal is just to construct an instance where every option is split on every day.

Unfortunately, this is not quite so straight-forward: whether or not an option is split depends on whether it lies inside an ironed interval in this $R_{\geq i}$ curve, which is itself the sum of revenue curves (some ironed and some not), and going back and forth between distributions and sums of revenue curves is somewhat of a mess. So really what we'd like to do is construct the $R_{\geq i}$ curves directly, and be able to claim that there exists a FedEx input inducing them. While not every profile $(R_{\geq 1}, \ldots, R_{\geq n})$ of curves is valid, we do provide a broad class of curves for which it is somewhat clean to show that there exists a FedEx input inducing them.

From here, it is then a matter of ensuring that we can find the revenue curve profiles we want (where for every day i, every menu option is split, because it is inside an ironed interval in $R_{\geq i}$) within our class. We'll highlight parts of our construction below.

[8] This is the standard transformation between "lotteries" and "distributions over prices" (e.g. [RZ83]).

LEMMA 3.1. *For any $v_{\max}$ and $n \le \log(v_{\max})$, there exists an input to the FedEx problem such that:*

- *R_1 is maximized at $v_{\max}/2$ (that is, $R_1(v_{\max}/2) \ge R_1(x)\ \forall x$) and has no ironed intervals.*
- *For all $i > 1$, $\tilde{R}_i$ has a maximizer at price $v \ge 2^i(2^{n-i} - 1)$ and has ironed intervals $[2^{n-i} + k2^{n-i+2}, 2^{n-i} + k2^{n-i+2} + 2^{n-i+1}]$ for $k \in \{0, \dots, 2^{i-2} - 1\}$.*
- *$\tilde{R}_i$ (the ironed revenue curve) is a constant function for all $i \ge 2$.*[9]
- *$R_{\ge i}$ has the same ironed intervals as R_i. In fact, $\forall x$, $R_{\ge i}(x) = R_i(x) + c$ for some constant c.*

As a result of this construction, we see that $R_{\ge i}$ has 2^{i-2} ironed intervals, whose endpoints themselves lie in ironed intervals of $R_{\ge i+1}$. This guarantees that all menu options from day i (which are guaranteed to be endpoints of ironed intervals) are split into two options on day $i + 1$.

THEOREM 3.1. *The optimal mechanism for any instance satisfying the conditons of Lemma 3.1 has i-deadline complexity 2^{i-1} for all i, and menu complexity $2^n - 1$.*

4 Approximately Optimal Mechanisms with Small Menus

In this section, we describe a mechanism that attains at least $1 - \varepsilon$ fraction of the optimal revenue for any FedEx instance with menu complexity $O\left(n\sqrt{\frac{n}{\varepsilon}\min\left(\frac{n}{\varepsilon}, \log v_{\max}\right)}\right)$, which proves Theorem 1.2.

Our main approach is to use the polygon approximation of concave functions applied to revenue curves. For a sequence of points X in the domain of a function f, the polygon approximation $\tilde{f}_X$ of a function with respect to X is the piecewise linear function formed by connecting the points $(x, f(x))$ for $x \in X$ by line segments. Thus, if the sequence X has n points, the function $\tilde{f}_X$ will have $n-1$ segments. For a concave function f, the line joining $(x_1, f(x_1))$ and $(x_2, f(x_2))$ for any two points x_1 and x_2, lies entirely below the function f. Thus, for concave functions f, we have for any sequence X, the value of $f(x) - \tilde{f}_X(x) \ge 0$. Typically, for a 'good' polygon approximation, one requires for $\varepsilon > 0$, that $f(x) - \varepsilon \le \tilde{f}_X(x) \le f(x)$.

It turns out that the question of approximating revenue with low menu complexity boils down to a question of approximating *revenue curves* with piecewise-linear functions of few segments. The connection isn't quite obvious, but isn't overly complicated. Without getting into formal details, here is a rough sketch of what's going on:

- Recall the Fiat et al. procedure to build the optimal mechanism: menu options from deadline $i-1$ might be "split" into two options for deadline i if they lie inside an ironed interval of $\tilde{R}_{\ge i}$. This might cause the menu complexity to double from one deadline to the next.
- Instead, we want to create at most k "anchoring points" on each revenue curve. For a menu option from deadline $i - 1$, instead of distributing it to the endpoints of its ironed interval, we distribute it to the two nearest anchor points.
- By Lemma 2.1, we know exactly how to evaluate the revenue lost by this change, and it turns out this is captured by the maximum gap between $\tilde{R}_{\ge i}(\cdot)$ and the polygon approximation obtained to $\tilde{R}_{\ge i}(\cdot)$ (this isn't obvious, but not hard.
- Finally, it turns out that the i-deadline menu complexity with at most k anchoring points is at most $2k$ (also not quite obvious, but also not hard). So the game is to find few anchoring points that obtain a good polygon approximation to each revenue curve.

COROLLARY 4.1. *Consider a FedEx instance with n deadlines. For all $i \in \{1, 2, \cdots, n\}$, let g_i be the function $\tilde{R}_{\ge i}$ defined in Definition 2.3, and let X_i be a sequence of k_i points in $[0, r_{\ge i}]$ such that for all $x \le r_{\ge i}$, we have $g_i(x) - \varepsilon_i \le \tilde{g}_{i X_i}(x) \le g_i(x)$. Then there exists a mechanism with i-deadline menu complexity $2k_i$ (and menu complexity $2\sum_i k_i$) whose revenue is at least $\mathsf{OPT} - \sum_{i=1}^n \varepsilon_i$.*

Here, OPT denotes the optimal revenue of the FedEx instance.

At this point, it seems like the right approach is to just set each $\varepsilon_i = \varepsilon \cdot \mathsf{OPT}/n$ and plug into the best existing bounds on polygon approximation. In some sense this is correct, but the menu complexity bounds one would obtain are far from optimal. The main insight is that we know something about the curves we wish to approximate: $\tilde{R}(x) \le \mathsf{OPT}$ for all x, and we want to leverage this fact if it can give us better guarantees. Additionally, if all values are integral in the range $\{1, \dots, v_{\max}\}$, we wish to leverage this fact as well, as it implies that an additive ε loss is also OK, as $\mathsf{OPT} \ge 1$. It turns out that both facts can indeed

[9] Note that it is possible for two disjoint ironed intervals to have the same slope.

be leveraged to obtain much stronger approximation guarantees than what are already known (essentially replacing $v_{\max}$ with $\ln(v_{\max})$ in previous bounds), stated in Theorem 4.1 below.

THEOREM 4.1. *For any $\varepsilon > 0$ and concave function $f : [0, v_{\max}] \to [0, \infty)$ such that $f(0) = 0$, $f^+(0) \leq 1$, $f^-(v_{\max}) \geq 0$*[10]*, there exists a sequence X of at most $O\left(\min\left\{1/\varepsilon, \sqrt{\frac{\log v_{\max}}{\varepsilon}}\right\}\right)$ points such that for all $x \in [0, v_{\max}]$,*

$$f(x) - \varepsilon\left(1 + f(v_{\max})\right) \leq \tilde{f}_X(x) \leq f(x).$$

The proof of Theorem 1.2 follows from Corollary 4.1 and Theorem 4.1 together with a little bit of algebra, and is presented in the full version of this paper.

Finally, we remark on some alternative terms that can be taken to replace $v_{\max}$ in Theorem 1.2. It will become clear why these replacements are valid after reading the proof of Theorem 1.2, but we will not further justify the validity of these replacements here.

- First, for instances with integral valuations, we may replace $v_{\max}$ everywhere with $\max_i r_{\geq i}$. This is essentially because we don't actually need to approximate $\tilde{R}_{\geq i}$ on the entire interval $[0, v_{\max}]$, but only the interval $[0, r_{\geq i}]$.
- We may further define $q = \max_i r_{\geq i}/\mathsf{OPT}$ for any (not necessarily integral, possibly continuous) instance, and replace $v_{\max}$ everywhere with q, even for non-integral instances. This is essentially because we only used the integrality assumption to guarantee that $\mathsf{OPT} \geq 1$.
- Finally, if $p_{\geq i}$ denotes the probability that the buyer has value at least $r_{\geq i}$ and deadline at least i, observe that $\mathsf{OPT} \geq r_{\geq i} \cdot p_{\geq i}$. So if the probability of sale at each $r_{\geq i}$ is at least p, we may observe that $q \geq 1/p$ (where q is defined as in the previous bullet) and replace $v_{\max}$ with $1/p$ everywhere.

The bullets above suggest that the "hard" instances (where some instance-specific parameter shows up in order to maintain optimal dependence on ε) are those where most of the revenue comes from very infrequent events where the buyer has an unusually high value. Due to the intricate interaction between different deadlines, these parameters can't be circumvented with simple discretization arguments, or by improved polygon approximations (provably, see Section 4.1), but it is certainly interesting to see if other arguments might allow one to replace $\log v_{\max}$ with (for example) something like $\log(n/\varepsilon)$.

4.1 A tight example for polygon approximation

It turns out that the guarantees provided by Theorem 4.1 are tight. Specifically, if no dependence on $v_{\max}$ is desired, then $1/\varepsilon$ is the best bound achievable. Also, if it's acceptable to depend on both $v_{\max}$ and ε, then the bound of $\sqrt{\frac{\log v_{\max}}{\varepsilon}}$ in Theorem 4.1 is tight. Taken together, this means that $O\left(\min\left\{1/\varepsilon, \sqrt{\frac{\log v_{\max}}{\varepsilon}}\right\}\right)$ lies at the Pareto frontier of the dependences achievable as a function of both $v_{\max}$ and ε. The examples proving tightness of these bounds are actually quite simple, and *provably the worst possible examples.*

PROPOSITION 4.1. *Let f be a concave function on $[0, v_{\max}]$, and let there be no polygon approximation of f using k segments for additive error ε. Then there exists a concave function g over $[0, v_{\max}]$ satisfying:*

- *There is no polygon approximation of g using k segments for additive error ε.*
- *$f(0) = g(0)$, $f^+(0) = g^+(0)$, $f(v_{\max}) = g(v_{\max})$, $f^-(v_{\max}) = g^-(v_{\max})$.*
- *g is piecewise-linear with $2k$ segments.*

5 Tightness of the approximation scheme

Finally, we construct an instance of the FedEx problem that is hard to approximate with small menu complexity. We try to reason similar to the example constructed in Section 3, but things are trickier here. In particular, the challenge in Section 3 was in mapping between distributions and revenue curves. But once we had the revenue curves, it was relatively straight-forward to plug through Fiat et al.'s algorithm [FGKK16] and ensure that the optimal auction had high menu complexity.

Already nailing down the behavior of an optimal auction was tricky enough, but we now have to consider *every* approximately optimal auction (almost all of which don't necessarily result from Fiat et al.'s algorithm (see, e.g. Section 4)). Indeed, one can imagine doing all sorts of strange things on any day i that are suboptimal, but might somehow avoid the gradual buildup in the i-deadline menu complexity.[11]

To cope with this, our approach has two phases: first, we characterize a restricted class of auctions that we call *clean.* At a very high level, clean auctions never

[10] We use f^+ to denote the right hand derivative and f^- to denote the left hand derivative.

[11] For example, an ε-approximate menu could set price 0 or ∞ with probability ε for shipment on any day, or something much more chaotic.

make "bizarre" choices on day i that both decrease the revenue gained on day i and *strictly* increase constraints on choices available for future days. To have an example in mind: if the revenue on day 1 is maximized by setting a price of 1, it *might* make sense to set price 2 to receive the item on day 1 instead, as this relaxes constraints on future days, and maybe this somehow helps when also constrained by menu complexity. But it makes *no sense* to instead set price 1/2: this only decreases the revenue achieved on day 1, *and* provides stricter constraints on future days (as now she has the option to get the item on day 1 at a cheaper price).

For our example, we first show that all clean auctions that maintain a good approximation ratio must have high menu complexity. We then follow up by making the claims in the previous paragraph formal: any arbitrary auction of low menu complexity can be derived by "muddling" a clean auction, a process which never increases the revenue. A little more specifically, cleaning the menu for deadline i can only increase the revenue and allow more options on later deadlines, without increasing the menu complexity. We conclude with a formal statement of our lower bound, which proves Theorem 1.3.

THEOREM 5.1. *There exists an instance of the FedEx problem such that for any mechanism that has at most $n/8$ menu options on a day $i \in (n/4, n/2]$, it has revenue at most* $\mathsf{OPT}\left(1 - \frac{1}{200000n^2}\right)$.

6 Conclusions and Future Work

We provide the first nearly-tight quantitative results on menu complexity in a multi-dimensional setting. Along the way, we design new polygon approximations for a hybrid additive-multiplicative guarantee that turns out to be just right for our application (as evidenced by the nearly-matching lower bounds obtained from the same ideas).

There remains lots of future work in the direction of menu complexity, most notably the push for tighter quantitative bounds in "truly" multi-dimensional settings, where the gaps between upper (exponential) and lower (polynomial) are vast. We believe that continuing a polygon approximation approach is likely to yield fruitful results. After all, there is a known connection between concave functions and any mechanism design setting via utility curves, and low menu complexity exactly corresponds to piece-wise linear utility curves with few segments. Still, there are two serious barriers to overcome: first, these utility curves are now multi-dimensional instead of single-dimensional revenue curves. And second, the relationship between utility curves and revenue is somewhat odd (expected revenue is equal to an integral over the support of $\vec{x} \cdot \Delta_f(\vec{x}) - f(\vec{x})$), whereas the relationship between revenue curves and revenue is more direct. There are also intriguing directions for future work along the lines of one-and-a-half dimensional mechanism design, the most pressing of which is understanding multi-bidder instances (as all existing work, including ours, is still limited to the single-bidder setting).

7 Instances with regular distributions may require randomness

For single-dimensional settings, it's well-understood that "the right" technical condition on value distributions to guarantee a simple optimal mechanism is *regularity*. This guarantees that "virtual values" are non-decreasing and removes the need for ironing, even for multi-bidder settings. Interestingly, "the right" technical condition on value distributions to guarantee a simple optimal mechanism for 1.5 dimensional settings is no longer regularity, but *decreasing marginal values*. For example, if all marginals satisfy decreasing marginal values, the optimal mechanism is deterministic for the FedEx problem [FGKK16], selling a single item to a budget-constrained buyer [CG00, DW17], and a capacity-constrained buyer [DHP17].

Still, regularity seems to buy *something* in these problems. For instance, Fiat et al. show that when there are only two possible deadlines, regularity suffices to guarantee that the optimal mechanism is deterministic. It has also been known since early work of Laffont and Robert that regularity suffices to guarantee that the optimal mechanism is deterministic when selling to a budget-constrained buyer with only one possible budget [LR96]. But the extent to which regularity guarantees simplicity remained open (and was explicitly stated as such in [DW17]). In this section, we show that regularity guarantees nothing beyond what was already known. In particular, there exists an instance of the FedEx problem with three possible deadlines where all marginals are regular but the optimal mechanism is randomized. This immediately implies an example for a budget-constrained buyer and three possible budgets as well (for instance, just set all three budgets larger than $v_{\max}$ so they will never bind).

In the full version of this paper we describe our instance of the FedEx problem where the optimal auction is randomized, despite all marginals being regular and there only being 3 possible deadlines (recall that Fiat et al. show that the optimal auction remains deterministic for regular marginals and 2 deadlines).

References

[BCKW10] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing Randomized Allocations. In *the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010.

[BGN17] Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. The menu-size complexity of revenue approximation. abs/1604.06580, 2017.

[BGO07] Jean-Daniel Boissonnat, Leonidas J. Guibas, and Steve Oudot. Learning smooth shapes by probing. *Comput. Geom. Theory Appl.*, 37(1):38–58, May 2007.

[BHR91] Rainer E. Burkard, Horst W. Hamacher, and Gnter Rote. Sandwich approximation of univariate convex functions with an application to separable convex programming. *Naval Research Logistics (NRL)*, 38(6):911–924, 1991.

[BILW14] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014.

[Bro08] E. M. Bronstein. Approximation of convex sets by polytopes. *Journal of Mathematical Sciences*, 153(6):727–762, 2008.

[CDP+14] Xi Chen, Ilias Diakonikolas, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. 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*, pages 1319–1328, 2014.

[CDW13] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding Incentives: Mechanism Design becomes Algorithm Design. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013.

[CDW16] Yang Cai, Nikhil Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. *SIGecom Exchanges*, 2016.

[CG00] Yeon-Koo Che and Ian Gale. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic theory*, 92(2):198–233, 2000.

[CHK07] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*, 2007.

[CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010.

[CMM11] Shuchi Chawla, David L. Malec, and Azarakhsh Malekian. 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*, pages 253–262, 2011.

[CR14] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing*, STOC '14, pages 243–252, New York, NY, USA, 2014. ACM.

[DDT13] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism Design via Optimal Transport. In *The 14th ACM Conference on Electronic Commerce (EC)*, 2013.

[DDT14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The Complexity of Optimal Mechanism Design. In *the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014.

[DDY16] Constantinos Daskalakis, Ilias Diakonikolas, and Mihalis Yannakakis. How good is the chord algorithm? *SIAM J. Comput.*, 45(3):811–858, 2016.

[DHN14] Shaddin Dughmi, Li Han, and Noam Nisan. Sampling and representation complexity of revenue maximization. In *Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*, pages 277–291, 2014.

[DHP16] Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 426–439, 2016.

[DHP17] Nikhil R. Devanur, Nima Hagpanah, and Christos-Alexander Psomas. Optimal multi unit mechanisms with private demands. Working paper, 2017.

[DJF16] Thodoris Lykouris Karthik Sridharan Eva Tardos Dylan J. Foster, Zhiyuan Li. Learning in games: robustness of fast convergence. In *the 30th annual conference on neural information processing systems (NIPS)*, 2016.

[DW17] Nikhil R. Devanur and S. Matthew Weinberg. The optimal mechanism for selling to a budget-constrained buyer: the general case. Working paper, 2017.

[FGKK16] Amos Fiat, Kira Goldner, Anna R. Karlin, and Elias Koutsoupias. The fedex problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 21–22, New York, NY, USA, 2016. ACM.

[GG09] S. Glasauer and P. Gruber. Asymptotic estimates for best and stepwise approximation of convex bodies iii. *Forum Mathematicum*, 9(9):383–404, 2009.

[GN17] Yannai A. Gonczarowski and Noam Nisan. Efficient empirical revenue maximization in single-parameter auction environments. abs/1604.06580, 2017.

[HMR15] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 45–60, New York, NY, USA, 2015. ACM.

[HN12] Sergiu Hart and Noam Nisan. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012.

[HN13] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *the 14th ACM Conference on Electronic Commerce (EC)*, 2013.

[HR15] Sergiu Hart and Philip J. Reny. Maximizing Revenue with Multiple Goods: Nonmonotonicity and Other Observations. *Theoretical Economics*, 10(3):893–922, 2015.

[LR96] Jean-Jacques Laffont and Jacques Robert. Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186, 1996.

[MR15] Jamie Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 136–144, 2015.

[MR16] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, pages 1298–1318, 2016.

[Rot92] G. Rote. The convergence rate of the sandwich algorithm for approximating convex functions. *Computing*, 48(3-4):337–361, March 1992.

[Rub16] Aviad Rubinstein. Settling the complexity of computing approximate two-player nash equillibria. In *FOCS*, 2016.

[RW15] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the 16th ACM Conference on Electronic Commerce*, 2015.

[RZ83] John Riley and Richard Zeckhauser. Optimal Selling Strategies: When to Haggle, When to Hold Firm. *Quarterly J. Economics*, 98(2):267–289, 1983.

[SBN17] Zhiyi Huang Sebastien Bubeck, Nikhil R. Devanur and Rad Niazadehd. Online auction sand multi-scale online learning. In *Proceedings of the eighteenth ACM conference on Economics and Computation (EC)*, 2017.

[Tha04] John Thanassoulis. Haggling over substitutes. *Journal of Economic Theory*, 117:217–245, 2004.

[WT14] Zihe Wang and Pingzhong Tang. Optimal mechanisms with simple menus. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 227–240, New York, NY, USA, 2014. ACM.

[YG97] XQ Yang and CJ Goh. A method for convex curve approximation. *European Journal of Operational Research*, 97(1):205–212, 1997.