

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

Raghuvansh R. Saxena^{*} Ariel Schwartzman[†] 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 $\text{poly}(n)$ 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

¹This 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.

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

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 $\varepsilon \rightarrow 0$, 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 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

²Actually our bounds can 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.

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

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.

⁴Interestingly (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].

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.

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 multivariate 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 $(1-\varepsilon)$ 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

2 Preliminaries

We consider a single bidder whose 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].

In Appendix A, we provide 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).

⁵Specifically, [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.

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 ($|\{p \mid \exists v, \pi(v, i) = p\}|$). 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 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 $v_1 < \dots < v_k$, and $a_i(1), \dots, a_i(k) \geq 0$ 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 $< i$.*

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

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

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 (a more detailed description appears in Appendix A, but the description below will suffice for our paper). 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 $(\pi_0, p_0), \dots, (\pi_k, p_k)$ (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 $p_1 < \dots < p_k$, 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 2.2 ([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 $(R_{\geq 1}, \dots, R_{\geq n})$ 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, but most details are in Appendix B.

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

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

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

We include in Figure 2 a picture of the generated revenue curves for $n = 4$. As a result of this construction, we see that $R_{\geq i}$ has 2^{i-2} ironed intervals, whose endpoints themselves lie in ironed intervals of $R_{\geq i+1}$. 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 $i + 1$. The proof of Theorem 3.2 (which implies Theorem 1.1) formalizes this.

Theorem 3.2. *The optimal mechanism for any instance satisfying the conditions of Lemma 3.1 has i -deadline complexity 2^{i-1} for all i , and menu complexity $2^n - 1$.*

The proof of Theorem 3.2 is presented at the end of Section B of the Appendix.

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. Most proofs appear in Appendix C, but we overview our approach here.

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 $f(x) - \tilde{f}_X(x) \geq 0$. Typically, for a ‘good’ polygon approximation, one requires for $\varepsilon > 0$, that $f(x) - \varepsilon \leq \tilde{f}_X(x) \leq f(x)$.

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 $\tilde{R}_{\geq i}$. 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 $\tilde{R}_{\geq i}(\cdot)$ and the polygon approximation obtained to $\tilde{R}_{\geq i}(\cdot)$ (this isn’t obvious, but not hard. See Appendix C).

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

- 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](#) of [Proposition C.7](#) describes the reduction formally, but all related proofs are in [Appendix C](#).

Corollary 4.1. *Consider a FedEx instance with n deadlines. For all $i \in \{1, 2, \dots, n\}$, let g_i be the function $\tilde{R}_{\geq i}$ defined in [2.3](#), and let X_i be a sequence of k_i points in $[0, r_{\geq i}]$ such that for all $x \leq r_{\geq i}$, we have $g_i(x) - \varepsilon_i \leq \tilde{g}_{i, X_i}(x) \leq g_i(x)$. 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 $\text{OPT} - \sum_{i=1}^n \varepsilon_i$.*

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

At this point, it seems like the right approach is to just set each $\varepsilon_i = \varepsilon \cdot \text{OPT}/n$ 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: $\tilde{R}(x) \leq \text{OPT}$ 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 $\text{OPT} \geq 1$. It turns out that both facts can indeed 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 4.2](#) below.

Theorem 4.2. *For any $\varepsilon > 0$ and concave function $f : [0, v_{\max}] \rightarrow [0, \infty)$ 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 $x \in [0, v_{\max}]$,*

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

The proof of [Theorem 1.2](#) follows from [Corollary 4.1](#) and [Theorem 4.2](#) together with a little bit of algebra, and is deferred until [Appendix C](#).

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 $q = \max_i r_{\geq i}/\text{OPT}$ 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 $\text{OPT} \geq 1$.
- 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 $\text{OPT} \geq r_{\geq i} \cdot p_{\geq i}$. 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 $\log(n/\varepsilon)$.

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

4.1 A tight example for polygon approximation

It turns out that the guarantees provided by [Theorem 4.2](#) 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 4.2](#) 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* (proof of the below claim appears in [Appendix C](#))

Proposition 4.3. *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 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. Formal definitions and claims related to this appear in [Appendix D](#). We conclude with a formal statement of our lower bound, which proves [Theorem 1.3](#).

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

Theorem 5.1. *Any mechanism for the FedEx instance described in [subsection D.1](#) that has at most $n/8$ menu options on a day $i \in (n/4, n/2]$ has revenue at most $\text{OPT} \left(1 - \frac{1}{200000n^2}\right)$.*

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).

We proceed by describing now 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). Throughout this section, instead of using revenue curves $R(v)$, we use $\Gamma(v) = -R(v)$. This is in accordance to [[FGKK16](#)].

7.1 The setting

Consider bidders with types distributed as F_1 on day 1, F_2 on day 2, and F_3 on day 3.

$$\begin{aligned} F_1(v) &= 1 - e^{-v}, & f_1(v) &= e^{-v}, \\ F_2(v) &= 1 - e^{-5v}, & f_2(v) &= 5e^{-5v}, \\ F_3(v) &= 1 - e^{-\frac{v}{5}}, & f_3(v) &= \frac{1}{5}e^{-\frac{v}{5}}. \end{aligned}$$

The distribution over days q is

$$q_1 = \frac{10}{21} \quad q_2 = \frac{10}{21} \quad q_3 = \frac{1}{21}.$$

Note that the three distributions are regular but don't have decreasing marginal revenues.

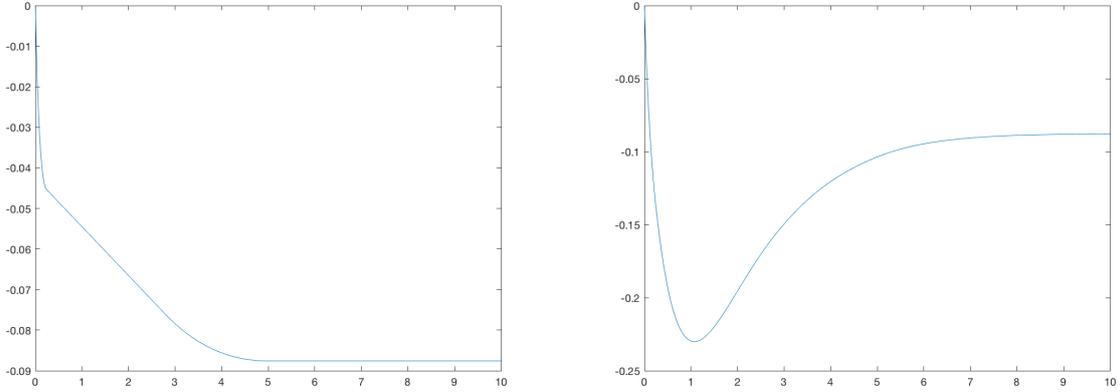


Figure 1: (Left) The curve $\hat{\Gamma}_{\geq 2}$ in our example. (Right) The curve $\Gamma_{\geq 1}$. Note that this curve is minimized at a point in the ironed interval of $\hat{\Gamma}_{\geq 2}$.

7.2 Analysis

We use the iterative procedure described in [FGKK16] to find the optimal auction.

$$\begin{aligned} \Gamma_{\geq 3} &= -\frac{v}{21}(1 - F_3(v)) = -\frac{v}{21}e^{-\frac{v}{5}}, \\ r_{\geq 3} &= 5. \\ \Gamma_{\geq 2} &= \begin{cases} \Gamma_2(v) + \hat{\Gamma}_{\geq 3}(v), & 0 \leq v \leq 5 \\ \Gamma_2(v) - \frac{5}{21e}, & 5 \leq v \end{cases} = \begin{cases} -\frac{10v}{21}e^{-5v} - \frac{v}{21}e^{-\frac{v}{5}}, & 0 \leq v \leq 5 \\ -\frac{10v}{21}e^{-5v} - \frac{5}{21e}, & 5 \leq v \end{cases} \\ r_{\geq 2} &\approx 5.0\dots \\ \Gamma_{\geq 1} &= \begin{cases} \Gamma_1(v) + \hat{\Gamma}_{\geq 2}(v), & 0 \leq v \leq 5.0\dots \\ \Gamma_1(v) - \frac{5}{21e}, & 5.0\dots \leq v \end{cases} = \begin{cases} -\frac{10v}{21}e^{-v} - \frac{10v}{21}e^{-5v} - \frac{v}{21}e^{-\frac{v}{5}}, & 0 \leq v \leq 0.245\dots \\ -\frac{10v}{21}e^{-v} - \mathbf{A \text{ line. See Figure 1}}, & 0.245\dots \leq v \leq 2.79\dots \\ -\frac{10v}{21}e^{-v} - \frac{10v}{21}e^{-5v} - \frac{v}{21}e^{-\frac{v}{5}}, & 2.79\dots \leq v \leq 5.0\dots \\ -\frac{10v}{21}e^{-v} - 0.0876\dots, & 5.0\dots \leq v \end{cases} \\ r_{\geq 1} &\approx 1.07. \end{aligned}$$

All real values written are approximate and end with “...”. The case in **boldface** denotes the ironed interval. The optimal price for day 1 is 1.07 which is in that interval. Hence the optimal auction has to be randomized.

Appendices

A Additional preliminaries

In this appendix, we summarize the approach in [FGKK16] to obtain the optimal mechanism for an arbitrary “FedEx” instance. We begin with the linear program that encodes this optimization problem.

$$\begin{aligned}
& \text{maximize} && \sum_i \sum_v q_i f_i(v) p(v, i) \\
& \text{subject to} && \pi(v, i)v - p(v, i) \geq \pi(v-1, i)v - p(v-1, i) && \forall v \geq 1, i && \text{(Leftwards IC)} \\
& && \pi(v, i)v - p(v, i) \geq \pi(v+1, i)v - p(v+1, i) && \forall v < v_{max}, i && \text{(Rightwards IC)} \\
& && \pi(v, i)v - p(v, i) \geq \pi(v, i-1)v - p(v, i-1) && \forall i > 1, v && \text{(Downwards IC)} \\
& && \pi(v, i) \leq 1 && \forall i, v && \text{(Feasibility)} \\
& && \pi(0, 1) = p(0, 1) = 0. && && \text{(Individual Rationality)}
\end{aligned} \tag{1}$$

Note that we have not included the constraints where the bidder misreports a higher deadline. No rational bidder would consider these deviations since they would always get non-positive utility. We now formally present the allocation curves described in Section 2. This, combined with the definitions of optimal revenue curves, provide a clean characterization of optimal auctions for any instance of the FedEx problem.

Definition A.1 (Optimal allocation curves [FGKK16]). *Let j^* be the largest j such that $v_j \leq r_{\geq i}$. For any $j \leq j^*$, consider two cases:*

- $\hat{R}_{\geq i}$ is not ironed at v_j . Then

$$a_{i,j}(v) = \begin{cases} 0 & v < v_j \\ 1 & \text{else.} \end{cases}$$

- $\hat{R}_{\geq i}$ is ironed at v_j . Let \underline{v}_j be the largest $v < v_j$ such that $R_{\geq i}$ is not ironed at v and \bar{v}_j be the smallest $v_j < v$ such that $R_{\geq i}$ is not ironed at v . Let $0 < \delta < 1$ be such that

$$v_j = \delta \underline{v}_j + (1 - \delta) \bar{v}_j.$$

Define

$$a_{i,j}(v) = \begin{cases} 0 & v < \underline{v}_j \\ \delta & \underline{v}_j \leq v \leq \bar{v}_j \\ 1 & v > \bar{v}_j. \end{cases}$$

Then set $a_i(v)$ as follows

$$a_i(v) = \begin{cases} \sum_{j=1}^{j^*} (\beta_j - \beta_{j-1}) a_{i,j}(v) & v < r_{\geq i} \\ 1 & v \geq r_{\geq i}, \end{cases}$$

where β_j is the probability of allocating the item with value a_j on day $i-1$.

Lemma A.1. [FGKK16] *The allocation curves a_i are monotone increasing from 0 to 1 and satisfy the Downwards IC from 1. Moreover, each a_i changes value only at points where $\hat{R}_{\geq i}$ is not ironed.*

Remark A.1. *Every solution to LP 1 is an optimal mechanism for the FedEx problem.*

In addition, we present this simple claim, which is unrelated to the FedEx problem itself, that will be useful in future sections.

Claim A.1. *Let f be a distribution over $[v_{\max}]$, and let ϕ_f, R_f be the corresponding virtual value and revenue curve, respectively. Then, for all $v \in [v_{\max} - 1]$*

$$R_f(v) - R_f(v - 1) = \phi_f(v)f(v).$$

Proof. This follows from some simple calculations:

$$\begin{aligned} R_f(v) - R_f(v - 1) &= v(1 - F(v - 1)) - (v + 1)(1 - F(v)) \\ &= v(F(v) - F(v - 1)) - (1 - F(v)) \\ &= vf(v) - (1 - F(v)) = f(v)\phi_f(v). \end{aligned}$$

□

B Omitted proofs and perturbed example of Section 3

B.1 Omitted proofs

Lemma B.1. *Fix $v > 0$. Let $(2v)^{-1} > \varepsilon > 0$ and let $R = \{r_i\}_{i=1}^v$ be a sequence of real numbers such that $|r_{i+1} - r_i| \leq \varepsilon, \forall i \in [v - 1]$. If $r_1 = 1$, there exists a distribution $f : [v] \rightarrow [0, 1]$ such that $R_f(i) = r_i, \forall i \in [v]$.*

Proof. By our choice of ε , all elements in the sequence are greater than $1/2$. Let $r_{v+1} = 0$ and define

$$F(i) = 1 - \frac{r_{i+1}}{i+1}, \forall i \in \{0, 1, 2, \dots, v\}.$$

We will show that F indeed corresponds to a valid distribution function by showing that it is non-negative and non-decreasing. Once we have shown this, it becomes clear that $R_f(i) = r_i$, proving the Lemma.

Claim B.1. *$F(i)$ is non-negative, non-decreasing.*

Proof. For $i = 0$, $F(i) = 1 - 1 = 0$. Suppose $i \geq 1$. Since $0 \leq r_i < 2, \forall i$, it follows that $1 \geq F(i) \geq 1 - \frac{r_{i+1}}{2} > 0$. To show it is non-decreasing, consider the difference between two consecutive terms. For all $v > i \geq 1$,

$$\begin{aligned} F(i) - F(i - 1) &= \frac{r_i}{i} - \frac{r_{i+1}}{i+1} = \frac{r_i}{i(i+1)} + \frac{r_i - r_{i+1}}{i+1} \\ &\geq \frac{1}{i+1} \left(\frac{r_i}{i} - \varepsilon \right) \\ &> \frac{1}{i+1} \left(\frac{r_i}{v} - \frac{1}{2v} \right) \geq 0. \end{aligned}$$

For $i = v$, $F(i) - F(i - 1) = \frac{r_i}{i} \geq 0$.

□

□

We are now ready to explicitly construct an example that achieves deadline menu complexity of 2^{n-1} where $v_{\max} = 2^n - 1$. Fix n and construct $n - 1$ sequences $S_i = \{s_{ij}\}_{j=1}^{2^n-1}$ of length $2^n - 1$, then

$$s_{ij} = \varepsilon \cdot (1 - 2b_{i+1}(j))b_i(j) \left(\prod_{k=1}^{i-1} (1 - b_k(j)) \right),$$

where $b_i(j)$ denotes the i -th least significant bit in the binary expansion of j , $\varepsilon = 4^{-n}$. Let $S_n = \{s_{nj}\}_{j=1}^{2^n-1}$ as $s_{nj} = \varepsilon(2b_n(j) - 1)$. Lemma B.1 implies that there exist distributions $f_j \forall j \in [n]$ such that $\phi_{f_j} f(j) = s_i \forall i \in [v - 1]$. In our construction the type distribution of day i corresponds to the distribution f_{n+1-i} . Let the distribution over days be uniform (i.e. $q_i = \frac{1}{n} \forall i \in [n]$). In order to show that this construction achieves a menu complexity of 2^{n-1} we first need to characterize the revenue curves for all days, and then show that an optimal auction exists where prices can be set so as to create a large menu. The intuition for these revenue curves is that their ironed intervals are nested: prices at the endpoints of ironed intervals on day i are the midpoints of new ironed intervals on day $i + 1$. In addition, after ironing the curves, they look like constants. Therefore, the optimal revenue curves will have the same ironed intervals as their original counterparts. We design the revenue of the first day to be maximized at the median value. On the next day this price will belong to an ironed interval, meaning that any optimal auction must offer a lottery over two prices that are not ironed for day 2. The size of the lottery offered directly translate into the minimum number of options a menu for that day must have. By the nesting construction, this will double the number of options offered on day $i + 1$ with respect to day i .

We can now use the above construction, combined with Claim A.1 and Lemma B.1, to show a more general result, Lemma 3.1.

Proof of Lemma 3.1. By simple examination, the revenue curve for day 1 is just a line that increases until it reaches $v_{\max}/2$ and then decreases, so it is maxed at $v_{\max}/2$. Consider day $i \geq 2$. The sequence that generates its revenue curve, $s_{n+1-i,j}$ is non-zero iff j has remainder $2^{n-i} \pmod{2^{n-i+1}}$. Since there are $2^n - 1$ values in the sequence, there will be 2^{i-1} non-zero values. These values will alternate between ε , $-\varepsilon$ and each alternation creates one ironed interval: the revenue decreases by ε , stays at that value for a while, and increases by ε again. This gives 2^{i-2} ironed intervals for the revenue curve of day i . Moreover, the last price where the revenue increases is at $2^i(2^{n-i+1} - 1)$, making it a valid maximizer. The revenue remains constant from there on so any higher value is also a maximizer.

For the third point, note that R_i takes values in $\{1, 1 - \varepsilon\}$, alternating between them for intervals whose lengths depend on i . \hat{R}_i , the upper concave hull of R_i , is then a constant function with value 1 everywhere.

We prove the last point by induction, starting from R_n . This holds because $R_n = R_{\geq n}$. Note that, from our previous claim, $\hat{R}_{\geq n} = \hat{R}_n$ is a constant function. Suppose that $R_{\geq i+1}(v) = R_{i+1}(v) + c$ for some c . For $v \leq r_{\geq i+1}$, $R_{\geq i}(v) = R_i(v) + \hat{R}_{\geq i+1}(v)$. For $v > r_{\geq i+1}$, $R_{\geq i}(v) = R_i(v) + \hat{R}_{\geq i+1}(r_{\geq i+1})$. In either case, the term added to $R_i(v)$ is a constant (and it is the same constant) by the inductive hypothesis, so the claim follows. \square

Lemma B.2. *The series of optimal revenue curves $R_{\geq i}$ induced by the distributions f_{n+1-i} is such that on any day i , an optimal allocation curve as constructed by [FGKK16] is a step function with 2^{i-1} jumps and takes the following form:*

$$a_i(v) = \begin{cases} 0 & v < 2^{n-i}, \\ \frac{k+1}{2^{i-1}} & 2^{n-i} + k2^{n-i+1} \leq v < 2^{n-i} + (k+1)2^{n-i+1}, \\ 1 & v \geq 2^n - 2^{n-i}, \end{cases}$$

where $0 \leq k \leq 2^{i-1} - 2$. Moreover, these prices where the function jumps on day i will belong to ironed intervals for the optimal revenue curve of the following day, $R_{\geq i+1}$ for all $i < n$.

Proof. We prove this also by induction going from the first day to the last. As stated before, $R_{\geq 1}$ is maximized at $v_{\max}/2$ and has no ironed intervals, therefore it is clear $r_{\geq 1} = v_{\max}/2$ and the optimal allocation curve a_1 corresponds to a step function at $v_{\max}/2$. By construction, this price belongs to an ironed interval of $R_{\geq 2}$. Thus the base case holds. Suppose that the statement is true for day i . We want to understand the optimal allocation curve for day $i + 1$.

First note that the places where the function jumps on day i , $P_i = \{2^{n-i} + k2^{n-i+1}\}_{k=0}^{2^{i-1}-1}$, belong to ironed intervals of $R_{\geq i+1}$ (which are the same intervals as R_{i+1}). This is because at these prices $j \in P_i$ the function $s_{n-i+1,j}$ takes a value of 0. The nearest place where they are non-zero are exactly at $j - 2^{n-i-1}$ and $j + 2^{n-i-1}$, where the function takes values ε and $-\varepsilon$, meaning we are inside an iron interval (the revenue decreases and then increase by ε). Thus by the optimal allocation curves suggested in [FGKK16] we observe that if on day i we offered a price of $2^{n-i} + k2^{n-i+1}$ with positive probability $\frac{k+1}{2^{i-1}}$, we must also allocate at prices $2^{n-i-1} + (2k)2^{n-i}$ and $2^{n-i-1} + (2k+1)2^{n-i}$ on day $i + 1$ with positive probability.

The probability for allocating the item with price $2^{n-i-1} + k2^{n-i} \leq v < 2^{n-i-1} + (k+1)2^{n-i}$ depends on the parity of k . If k is odd then the values in this range correspond to a non-ironed interval on day $i + 1$, meaning they preserve the probability of allocation from day i . The probability of allocating on the interval with endpoints $2^{n-i} + \frac{(k-1)2^{n-i+1}}{2}$ and $2^{n-i} + \frac{k+1}{2}2^{n-i+1}$, which contains our new interval of interest, is $\frac{k+1}{2^{i+1}}$. If k is even then we belong to an ironed interval on day $i + 1$, meaning that the probability of allocation is going to be the average of allocating on the two intervals on day i that intersect this one. These intervals are $[2^{n-i} + \frac{k-2}{2}2^{n-i+1}, 2^{n-i} + \frac{k}{2}2^{n-i+1}]$ and $[2^{n-i} + \frac{k}{2}2^{n-i+1}, 2^{n-i} + \frac{(k+2)}{2}2^{n-i+1}]$. Thus the probability of allocating at v on day $i + 1$ is just $\frac{k+1}{2^{i+1}}$. \square

With this understanding of the revenue curves of the instance we construct, we are ready to prove Theorem 3.2, the main result of Section 3.

Proof of Theorem 3.2. By Lemma B.2, the optimal allocation curves for day i are step-wise increasing functions with 2^{i-1} steps. This means that for any deadline i , the i -deadline menu complexity of this instance is 2^{i-1} . Combining the number of options offered over all n days gives a menu complexity of $2^n - 1$. \square

B.2 Perturbed case

In this appendix, we show how to tweak the example in Section 3 so that *no* optimal auction has a menu complexity less than 2^n . The problem with the example in Section 3 is that we don't have to follow the allocation suggestion of [FGKK16] in order to achieve an optimal auction because we could simply choose larger ironed intervals on every day that spanned the whole spectrum of prices and (because of the simplicity of the construction) still recover all the revenue. We add a small non-linear term to the revenue curve of each day to dissuade from this while still preserving the 'nested' structure of ironed intervals. Consider now the $n - 1$ sequences

$$s_{ij} = \frac{\varepsilon}{2} \cdot (1 - 2b_{i+1}(j))b_i(j) \left(\prod_{k=1}^{i-1} (1 - b_k(j)) \right) + \delta j \cdot (2v_{\max} - j),$$

where $v_{\max} = 2^n$, $\varepsilon = \frac{1}{4v}$ and δ , the weight of the non-linear term, is v^{-10} . The distribution for the first day is the same as the one we used for the previous case. Note that the non-linear term added is maxed at v_{\max} . Again, Lemma B.1, Claim A.1 allow us to conclude there are revenue curves R_i and valid distributions f_{n+1-i} , whose changes are dictated by the sequence s_{n+1-i} . We restate the results from 3 with the appropriate adjustments.

Lemma B.3. *The series of revenue curves R_i induced by the distributions f_{n+1-i} satisfy the following properties:*

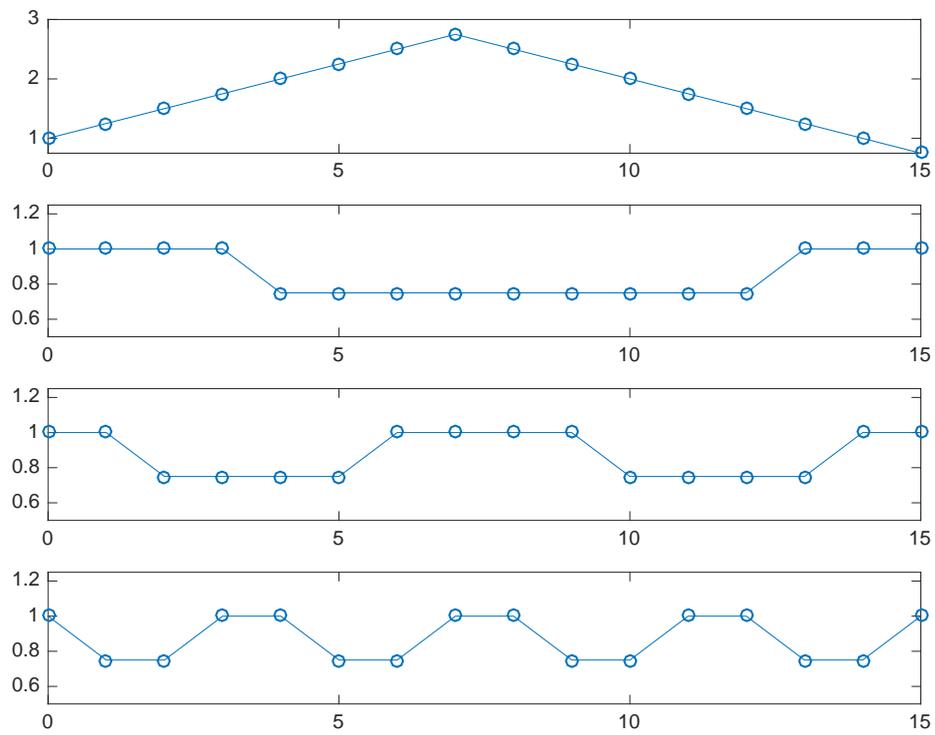


Figure 2: Revenue curves corresponding to an instantiation of our construction with $n = 4$ and $\varepsilon = .25$.

- R_1 is maximized at $v_{\max}/2$ and has no ironed intervals.
- R_i for $1 < i \leq d$ has a unique maximizer at price v_{\max} and has 2^{i-2} ironed intervals.
- $R_{\geq i}$ has the same ironed intervals as R_i .

Proof. The first point remains true since we haven't changed R_1 . Since v_{\max} was a maximizer of the function before adding the non-linear term, it will remain a maximizer since it's also optimal for that function. In fact, it is now the *unique* maximizer. This implies that $r_{\geq n} = v_{\max}$. The last point follows from a similar inductive proof to the one in the original case. It is easy to see that now $R_{\geq i}$ will be maximized at v_{\max} , for all i . Then by the characterization from [FGKK16] we have that for any $v < v_{\max}$, $R_{\geq i}(v) = R_i(v)$. \square

The key property of Lemma B.2 is that the ironed intervals for day i are nested in those of day $i + 1$. This property is kept despite the small adjustment of the function, so the statement of the Lemma remains and the proof of the main theorem of Section 3 follows. After the adjustment the ironed intervals become unique and no clever algorithm can take large ironed intervals to avoid a large menu complexity.

C Omitted proofs of Section 4

C.1 Polygon approximation of concave functions

Polygon approximation of functions is a classic problem in approximation theory. The central problem in approximation theory is to see how 'well' does a class of 'simple' functions approximate an arbitrary function. Usually, the 'simple' functions have a property that makes them easy to study and the approximation scheme ensures that results carry over to arbitrary functions. In the subfield of polygon approximation, the class of 'simple' functions is the set of all continuous piecewise linear functions. Different error metrics (an error metric defines the properties of a good approximation) are studied, the most common being an additive error defined by a parameter ε . We define polygon approximation formally and state a celebrated result for additive ε error.

Definition C.1 (Polygon approximation). *Given a function $f : [a, b] \rightarrow \mathbb{R}$ and a sequence X of points $a = x_0, x_1, \dots, x_{n-1}, x_n = b$, define \tilde{f}_X , the X -approximation of f , to be the continuous piecewise linear function that passes through $(x_0, f(x_0)), (x_1, f(x_1)), \dots, (x_n, f(x_n))$.*

Before continuing, we briefly state the notation that we are going to use throughout this section. For a concave function f , we use f^+ to denote the right-derivative and f^- to denote the left derivative. We remind the reader that for a concave function over an interval, both f^- and f^+ are well defined at every point (and there are only countably many points for which f is not differentiable). The following theorem has been taken from [Rot92] and has been reworded to fit our notation.

Theorem C.1 ([Rot92, BHR91]). *For any $\varepsilon > 0$ and concave $f : [a, b] \rightarrow \mathbb{R}$ such that $b - a = L$ and $f^+(a) - f^-(b) \leq \Delta$, there exists a sequence X of at most $3 + \sqrt{\frac{9L\Delta}{8\varepsilon}}$ points such that $\forall x \in [a, b], 0 \leq f(x) - \tilde{f}_X(x) \leq \varepsilon$.*

Theorem C.1, as stated above, deals with arbitrary concave functions and is known to be tight [Rot92]. In our setting of the FedEx problem, the concave functions that we wish to approximate are monotone revenue curves defined over $[0, r_{\geq i}]$ (for the rest of this section, we will use v_{\max} instead of $r_{\geq i}$ to refer to the right end-point of the interval to keep the notation non-specific to the FedEx problem). Being monotone revenue curves, they satisfy many other properties, e.g., $f(0) = 0$, $f^+(x) \leq 1$, and $f^-(x) \geq 0$. Also, the error that we are interested in

is *not* the additive ε error that is dealt with in the theorem. **Lemma C.3** formalizes the exact guarantees we desire. Note that the error metric that we use is an additive multiplicative hybrid that increases with $f(v_{\max})$. To prove this lemma, we use **Theorem C.1** separately on each of the intervals $[0, 1], [1, 2], \dots, [v_{\max}/4, v_{\max}/2], [v_{\max}/2, v_{\max}]$. We show that when the size of the approximating sequence X is large in these intervals, then $f(v_{\max})$ must be large. In turn, this implies that we have more wiggle room in our error. We exploit this room to improve the dependence on v_{\max} to logarithmic from the polynomial dependence on L in **Theorem C.1**.

C.2 Proof of Theorem 4.2

We begin with an easy lemma: **Lemma C.2** shows how to lower bound $f(v_{\max})$ using the derivatives of various points along the way.

Lemma C.2. *For any concave $f : [0, v_{\max}] \rightarrow [0, \infty)$ such that $f^-(v_{\max}) \geq 0$, we have*

$$f(v_{\max}) \geq \sum_{i=0}^{\lceil \log v_{\max} \rceil} f^+ \left(\frac{v_{\max}}{2^i} \right) \frac{v_{\max}}{2^{i+1}}.$$

Proof. Since the function f is concave, we get that for all $i > 0$,

$$f \left(\frac{v_{\max}}{2^{i-1}} \right) - f \left(\frac{v_{\max}}{2^i} \right) \geq f^+ \left(\frac{v_{\max}}{2^{i-1}} \right) \frac{v_{\max}}{2^i}.$$

Adding these inequalities for $i = 1, 2, \dots, \lceil \log v_{\max} \rceil + 1$, we get

$$f(v_{\max}) \geq \sum_{i=1}^{\lceil \log v_{\max} \rceil + 1} f^+ \left(\frac{v_{\max}}{2^{i-1}} \right) \frac{v_{\max}}{2^i} \geq \sum_{i=0}^{\lceil \log v_{\max} \rceil} f^+ \left(\frac{v_{\max}}{2^i} \right) \frac{v_{\max}}{2^{i+1}}.$$

□

And now, we show how to improve the depends on v_{\max} from linear to logarithmic for our relaxed hybrid guarantee.

Lemma C.3. *For any $\varepsilon > 0$ and concave function $f : [0, v_{\max}] \rightarrow [0, \infty)$ such that $f(0) = 0$, $f^+(0) \leq 1$, $f^-(v_{\max}) \geq 0$, there exists a sequence X of at most $O \left(\log v_{\max} + \sqrt{\frac{\log v_{\max}}{\varepsilon}} \right)$ points such that for all $x \in [0, v_{\max}]$,*

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

Proof. We divide the domain of the function f into the intervals $I_i = \left[\frac{v_{\max}}{2^i}, \frac{v_{\max}}{2^{i-1}} \right]$ for $i \in \{1, 2, \dots, \lceil \log v_{\max} \rceil\}$ and the interval $I_0 = \left[0, \frac{v_{\max}}{2^{\lceil \log v_{\max} \rceil}} \right]$. For each interval I_i , we define a sequence X_i that approximates f well over this interval. The final sequence X is obtained by combining all the sequences X_i .

We start with the interval I_0 . The length of I_0 is $\frac{v_{\max}}{2^{\lceil \log v_{\max} \rceil}} \leq 1$. Since $f^+(0) \leq 1$, we can apply **Theorem C.1** with $\Delta = 1$ and conclude that to approximate f up to an additive $\varepsilon \leq \varepsilon(1 + f(v_{\max}))$ error on this interval, we need a sequence X_0 of at most $3 + \sqrt{\frac{9}{8\varepsilon}}$ points.

We now argue for the rest of the domain of f . The length L_i of the interval I_i is $\frac{v_{\max}}{2^i}$. Let $\Delta_i = f^+ \left(\frac{v_{\max}}{2^i} \right) - f^- \left(\frac{v_{\max}}{2^{i-1}} \right)$. Applying **Theorem C.1** on I_i with the error parameter $\varepsilon f(v_{\max}) \leq \varepsilon(1 + f(v_{\max}))$ gives a sequence X_i at most $3 + \sqrt{\frac{9}{8} \frac{v_{\max} \Delta_i}{2^i \varepsilon f(v_{\max})}}$ points such that \tilde{f}_{X_i} approximates f on I_i up to an additive $\varepsilon f(v_{\max})$ error.

Combine all the sequences X_i for $i = 0, 1, \dots, \lceil \log v_{\max} \rceil$ to get a single sequence X such that \tilde{f}_X approximates f up to an additive $\varepsilon(1 + f(v_{\max}))$ throughout its domain $[0, v_{\max}]$. The number of points in X is at most

$$\begin{aligned}
& 3 + \sqrt{\frac{9}{8\varepsilon}} + \sum_{i=1}^{\lceil \log v_{\max} \rceil} 3 + \sqrt{\frac{9}{8} \frac{v_{\max} \Delta_i}{2^i \varepsilon f(v_{\max})}} \\
& \leq 3(1 + \lceil \log v_{\max} \rceil) + \sqrt{\frac{9}{8\varepsilon}} + \frac{1}{\sqrt{\varepsilon f(v_{\max})}} \sum_{i=1}^{\lceil \log v_{\max} \rceil} \sqrt{\frac{9}{8} \frac{v_{\max} \Delta_i}{2^i}} \\
& \leq 3(1 + \lceil \log v_{\max} \rceil) + \sqrt{\frac{9}{8\varepsilon}} + \frac{1}{\sqrt{\varepsilon f(v_{\max})}} \sqrt{\sum_{i=1}^{\lceil \log v_{\max} \rceil} \frac{9}{8} \frac{v_{\max} \Delta_i}{2^i}} \quad (\text{Cauchy-Schwarz}) \\
& \leq 3(1 + \lceil \log v_{\max} \rceil) + \sqrt{\frac{9}{8\varepsilon}} + \frac{1}{\sqrt{\varepsilon f(v_{\max})}} \sqrt{\sum_{i=1}^{\lceil \log v_{\max} \rceil} \frac{9}{8} \frac{v_{\max} f^+\left(\frac{v_{\max}}{2^i}\right)}{2^i}} \quad (\text{Definition of } \Delta_i) \\
& \leq 3(1 + \lceil \log v_{\max} \rceil) + \sqrt{\frac{9}{8\varepsilon}} + \sqrt{\frac{18}{8} \frac{\lceil \log v_{\max} \rceil}{\varepsilon}} \quad (\text{Lemma C.2}) \\
& = O\left(\log v_{\max} + \sqrt{\frac{\log v_{\max}}{\varepsilon}}\right).
\end{aligned}$$

□

It is noteworthy that [Lemma C.3](#) preserves the quadratic dependence on ε in [Theorem C.1](#), and simply replaces v_{\max} with $\ln(v_{\max})$. Note that essentially the main idea is that we wish to target additive error $\varepsilon \cdot f(v_{\max})$ in the domain of “large x ” (as targeting an additive ε in this range requires linear dependence on v_{\max} that can be avoided if $f(v_{\max})$ is big) but additive error ε in the domain of “small x ” (as targeting an additive ε requires linear dependence on v_{\max} which is small). So in order to get the improved guarantee, our bound needs the flexibility to take either the additive or multiplicative approximation in the different ranges.

To complete the proof of [Theorem 4.2](#), we additionally provide a much simpler argument showing that the dependence on v_{\max} can be removed at the cost of suboptimal dependence on ε .

Lemma C.4. *For any $\varepsilon > 0$ and any increasing concave function $f : [0, v_{\max}] \rightarrow [0, \infty)$ such that $f(0) = 0$, there exists a sequence X of at most $O(1/\varepsilon)$ points such that for all $x \in [0, v_{\max}]$,*

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

Proof. We divide the range $[0, f_{\max}]$ of the function f into $1/\varepsilon$ intervals and take one point from each interval. Formally, let $x_i = \min_x \{x \mid f(x) \geq i\varepsilon f(v_{\max})\}$, for all $i \in \{0, \dots, \lfloor 1/\varepsilon \rfloor\}$. Note that since f is continuous (it is concave), $f(x_i) = i\varepsilon f(v_{\max})$.

We define the sequence X to be the points $\{x_i\}_{i=0}^{\lfloor 1/\varepsilon \rfloor}$ followed by $x_{\lfloor 1/\varepsilon \rfloor + 1} = v_{\max}$. We prove that \tilde{f}_X is the required polygon approximation. For any point $x \in [0, v_{\max}]$, let $i \in \{0, 1, \dots, \lfloor 1/\varepsilon \rfloor\}$ be such that $x_i \leq x < x_{i+1}$. Since f is increasing, we have $f(x_i) \leq f(x) \leq f(x_{i+1})$. This gives $i\varepsilon f(v_{\max}) \leq f(x) \leq (i+1)\varepsilon f(v_{\max})$.

Since $\tilde{f}_X(x)$ is a piecewise linear function, we have $f(x_i) = \tilde{f}_X(x_i) \leq \tilde{f}_X(x) \leq \tilde{f}_X(x_{i+1}) = f(x_{i+1})$. Again, we get $i\varepsilon f(v_{\max}) \leq \tilde{f}_X(x) \leq (i+1)\varepsilon f(v_{\max})$.

Combining the two inequalities, we get $f(x) - \tilde{f}_X(x) \leq \varepsilon f(v_{\max})$. Since f is concave, any piecewise linear segment joining two points on the function lies below the function and we have:

$$0 \leq f(x) - \tilde{f}_X(x) \leq \varepsilon f(v_{\max}).$$

□

Given any function f , we can use either of the two polygon approximation schemes defined by [Lemma C.3](#) and [Lemma C.4](#). Taking the better of these two schemes we get a proof of [Theorem 4.2](#).

Proof of Theorem 4.2. Let X_1 and X_2 be the sequences given by [Lemma C.3](#) and [Lemma C.4](#) respectively. Suppose $1/\varepsilon \leq \log v_{\max}$. In this case, the size of X_1 is at most $O(1/\varepsilon) = O\left(\min\left\{1/\varepsilon, \sqrt{\frac{\log v_{\max}}{\varepsilon}}\right\}\right)$. Now suppose $1/\varepsilon \geq \log v_{\max}$. In this case, the size of X_2 is at most $O\left(\log v_{\max} + \sqrt{\frac{\log v_{\max}}{\varepsilon}}\right) = O\left(\sqrt{\frac{\log v_{\max}}{\varepsilon}}\right) = O\left(\min\left\{1/\varepsilon, \sqrt{\frac{\log v_{\max}}{\varepsilon}}\right\}\right)$. Thus, the smaller of the two sequence is of size at most $O\left(\min\left\{1/\varepsilon, \sqrt{\frac{\log v_{\max}}{\varepsilon}}\right\}\right)$. \square

C.3 Proof of Proposition 4.3

Proof of Proposition 4.3. Consider any increasing concave function f and a polygon approximation \tilde{f}_X of f . Let $X = \{0 = x_0, \dots, x_k = v_{\max}\}$. Since f is concave, the function \tilde{f}_X lies ‘below’ the function f . Consider now the tangents to the function f drawn at points in X . Since f is concave, these tangents are always ‘above’ the function f . The tangents can be stitched together to define a function g that is always above the function f . Formally, define $g(x) = \min_{i \in \{0, \dots, k\}} \{f(x_i) + \min\{f^-(x_i) \cdot (x - x_i), f^+(x_i) \cdot (x - x_i)\}\}$. Observe that $f(x_i) = g(x_i)$, $f^-(x_i) = g^-(x_i)$, and $f^+(x_i) = g^+(x_i)$ for all i , giving us the second bullet. Observe also that g has at most $2k$ different slopes (because x_0 and x_k each only contribute one possibility, while the rest contribute two possibilities), so it is piecewise-linear with at most $2k$ segments.

It remains to show that there exists a polygon approximation \tilde{f}_X we can start with so that there is no polygon approximation of g using k segments for additive error ε . Let X be $k + 1$ points so as to minimize $\max_x \{f(x) - \tilde{f}_X(x)\}$, and call this value δ (by hypothesis, $\delta > \varepsilon$). This guarantees us the following property on X :

$$\max_{x \in [x_i, x_{i+1}]} \{f(x) - \tilde{f}_X(x)\} = \delta \text{ for all } i.$$

If not, then there is some i where $\max_{x \in [x_i, x_{i+1}]} \{f(x) - \tilde{f}_X(x)\} < \delta$. x_i could then be moved a teeny bit to the left (making the approximation within the interval $[x_i, x_{i+1}]$ a teeny bit worse) so that $\max_{x \in [x_i, x_{i+1}]} \{f(x) - \tilde{f}_X(x)\}$ remains $< \delta$, but $\max_{x \in [x_{i-1}, x_i]} \{f(x) - \tilde{f}_X(x)\}$ becomes $< \delta$ as well (as it also makes the approximation within the interval $[x_{i-1}, x_i]$ a teeny bit better, and it was $\leq \delta$ to begin with). Iterating this procedure would then let us make $\max_{x \in [x_j, x_{j+1}]} \{f(x) - \tilde{f}_X(x)\} < \delta$ for all $j \leq i$. And a similar argument for moving x_{i+1} a teeny bit to the right allows us to claim the same for $j \geq i$. At the end, we have now made $\max_x \{f(x) - \tilde{f}_X(x)\}$ strictly less than δ , a contradiction.

Now that we have $\max_{x \in [x_i, x_{i+1}]} \{f(x) - \tilde{f}_X(x)\} = \delta > \varepsilon$ for all i , we can immediately see that for any function $g(\cdot)$ with $g(x_i) = f(x_i)$ for all i , and $g(x) \geq f(x)$ for all x , any Y such that $\tilde{g}_Y(\cdot)$ is an ε -approximation polygon approximation requires at least one point $y \in Y \cap [x_i, x_{i+1}]$. So we can immediately see that g also has no polygon approximation using k segments for additive error ε . \square

With this in mind, we’ll now restrict attention to piecewise-linear functions with not too many segments. To this end, first consider a piecewise-linear function with just two segments. The first segment is of length l_1 with slope m_1 . Similarly, the second segment is of length l_2 and slope $m_2 < m_1$. If we wish to approximate this function using just one segment, the segment should have length $l_1 + l_2$ and slope $\frac{m_1 l_1 + m_2 l_2}{l_1 + l_2}$. The error in this approximation will be $\frac{(m_1 - m_2) l_1 l_2}{l_1 + l_2}$. This calculation says that the error roughly depends on the product of the slope difference and the lengths. For our lower bound, we would wish to have many such segments

combined together so that the error in all of the segments is the same. For a concave function, the slope keeps decreasing as x grows. Thus, if the error has to be the same in every segment, the lengths of each segment have to keep increasing so that the product is the same. Observe that the function $\log(x)$ satisfies this property.

With this direction in mind, we now demonstrate a function that we call LPL_k (for logarithmically-piecewise-linear) that can't be approximated using less than k segments. We chose this name because the function has derivative $1/2^i$ in the range $[3(2^{i+1} - 2), 3(2^{i+2} - 2)]$ which is (roughly) a constant factor away from the derivative $1/(x + 6)$ of the logarithmic function $\ln(x + 6)$. In fact, what we do can be seen as stitching together appropriate tangents to the logarithmic curve. For a value of $\varepsilon < 1/100$, let $k = \frac{1}{100\varepsilon}$. Consider the following function defined on $[0, 3(2^{k+1} - 2) = v_{\max}]$:

$$LPL_k(x) = \begin{cases} x, & 0 \leq x \leq 3(2^2 - 2) \\ \frac{1}{2}(x - 3(2^2 - 2)) + 6, & 3(2^2 - 2) < x \leq 3(2^3 - 2) \\ \frac{1}{4}(x - 3(2^3 - 2)) + 12, & 3(2^3 - 2) < x \leq 3(2^4 - 2) \\ \vdots \\ \frac{1}{2^{k-1}}(x - 3(2^k - 2)) + 6(k - 1), & 3(2^k - 2) < x \leq 3(2^{k+1} - 2). \end{cases}$$

The function LPL_k satisfies $\varepsilon(1 + LPL_k(v_{\max})) = \varepsilon + 6k\varepsilon \leq \frac{1}{2}$, and that $f'(0) \leq 1$, $f'(v_{\max}) \geq 0$. We prove that any good polygon approximation for $LPL_k(x)$ must have at least $\Omega(k)$ segments. Since $k = \frac{1}{100\varepsilon}$, this proves that a dependence of $1/\varepsilon$ is unavoidable if we want to be independent of v_{\max} . Since in this example, $\sqrt{\frac{\log v_{\max}}{\varepsilon}} = \Theta(k)$, this also shows that the bound in [Theorem 4.2](#) is tight. The idea of our proof is to identify disjoint intervals such that any sequence X that defines a good polygon approximation should have at least one point in every such interval. We will identify $\Theta(k)$ such intervals for LPL_k .

Lemma C.5. *Suppose there exists a sequence $X = \{x_i\}_{i \geq 1}$ of points such that $LPL_k(x) - 1/2 \leq \widetilde{LPL}_{kX}(x) \leq LPL_k(x)$ for all $x \in [0, v_{\max}]$. The sequence X has at least one point in the range $I_i = [3(2^i - 2) - 2^i, 3(2^i - 2) + 2^i]$ for all $2 \leq i \leq k$.*

Proof of Lemma C.5. Suppose for the sake of contradiction there is an i for which $X \cup I_i = \phi$. We prove that the approximation doesn't hold at the midpoint of I_i (call this m). Since no point of X is in I_i and LPL_k is concave, $\widetilde{LPL}_{kX}(m)$ is at most $\frac{1}{2}(LPL_k(3(2^i - 2) - 2^i) + LPL_k(3(2^i - 2) + 2^i))$. Thus, $LPL_k(m) - \widetilde{LPL}_{kX}(m) \geq 6(i - 1) - \frac{1}{2}(6(i - 2) + 2 + 6(i - 1) + 2) = 1$. \square

The following corollary follows immediately from [Lemma C.5](#).

Corollary C.6. *An $\varepsilon(1 + LPL_k(v_{\max}))$ additive polygon approximation of LPL_k above requires at least $k - 1$ segments.*

[Corollary C.6](#) shows that the bounds of [Theorem 4.2](#) are tight, both in the case that there is no dependence on v_{\max} , and in the case that we allow dependence v_{\max} . Moreover, [Proposition 4.3](#) shows that "LPL-like" functions exhibit the worst-case gaps for whatever kinds of guarantees are desired.

C.4 From Polygon Approximation to Approximate Auctions

In this subsection we prove [Proposition C.7](#) and [Corollary 4.1](#) establishing the connection between "small" polygon approximations and approximate auctions with low menu complexity. The high level idea of the proof of [Proposition C.7](#) is that the revenue generated by a mechanism on days i through n is equal to a weighted sum of the function $R_{\geq i}(\cdot)$ evaluated at points where the allocation curve changes (see [Lemma 2.1](#)). If a good polygon approximation of $R_{\geq i}$ exists,

then every point where the allocation curve changes can be substituted by the endpoints of the segment of polygon approximation that contains it. The error in the polygon approximation would become loss in revenue after this transformation. The number of segments (or endpoints) in the polygon approximation would be the menu complexity of the transformed mechanism. Thus, a good and small polygon approximation would generate an almost optimal and low menu complexity mechanism. The proof of [Corollary 4.1](#) relies on iteratively calling on [Proposition C.7](#) from the first day to the last while keeping a check on the total revenue lost compared to the optimal over all days.

Proposition C.7. *Consider a FedEx instance with n deadlines. For $i \in \{1, 2, \dots, n\}$, let g be the function $\tilde{R}_{\geq i}$ defined in [2.3](#), and let X be a sequence of k points in $[0, r_{\geq i}]$ such that for all $x \leq r_{\geq i}$, we have $g(x) - \varepsilon \leq \tilde{g}_X(x) \leq g(x)$. Then for any mechanism M , there exists a mechanism with the following properties:*

- The menu offered on days $1, 2, \dots, i - 1$ is the same as M .
- The i -deadline menu complexity is at most $2k$.
- The revenue is at least $\text{OPT}_{M,i} - \varepsilon$.

Here, $\text{OPT}_{M,i}$ denotes the optimal revenue for any mechanism that offers the same menu as M on deadlines $1, 2, \dots, i - 1$.

Intuitively, every segment of \tilde{g}_X corresponds to a different menu option in the approximate mechanism, and the fact that $\tilde{g}_X(x)$ and $g(x) = \tilde{R}_{\geq i}(x)$ are close for all x implies that the revenue of the mechanism based on \tilde{g}_X isn't far from optimal (by [Lemma 2.1](#)). The additional factor of two comes from the fact that $\tilde{R}_{\geq i}$ is ironed, so to “set price p ” in $\tilde{R}_{\geq i}$, we might need to randomize over two prices in $R_{\geq i}$.

Proof of Proposition C.7. We assume without loss of generality that the mechanism M does the optimal pricing on days i through n (given the prices it sets on days 1 through $i - 1$). Thus, the revenue generated by M is $\text{OPT}_{M,i}$. The menu offered by any such mechanism on day i can be seen as a distribution f over prices $\{p_j\}$. The prices p_j are all at most $r_{\geq i}$ as the M behaves optimally on day i (see [A.1](#)). Since the mechanism behaves optimally day $i + 1$ onwards, the revenue generated by M on days i through n is $\sum_j f(p_j)R_{\geq i}(p_j)$ by [Lemma 2.1](#).

Let X be the sequence of k points such that $g(x) - \varepsilon \leq \tilde{g}_X(x) \leq g(x)$ for all $x \leq r_{\geq i}$. Since $g = \tilde{R}_{\geq i}$ is the upper concave envelope of $R_{\geq i}$, for every point $x \leq r_{\geq i}$, there exists points $x_l \leq x$ and $x_r \geq x$ in $[0, r_{\geq i}]$ such that $g(x_l) = R_{\geq i}(x_l)$, $g(x_r) = R_{\geq i}(x_r)$, and for any $\lambda \in [0, 1]$, we have $g(\lambda x_l + (1 - \lambda)x_r) = \lambda g(x_l) + (1 - \lambda)g(x_r)$. For all points $x \in X$, we add the points x_l and x_r to X and remove the point x . Adding the points x_l and x_r will only improve the quality of the polygon approximation. Further, since g is linear on the interval $[x_l, x_r]$ that contains x , removing x doesn't affect the quality of the polygon approximation. Thus, the new set X satisfies

$$g(x) - \varepsilon \leq \tilde{g}_X(x) \leq g(x). \quad (2)$$

Further, by construction all elements $v \in X$ satisfy

$$g(v) = R_{\geq i}(v). \quad (3)$$

The set X now has at most $2k$ points. For any $x \leq r_{\geq i}$, we let \underline{x} (resp. \bar{x}) denote the largest (resp. smallest) value in X that is at most (at least) x . With this notation, if $x = \lambda_x \underline{x} + (1 - \lambda_x) \bar{x}$, then we have that

$$\tilde{g}_X(x) = \lambda_x \tilde{g}_X(\underline{x}) + (1 - \lambda_x) \tilde{g}_X(\bar{x}) = \lambda_x g(\underline{x}) + (1 - \lambda_x) g(\bar{x}). \quad (4)$$

We are now ready to define a mechanism M' with high revenue and low i -deadline menu complexity. The mechanism M' is defined as follows:

- Mimic M on days $1, 2, \dots, i - 1$.
- On day i , if p_j has mass $f(p_j)$, add mass $f(p_j)\lambda_{p_j}$ to \underline{p}_j and mass $f(p_j)(1 - \lambda_{p_j})$ to \bar{p}_j . This generates a new distribution f' on X which defines our menu for day i .
- Do the optimal pricing days $i + 1$ onwards given the prices on day 1 through i .

Since the menu offered on day i is defined by a distribution over X , its size is at most $2k$, the size of X . We finish the proof by showing that M' is feasible and has high revenue.

Claim C.1. *The mechanism M' defined above is feasible.*

Proof. We only need to check the intra-day incentive compatibility constraints. For days 1 through $i - 1$, M' mimics M and thus, these constraints are satisfied. Days $i + 1$ onwards, M' prices optimally and thus, satisfies the constraints. We only argue about the constraints for day i . We wish to prove that the utility of any bidder with type (x, i) in mechanism M' is at least that of the bidder with type $(x, i - 1)$. We know that the utility of any bidder with type (x, i) in mechanism M is at least that of the bidder with type $(x, i - 1)$. Since the mechanisms M and M' are the same on days 1 through $i - 1$, it is sufficient to prove that the utility of any bidder with type (x, i) in mechanism M' is at least that of the bidder with type (x, i) in mechanism M . This follows directly from the expression of utility $\sum_{p_j \leq x} f(p_j)(x - p_j)$. We have

$$\begin{aligned}
& \sum_{X \ni v \leq x} f'(v)(x - v) \\
&= \sum_{p_j \leq x} f(p_j)\lambda_{p_j}(x - \underline{p}_j) + f(p_j)(1 - \lambda_{p_j})(x - \bar{p}_j) + \sum_{p_j \in (\underline{x}, x]} f(p_j)\lambda_{p_j}(x - \underline{p}_j) \quad (\text{Definition of } f') \\
&= \sum_{p_j \leq x} f(p_j)(x - p_j) + \sum_{p_j \in (\underline{x}, x]} f(p_j) \left(1 - \frac{p_j - \underline{p}_j}{\bar{p}_j - \underline{p}_j}\right) (x - \underline{p}_j) \quad (\text{Definition of } \lambda_{p_j}) \\
&\geq \sum_{p_j \leq x} f(p_j)(x - p_j) + \sum_{p_j \in (\underline{x}, x]} f(p_j) \left(1 - \frac{p_j - \underline{p}_j}{x - \underline{p}_j}\right) (x - \underline{p}_j) \quad (\text{Since } \bar{p}_j > x) \\
&= \sum_{p_j \leq x} f(p_j)(x - p_j) + \sum_{p_j \in (\underline{x}, x]} f(p_j)(x - p_j) = \sum_{p_j \leq x} f(p_j)(x - p_j).
\end{aligned}$$

□

Claim C.2. *The revenue generated by M' is at least $\text{OPT}_{M,i} - \varepsilon$.*

Proof. First, we note M and M' are identical for the first $i - 1$ days. Thus, the revenue generated on the first $i - 1$ days is the same for both M and M' . Since M' prices optimally day $i + 1$ onwards, the revenue generated on day i through n is $\sum_{v \in X} f'(v)R_{\geq i}(v)$. We have

$$\begin{aligned}
\sum_{v \in X} f'(v)R_{\geq i}(v) &= \sum_{v \in X} f'(v)g(v) && (\text{Equation 3}) \\
&= \sum_{v_j} f(v_j)\lambda_{v_j}g(\underline{v}_j) + f(v_j)(1 - \lambda_{v_j})g(\bar{v}_j) && (\text{Definition of } f') \\
&= \sum_{v_j} f(v_j)\tilde{g}_X(v_j) && (\text{Equation 4})
\end{aligned}$$

$$\begin{aligned}
&\geq \sum_{v_j} f(v_j) (g(v_j) - \varepsilon) && \text{(Equation 2)} \\
&\geq \sum_{v_j} f(v_j) R_{\geq i}(v_j) - \varepsilon.
\end{aligned}$$

This analysis proves that the total revenue generated by M' is at most ε less than that generated by M . In other words, it is at least $\text{OPT}_{M,i} - \varepsilon$. □

These two claims combined suffice to prove the proposition. □

Having described the construction of low i -deadline menu complexity mechanisms, we repeat this construction n times (once for each deadline) to get our approximately optimal mechanism. We also show how to use [Proposition C.7](#) to go from the optimal mechanism to an entire menu of low menu complexity. Thus, the total revenue loss on days i through n is at most ε . Using the above result, we prove [Corollary 4.1](#).

Proof of Corollary 4.1. We prove that for all $i \in \{0, 1, \dots, n\}$, there exists a mechanism with j -deadline menu complexity of k_j for all $1 \leq j \leq i$ whose revenue is at least $\text{OPT} - \sum_{j=1}^i \varepsilon_j$. For $i = n$, this is the same as the statement of the corollary. This proof will proceed via induction on i . For $i = 0$, the statement is trivial. We assume the statement for $i - 1$ and prove it for i .

Let M_{i-1} be the mechanism that is promised by the induction hypothesis and let $\text{Rev}_{i-1} \geq \text{OPT} - \sum_{j=1}^{i-1} \varepsilon_j$ be its revenue.

We invoke [Proposition C.7](#) on X_i and M_{i-1} to get a mechanism M_i with the following properties:

- The menu offered by M_i on days $1, 2, \dots, i - 1$ is the same as M_{i-1} .
- The i -deadline menu complexity is at most $2k_i$.
- The revenue is at least $\text{OPT}_{M_{i-1},i} - \varepsilon_i$.

Property 1 and 2 imply that the j -deadline menu complexity of M_i is at most k_j for all $1 \leq j \leq i$. The revenue of M_i is at least $\text{OPT}_{M_{i-1},i} - \varepsilon_i \geq \text{Rev}_{i-1} - \varepsilon_i \geq \text{OPT} - \sum_{j=1}^{i-1} \varepsilon_j - \varepsilon_i = \text{OPT} - \sum_{j=1}^i \varepsilon_j$. This completes the proof. □

Finally, we may complete the proof of [Theorem 1.2](#).

Proof of Theorem 1.2. For any FedEx instance where the bidder's types have integral support $1, 2, \dots, v_{\max}$, observe that the optimal revenue is at least 1. This is because 1 is the revenue generated by the auction that offers the price 1 on all deadlines. Further, for all deadlines i , $1 + \tilde{R}_{\geq i}(v_{\max}) \leq \text{OPT} + \text{OPT} = 2\text{OPT}$.

Note that all revenue curves $\tilde{R}_{\geq i} = g_i$ when restricted to $[0, r_{\geq i}]$ satisfy the requirements of [Theorem 4.2](#). This implies that for all i , there exists sequences X_i of at most $O(\min\{n/\varepsilon, \sqrt{\frac{n}{\varepsilon} \log v_{\max}}\})$ points such that for all $x \in [0, r_{\geq i}]$,

$$g_i(x) - \frac{\varepsilon}{2n} (1 + g_i(v_{\max})) \leq \tilde{g}_{i X_i}(x) \leq g_i(x).$$

We set $\varepsilon_i = \frac{\varepsilon}{2n} (1 + g_i(v_{\max}))$ and use [Corollary 4.1](#) to get that there exists a mechanism with menu complexity at most $O(n \min\{n/\varepsilon, \sqrt{\frac{n}{\varepsilon} \log v_{\max}}\})$ whose revenue is at least $\text{OPT} - \sum_{i=1}^n \varepsilon_i = \text{OPT} - \sum_{i=1}^n \frac{\varepsilon}{2n} (1 + g_i(v_{\max})) \geq \text{OPT} - \sum_{i=1}^n \varepsilon \frac{\text{OPT}}{n} = (1 - \varepsilon)\text{OPT}$. □

D Analysis omitted in Section 5

D.1 The instance

In this section, we describe our instance of the FedEx problem that is hard to approximate using small menus. Before delving into the details of our parameters, we give a brief high level intuition.

In Section 4.1, we described a function LPL_k that is hard to polygon approximate. The function LPL_k provided a lower bound for Theorem 4.2 that was tight (modulo constant factors). The idea behind this example was to have a piecewise linear function such that the slope of the constituent segments decreases gradually in a controlled way. Because of the intimate connection between polygon approximation and menu complexity (Proposition C.7), similar ideas should also help in constructing a FedEx instance that is hard to approximate using small menus.

The construction of the hard FedEx instance can be expected to have many additional complications. The complications arise from the fact that each revenue curve $R_{\geq i}$ in a FedEx instance is formed by appropriately adding revenue curves of multiple distributions (see Definition 2.3) and $\tilde{R}_{\geq i}$ is obtained by ‘ironing’ $R_{\geq i}$. Since the size of the optimal menu only increases when prices ‘split’ across ironed intervals, it is necessary to have multiple (polynomial, in our case) ironed intervals in the function $\tilde{R}_{\geq i}$.

The presence of multiple ironed intervals in $\tilde{R}_{\geq i}$ is, however, not sufficient. It is easy to see why. Imagine an example where all the revenue curves R_i are generated from the same underlying distribution. If this is the case, all the curves $R_{\geq i}$ might have the same ironed intervals.¹² If the ironed intervals in the curves $\tilde{R}_{\geq i}$ for all i correspond, then even the optimal revenue auction would require only 1 price on each day. This is because only those prices ‘split’ that are inside an ironed interval. If all the ironed intervals correspond, then no price can ever lie inside a interval. The solution is to have ironed intervals ‘nested’ so that the ends of the interval on day i lie inside the intervals on day $i + 1$.

The two points raised above concern splitting in general and were also true of our worst case example described in Section 3. If we want a lower bound for all approximate auctions, we also want the slope between two consecutive segments to be very different (as in LPL_k). This ensures that we need at least one menu option for each segment for a good approximation.

With these ideas in mind, we describe our hard instance. We assume that the bidder’s type is drawn from a distribution supported on $\{1, 2, \dots, v_{\max}\} \times \{1, 2, \dots, n\}$. The value of v_{\max} in our example is $5n$ and the distribution q over days is uniform. We make explicit that we don’t include the q in our calculation. Since they only scale all revenue curves by $1/n$, they don’t affect the quality of a $1 - \epsilon$ approximation. The marginal distribution F_i on day i is designed so that the revenue curve $\tilde{R}_{\geq i}$ has $i + 1$ segments. The slope of the first i segments is geometrically decreasing with common ratio $\lambda = 1 + \frac{1}{4n}$. This value is large enough so that a $1 - O(1/n^2)$ approximation will require one menu option for each segment. It also satisfies,

Claim D.1. For $n \geq i \geq 0$, $\frac{1}{\lambda^i} > \frac{3}{4}$.

Proof.

$$\frac{3}{4}\lambda^i < \frac{3}{4}\left(1 + \frac{1}{4n}\right)^n \leq \frac{3}{4}e^{\frac{1}{4}} < 1.$$

□

As we define the slopes to be in geometric progression, the geometric sum $1 + \frac{1}{\lambda} \dots$ will appear multiple times in our description and its analysis. To avoid writing the all the terms everywhere, we define $S_0 = n$ and $S_i = n + \sum_{j=0}^{i-1} \lambda^{-j}$ for all $i \geq 1$.

¹²This is not always true as whether a particular interval is ironed or not for a given deadline depends on many other factors, e.g. the distribution q_i across the various days.

The $(i+1)^{\text{th}}$ segment of $\tilde{R}_{\geq i}$ will have slope $(n+1-i)C$ where $C = 1/2$. This last segment will split into two segments on day $i+1$. The first segment will have slope $1/\lambda^i$, continuing the geometric progression. The last segment will have slope $(n-i)C$, allowing it to be split similarly in the future. In order to ensure this behavior of $\tilde{R}_{\geq i}$, we need the slope of R_i to be β_i for this segment, where

$$\beta_i = \frac{1}{3n+i} \left(\frac{3n-i}{2} - \frac{n-i}{\lambda^i} + S_i \right). \quad (5)$$

This value of β_i satisfies the following two properties:

Claim D.2. β_i is an increasing sequence for $n \geq i \geq 0$.

Proof. Fix $n > i \geq 0$.

$$\begin{aligned} & \beta_{i+1} - \beta_i \\ &= \frac{1}{3n+i+1} \left(\frac{3n-i-1}{2} - \frac{n-i-1}{\lambda^{i+1}} + S_{i+1} \right) - \frac{1}{3n+i} \left(\frac{3n-i}{2} - \frac{n-i}{\lambda^i} + S_i \right) \quad (\text{Equation 5}) \\ &= \frac{1}{(3n+i)(3n+i+1)} \left(-3n + \frac{(\lambda-1)(3n+i+1)(n-i)}{\lambda^{i+1}} + \frac{4n}{\lambda^{i+1}} + \frac{3n+i}{\lambda^i} - S_i \right) \\ &\geq \frac{1}{(3n+i)(3n+i+1)} \left(\frac{3(3n+i+1)(n-i)}{16n} + \frac{3(3n+i)}{4} - S_i \right) \quad (\text{Claim D.1}) \\ &\geq \frac{1}{(3n+i)(3n+i+1)} \left(\frac{9n}{4} - 2n \right) > 0. \end{aligned}$$

□

Claim D.3. For $n \geq i \geq 1$, $\frac{1}{2} < \beta_1 = \frac{1}{2} + \frac{5}{4\lambda(3n+1)} \leq \beta_i \leq \frac{3}{4}$.

Proof. Verify that β_1 is indeed correct. After proving Claim D.2, it is sufficient to establish $\beta_n \leq \frac{3}{4}$.

$$\beta_n = \frac{1}{4n}(n + S_n) \leq \frac{3n}{4n}.$$

□

We now define the marginal distribution F_i on the i^{th} deadline. We write an expression for the distribution and then, calculate the revenue curve R_i of the distribution F_i . The part of the analysis where we prove that F_i is a valid distribution is omitted. This would use $\frac{S_i}{n+i} \geq \frac{S_{i+1}}{n+i+1}$ and $\beta_i \leq \frac{S_i}{n+i}$ (where Claim D.1 and Claim D.3 help prove the latter) and can easily be done by the interested reader. The distribution over types on the i^{th} day is given by :

$$F_i(x) = \begin{cases} 0, & 0 \leq x \leq n \\ \left(1 - \frac{S_1}{n+1}\right), & n \leq x \leq n+1 \\ \left(1 - \frac{S_2}{n+2}\right), & n+1 \leq x \leq n+2 \\ \left(1 - \frac{S_3}{n+3}\right), & n+2 \leq x \leq n+3 \\ \vdots \\ \left(1 - \frac{S_i}{n+i}\right), & n+i-1 \leq x \leq n+i \\ (1 - \beta_i), & n+i \leq x \leq 3n+i \\ 1, & 3n+i \leq x \leq 5n. \end{cases} \quad (6)$$

The revenue curve for this distribution is given by $R_i(v) = v(1 - F_i(v))$:

$$R_i(x) = \begin{cases} x, & 0 \leq x \leq n \\ \left(\frac{S_1}{n+1}\right)x, & n < x \leq n+1 \\ \left(\frac{S_2}{n+2}\right)x, & n+1 < x \leq n+2 \\ \left(\frac{S_3}{n+3}\right)x, & n+2 < x \leq n+3 \\ \vdots \\ \left(\frac{S_i}{n+i}\right)x, & n+i-1 < x \leq n+i \\ \beta_i x, & n+i < x \leq 3n+i \\ 0, & 3n+i < x \leq 5n. \end{cases} \quad (7)$$

Using Definition 2.3, we now calculate the combined revenue curve $R_{\geq i}$ for days i through n . We prove that $R_{\geq i}$ has the form we desire using induction.

Theorem D.1.

$$R_{\geq i}(x) = \begin{cases} x + (n-i)x, & 0 \leq x \leq n \\ \left(\frac{S_1}{n+1}\right)x + (n-i)(x - n + S_0), & n < x \leq n+1 \\ \left(\frac{S_2}{n+2}\right)x + (n-i)\left(\frac{x-n-1}{\lambda} + S_1\right), & n+1 < x \leq n+2 \\ \left(\frac{S_3}{n+3}\right)x + (n-i)\left(\frac{x-n-2}{\lambda^2} + S_2\right), & n+2 < x \leq n+3 \\ \vdots \\ \left(\frac{S_i}{n+i}\right)x + (n-i)\left(\frac{x+1-n-i}{\lambda^{i-1}} + S_{i-1}\right), & n+i-1 < x \leq n+i \\ \beta_i x + (n-i)\left(\frac{x-n-i}{\lambda^i} + S_i\right), & n+i < x \leq n+i+1 \\ \beta_i x + (n-i)(C(x-n-i-1) + S_{i+1}), & n+i+1 < x \leq 3n+i \\ (n-i)(C(x-n-i-1) + S_{i+1}), & 3n+i < x \leq 3n+i+1 \\ (n-i)(2nC + S_{i+1}), & 3n+i+1 < x \leq 5n. \end{cases} \quad (8)$$

Proof. Proof by backwards induction. The base case $i = n$ is easily verified. For the inductive step, we first calculate $\tilde{R}_{\geq i}$. This is done by ironing the curve $R_{\geq i}$ defined in Equation 8. We provide an expression for $\tilde{R}_{\geq i}$ and prove that it is correct in Lemma D.4 in Subsection D.2 of this Appendix. The expression is:

$$\tilde{R}_{\geq i}(x) = \begin{cases} (n+1-i)x, & 0 \leq x \leq n \\ (n+1-i)(x - n + S_0), & n < x \leq n+1 \\ (n+1-i)\left(\frac{x-n-1}{\lambda} + S_1\right), & n+1 < x \leq n+2 \\ (n+1-i)\left(\frac{x-n-2}{\lambda^2} + S_2\right), & n+2 < x \leq n+3 \\ \vdots \\ (n+1-i)\left(\frac{x+1-n-i}{\lambda^{i-1}} + S_{i-1}\right), & n+i-1 < x \leq n+i \\ (n+1-i)(C(x-n-i) + S_i), & n+i < x \leq 3n+i \\ (n-i)(2nC + S_{i+1}) + \frac{(n-i)C - \beta_i(3n+i)}{2n-i}(x - 5n), & 3n+i < x \leq 5n. \end{cases} \quad (9)$$

Note that this function is maximized at $r_{\geq i} = 3n+i$. We now use Definition 2.3 which says:

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

Using Equation 7 and Equation 9 in this gives:

$$R_{\geq i-1}(x) = \begin{cases} x + (n+1-i)x, & 0 \leq x \leq n \\ \left(\frac{S_1}{n+1}\right)x + (n+1-i)(x-n+S_0), & n < x \leq n+1 \\ \left(\frac{S_2}{n+2}\right)x + (n+1-i)\left(\frac{x-n-1}{\lambda} + S_1\right), & n+1 < x \leq n+2 \\ \left(\frac{S_3}{n+3}\right)x + (n+1-i)\left(\frac{x-n-2}{\lambda^2} + S_2\right), & n+2 < x \leq n+3 \\ \vdots \\ \left(\frac{S_{i-1}}{n+i-1}\right)x + (n+1-i)\left(\frac{x+2-n-i}{\lambda^{i-2}} + S_{i-2}\right), & n+i-2 < x \leq n+i-1 \\ \beta_{i-1}x + (n+1-i)\left(\frac{x+1-n-i}{\lambda^{i-1}} + S_{i-1}\right), & n+i-1 < x \leq n+i \\ \beta_{i-1}x + (n+1-i)(C(x-n-i) + S_i), & n+i < x \leq 3n+i-1 \\ (n+1-i)(C(x-n-i) + S_i), & 3n+i-1 < x \leq 3n+i \\ (n+1-i)(2nC + S_i), & 3n+i < x \leq 5n. \end{cases}$$

which is of the form required by the induction hypothesis. \square

It is relatively easy to find the optimal auction in our instance using Fiat et al.'s algorithm [FGKK16]. The optimal auction closely follows the behavior of the curves $\tilde{R}_{\geq i}$. It has one menu option $= 3n+1$ on day 1. On each subsequent day, the last menu option splits into two options while all the other options are carried as is. For example, the price $3n+1$ splits into two prices of $n+2$ and $3n+2$ on day 2. The price $n+2$ is pushed forward to all of the days while the price $3n+2$ is split into $n+3$ and $3n+3$. The price $n+3$ is pushed through while the price $3n+3$ is again split into two on day 4. This process goes on till the n^{th} day. The menu offered on day n has options $n+2, n+3, \dots, n+n$ and $4n$.

Remark D.1. *The optimal revenue in the setting described is at most $\max R_{\geq 1} = n(2n+1)$ which is between $2n^2$ and $3n^2$. This is denoted by OPT throughout.*

D.2 Omitted details from the proof of Theorem D.1

In this subsection we show that $\tilde{R}_{\geq i}(x)$ is the upper concave envelope of $R_{\geq i}(x)$, where both functions are as defined in Appendix D. The proof consists of showing that the two functions agree on the points where $\tilde{R}_{\geq i}(x)$ changes, and that $\tilde{R}_{\geq i}(x)$ is greater than $R_{\geq i}(x)$ on the other points. This, combined with the fact that $\tilde{R}_{\geq i}(x)$ is linear in between the points where it changes, suffices to show the claim.

Lemma D.2. *For all $x \in \{0, n, n+1, n+2, \dots, n+i-1, n+i, 3n+i, 5n\}$, we have $\tilde{R}_{\geq i}(x) = R_{\geq i}(x)$.*

Proof. We prove this lemma using four claims, all of which establish $\tilde{R}_{\geq i}(x) = R_{\geq i}(x)$ for different domains of x .

Claim D.4. *For all $x \in \{0, n\}$, we have $\tilde{R}_{\geq i}(x) = R_{\geq i}(x)$.*

Proof. By verification: $x + (n-i)x = (n+1-i)x$. \square

Claim D.5. *For all $x \in \{n+1, n+2, \dots, n+i-1, n+i\}$, we have $\tilde{R}_{\geq i}(x) = R_{\geq i}(x)$.*

Proof. By verification:

$$\tilde{R}_{\geq i}(x) - R_{\geq i}(x) = \left(\frac{1}{\lambda^{x-n-1}} + S_{x-n-1}\right) - S_{x-n} = 0.$$

\square

Claim D.6. $\tilde{R}_{\geq i}(3n+i) = R_{\geq i}(3n+i)$.

Proof.

$$\begin{aligned}
R_{\geq i}(3n+i) &= \beta_i(3n+i) + (n-i)(C(2n-1) + S_{i+1}) \\
&= \left(\frac{3n-i}{2} - \frac{n-i}{\lambda^i} + S_i \right) + (n-i)(C(2n-1) + S_{i+1}) \quad (\text{Equation 5}) \\
&= n + \left((n-i)C - \frac{n-i}{\lambda^i} + S_i \right) + (n-i)(C(2n-1) + S_{i+1}) \\
&= n + S_i + (n-i)(2nC + S_i) \\
&= (n+1-i)(2nC + S_i) = \tilde{R}_{\geq i}(3n+i).
\end{aligned}$$

□

Claim D.7. $\tilde{R}_{\geq i}(5n) = R_{\geq i}(5n)$.

Proof. This is easily verifiable from the description of the functions.

□

□

Lemma D.3. For all $0 \leq x \leq 5n$, we have $\tilde{R}_{\geq i}(x) \geq R_{\geq i}(x)$.

Proof. Note that both functions are piecewise linear. The function $\tilde{R}_{\geq i}$ changes form at points $S_1 = \{0, n, n+1, n+2, \dots, n+i-1, n+i, 3n+i, 5n\}$. The function $R_{\geq i}$ changes form at points in $S_2 = S_1 \cup \{n+i+1, 3n+i+1\}$. To prove this lemma, it is sufficient to show that $\tilde{R}_{\geq i}(x) \geq R_{\geq i}(x)$ for all $x \in S_2$. If $x \in S_1$, this follows by [Lemma D.2](#). We only prove for $x \in \{n+i+1, 3n+i+1\}$.

If $x = n+i+1$, we get

$$\begin{aligned}
&\tilde{R}_{\geq i}(n+i+1) - R_{\geq i}(n+i+1) \\
&= (n+1-i)(C + S_i) - \beta_i(n+i+1) - (n-i) \left(\frac{1}{\lambda^i} + S_i \right) \\
&= C + S_i + (n-i) \left(C - \frac{1}{\lambda^i} \right) - \beta_i(n+i+1) \\
&= C + \beta_i(3n+i) - 2nC - \beta_i(n+i+1) \quad (\text{Equation 5}) \\
&= (\beta_i - C)(2n-1) \\
&> 0. \quad (\text{D.3})
\end{aligned}$$

If $x = 3n+i+1$, we get

$$\tilde{R}_{\geq i}(3n+i+1) \geq (n-i)(2nC + S_{i+1}) = R_{\geq i}(3n+i+1),$$

as the other term in the definition of $\tilde{R}_{\geq i}$ is positive.

□

Lemma D.4. The upper concave envelope of the function $R_{\geq i}$ defined in [Equation 8](#) is the function $\tilde{R}_{\geq i}$ defined in [Equation 9](#).

Proof. By definition, $\tilde{R}_{\geq i}$ is a continuous piecewise linear function whose successive segments have decreasing slope. Thus, it is concave. [Lemma D.3](#) shows that the function $\tilde{R}_{\geq i}$ is always larger than $R_{\geq i}$. Note that the function $\tilde{R}_{\geq i}$ is piecewise linear and changes form at points $0, n, n+1, n+2, \dots, n+i-1, n+i, 3n+i$, and $5n$. [Lemma D.2](#) shows that at all these points, the value of $\tilde{R}_{\geq i}$ is that same as that of $R_{\geq i}$. This finishes the proof. □

D.3 Cleaning

We start by fixing notation that we employ in this subsection. The letters R , \tilde{R} , q , and F are reserved for the revenue curves and the type distribution of our FedEx instance. These are described in [Equation 8](#), [Equation 9](#), and [Equation 6](#) respectively (q is assumed to be uniform and is omitted throughout). The letters A and B would denote the allocation curves of mechanisms, *i.e.*, the allocation curve on day i would be denoted by A_i or B_i . Thus, for $x \in [0, v_{\max}]$, a bidder that reports type (x, i) gets the item with probability $A_i(x)$. We assume without loss of generality that $A_i(0) = 0$ and $A_i(v_{\max}) = 1$. In line with the notation used for probability distributions, we reserve the lower case a_i for the ‘allocation density’, *i.e.*, $a_i(x) = A_i(x) - A_i(x - 1)$. With this notation, the menu complexity of a mechanism M_A is the number of (x, i) such that $a_i(x) \neq 0$. Unless specified otherwise, all statements about clean mechanisms pertain our particular instance of the FedEx problem defined in the previous section.

In the FedEx problem, menus on day i are constrained by the menu offered on day $i - 1$. These constraints are represented by the downwards IC constraints in [Equation 1](#). In this sense, some menus on day $i - 1$ are strictly more constraining than others. For example, a price of p_1 is offered on day $i - 1$ is strictly more constraining than a price of $p_2 > p_1$. This is because the utility $\pi(v, i - 1)v - p(v, i - 1)$ is strictly higher. Similarly, a distribution over two prices of p_1 and p_2 with probabilities q and $1 - q$ respectively is strictly more constraining than offering the single price $p = qp_1 + (1 - q)p_2$. This reasoning also applies in case the price p has a small mass in a larger distribution. In the case a feasible less constraining menu generates a higher revenue on a particular day than a more constraining one, the latter can be changed to the former without violating any of the feasibility constraints. This change will only increase the revenue generated.

In the problem instance we consider, the revenue curve R_i for a given day i is increasing and touches the upper concave envelope at all points in $\mathbb{Z} \cap [0, n + i]$. Thus, if a price p in this range is offered with probability $f(p)$ on day $i - 1$, the least constraining and revenue maximizing option on day i is to offer the same price with the same probability. We define ‘clean’ mechanisms to be mechanisms that satisfy this property. Formally,

Definition D.1 (Clean mechanisms). *Let a_i be the allocation density for day i of a mechanism M . The mechanism M is said to be clean if for all days $i \geq 1$ and any point $0 \leq x \leq n + i$, we have*

$$a_{i+1}(x) \geq a_i(x).$$

The reason we do not require equality in the expression above is that the (allocation) mass from points higher than $n + i$ may be moved to lower points causing an increase in the value of $a_i(\cdot)$ there. We prove in [Theorem D.8](#) that an arbitrary mechanism can be converted to a clean mechanism that is ‘intimately’ connected to the original mechanism. We state a direct corollary here that is obtained by setting $k = n$ and $x_i = n + i$ for all i in the statement of the theorem.

Corollary D.5 (Cleaning). *Consider any mechanism M_A for the FedEx instance described in [subsection D.1](#). There exists a mechanism M_B such that:*

- M_B mimics M_A on day 1.
- M_B is clean.
- For all days i , we have $B_i \succeq A_i$.
- For all days i , $\sum_{j=0}^{v_{\max}} R_i(j)b_i(j) \geq \sum_{j=0}^{v_{\max}} R_i(j)a_i(j)$.

Here, \succeq denotes stochastic dominance of the second-order.

D.4 The statement and proof of **Theorem D.8** ¹³

Consider any mechanism M . Let $A_i(v) : \{0, 1, \dots, V\} \rightarrow [0, 1]$ denote the allocation function of mechanism M on day i . We know from the IC constraints that $A_i(v)$ is monotone non-decreasing for all i . Without loss of generality, assume that $A_i(V) = 1$. We interpret A_i as the cumulative distribution function of some random variable. The leftwards IC constraints say that for all $i > 1, v \leq V$, we have,

$$\sum_{j=0}^v A_i(j) \geq \sum_{j=0}^v A_{i-1}(j). \quad (10)$$

We define stochastic dominance and state a well-known result

Definition D.2. Consider two discrete random variables X and Y supported on the non-negative integers. (F, G are the cdfs of the random variables). We say X has (second-order) stochastic dominance over Y if and only if for all $x \in \mathbb{Z}$,

$$\sum_{j=0}^x F(j) \leq \sum_{j=0}^x G(j).$$

We denote this by $X \succeq Y$.

We have the following well-known result on necessary and sufficient conditions for second-order stochastic dominance:

Theorem D.6. Let X_A and X_B be two random variables with distributions A and B respectively. The following statements are equivalent:

- $X_A \succeq X_B$.
- There exist random variables Y and Z such that $X_B \stackrel{d}{=} X_A + Y + Z$, with Y always at most 0 and Z such that $\mathbb{E}[Z | X_A + Y] = 0$.
- For any concave and increasing function f , it holds that $\sum_{j=0}^V f(j)a(j) \geq \sum_{j=0}^V f(j)b(j)$.

Here, $\stackrel{d}{=}$ denotes equivalence in distribution.

We proceed to prove our main technical lemma.

Lemma D.7. If any two random variables A and C defined on $[0, V] \cap \mathbb{Z}$ satisfy $A \succeq C$, then for any x , there exists a B such that $A \succeq B \succeq C$ and for all $v \leq x$, we have $\Pr(B = v) \geq \Pr(A = v)$.

Furthermore, if f is a function that is increasing on $[0, x]$ such that $f(j) = \tilde{f}(j)$ for all $0 \leq j \leq x$, then $\sum_{j=0}^V f(j) \Pr(B = j) \geq \sum_{j=0}^V \tilde{f}(j) \Pr(C = j)$.

Proof. **Theorem D.6** promises random variables Y and Z such that $C \stackrel{d}{=} A + Y + Z$, where $Y \leq 0$ and $\mathbb{E}[Z | A + Y] = 0$. Our goal is to construct a random variable B such that $A \succeq B \succeq C$. We wish to establish both the stochastic dominance results via **Theorem D.6**. To this end, we first construct random variables Y_1 and Z_1 and use these to define B . Let

$$Y_1 = \begin{cases} 0, & \text{if } A < x \\ \max(Y, x - A), & \text{if } A \geq x. \end{cases} \quad (11)$$

¹³Auxiliary results that can be skipped without loss of comprehension.

Note that $Y_1 \leq 0$. Next, define Z_1

$$Z_1 = \begin{cases} Z, & \text{if } x - A < Y_1 \\ 0, & \text{if } x - A \geq Y_1. \end{cases} \quad (12)$$

We prove that $\mathbb{E}[Z_1 | A + Y_1] = 0$ in [D.9](#) paving the way for [Theorem D.6](#). For that, we need another result.

Claim D.8. *For any $k > x$, the event $A + Y_1 = k$ happens if and only if $A + Y = k$.*

Proof. First, assume $A + Y_1 = k > x$. This means that $A \geq A + Y_1 = k > x$ and hence, by [Equation 11](#), $Y_1 = \max(Y, x - A)$. However, $Y_1 > x - A$ implying that $Y_1 = Y$.

Second, assume that $A + Y = k$. Again, we have $A \geq A + Y = k > x$ and hence, by [Equation 11](#), $Y_1 = \max(Y, x - A) = Y$. \square

Claim D.9. $\mathbb{E}[Z_1 | A + Y_1] = 0$.

Proof. Fix $A + Y_1 = k$ where k is arbitrary. If $k \leq x$, then $Z_1 = 0$ and the claim follows.

If, on the other hand, $k > x$, then [D.8](#) applies. Thus, by [Equation 12](#), we have

$$\mathbb{E}[Z_1 | A + Y_1 = k] = \mathbb{E}[Z | A + Y = k] = 0. \quad \square$$

Define $B = A + Y_1 + Z_1$. Then, [D.9](#) means that $A \succeq B$. Now, we wish to prove that $B \succeq C$. Our strategy is to again employ [Theorem D.6](#). For this, we define Y_2 and Z_2 .

$$Y_2 = Y - Y_1. \quad (13)$$

Note that

Claim D.10. $Y_2 \leq 0$.

Proof. We prove that $Y \leq Y_1$. If $A < x$, the $Y_1 = 0$ and this is trivial. Otherwise, $Y_1 = \max(Y, x - A) \geq Y$. \square

We next define Z_2 as

$$Z_2 = Z - Z_1. \quad (14)$$

Claim D.11. $\mathbb{E}[Z_2 | B + Y_2] = 0$.

Proof. Fix $B + Y_2 = k$ where k is arbitrary. Since $B = A + Y_1 + Z_1$, this is equivalent to saying that $A + Y + Z_1 = k$.

We prove that $\mathbb{E}[Z_2 | B + Y_2] = 0$ by arguing $\mathbb{E}[Z_2 | B + Y_2, x - A < Y_1] = 0$ and $\mathbb{E}[Z_2 | B + Y_2, x - A \geq Y_1] = 0$. We first deal with the former. If $x - A < Y_1$, then $Z_1 = Z$ ([Equation 12](#)) and hence $Z_2 = 0$ and we are done.

Now, suppose $x - A \geq Y_1$. In this case, $Z_1 = 0$ ([Equation 12](#)) and hence $B + Y_2 = k$ if and only if $A + Y = k$.

$$\mathbb{E}[Z_2 | B + Y_2 = k, x - A \geq Y_1] = \mathbb{E}[Z | A + Y = k, x - A \geq Y_1].$$

Note that $x - A \geq Y_1$ is the same as $x \geq A + Y_1$ which happens if and only if $x \geq A + Y = k$ ([D.8](#)). Thus,

$$\mathbb{E}[Z | A + Y = k, x - A \geq Y_1] = \mathbb{E}[Z | A + Y = k, x \geq k] = 0. \quad \square$$

Since $B + Y_2 + Z_2 = A + Y_1 + Z_1 + Y_2 + Z_2 = A + Y + Z = C$ by Equation 13 and Equation 14, we have that $B \succeq C$ by D.10 and D.11.

Consider the event $A = v$ for $v < x$. Since $v \leq x$, both Y_1 and Z_1 are 0. Thus, $B = v$. Hence, $\Pr(B = v) \geq \Pr(A = v)$.

Finally, consider a function f such that $f(j) = \tilde{f}(j)$ for all $0 \leq j \leq x$. We wish to prove that $\sum_{j=0}^V f(j) \Pr(C = j) \leq \sum_{j=0}^V f(j) \Pr(B = j)$. We break the proof into 2 cases. Each one of the following claims handles one such case.

Claim D.12. *Let E_1 be the event $A + Y > x$. We have $\sum_{j=0}^V f(j) \Pr(C = j \mid E_1) \leq \sum_{j=0}^V f(j) \Pr(B = j \mid E_1)$*

Proof. By D.8, E_1 can equivalently be described as $A + Y_1 > x$. We use these two forms interchangeably in this proof. If E_1 occurs, $Y_1 = Y$ and $Z_1 = Z$ (Equation 11 and Equation 12). Thus, $B = C$ and the claim follows. \square

Claim D.13. *Let E_2 be the event $A + Y \leq x$. We have $\sum_{j=0}^V f(j) \Pr(C = j \mid E_2) \leq \sum_{j=0}^V f(j) \Pr(B = j \mid E_2)$*

Proof. By D.8, E_2 can equivalently be described as $A + Y_1 \leq x$. We use these two forms interchangeably in this proof. If E_2 occurs, $Z_1 = 0$ (Equation 12). Thus, $B = A + Y_1 + Z_1 = A + Y_1$ implying $x \geq B \geq A + Y$. Thus, for all $j_1 \leq x$, we have

$$\sum_{j=0}^V f(j) \Pr(B = j \mid A + Y = j_1) = \sum_{j=j_1}^x f(j) \Pr(B = j \mid A + Y = j_1) \geq f(j_1).$$

We have

$$\begin{aligned} \sum_{j=0}^V f(j) \Pr(C = j \mid E_2) &\leq \sum_{j=0}^V \tilde{f}(j) \Pr(C = j \mid E_2) \\ &\leq \sum_{j_1=0}^x \Pr(A + Y = j_1 \mid E_2) \sum_{j=0}^V \tilde{f}(j) \Pr(C = j \mid A + Y = j_1) \\ &\leq \sum_{j_1=0}^x \Pr(A + Y = j_1 \mid E_2) \tilde{f} \left(\sum_{j=0}^V j \Pr(C = j \mid A + Y = j_1) \right) \\ &= \sum_{j_1=0}^x \Pr(A + Y = j_1 \mid E_2) \tilde{f}(j_1) \\ &\leq \sum_{j_1=0}^x \Pr(A + Y = j_1 \mid E_2) \sum_{j=0}^V \tilde{f}(j) \Pr(B = j \mid A + Y = j_1) \end{aligned}$$

If $A + Y = j_1 \leq x$, we have $B \in [j_1, x]$ and hence,

$$\begin{aligned} \sum_{j=0}^V f(j) \Pr(C = j \mid E_2) &\leq \sum_{j_1=0}^x \Pr(A + Y = j_1 \mid E_2) \sum_{j=0}^V \tilde{f}(j) \Pr(B = j \mid A + Y = j_1) \\ &= \sum_{j_1=0}^x \Pr(A + Y = j_1 \mid E_2) \sum_{j=0}^V f(j) \Pr(B = j \mid A + Y = j_1) \end{aligned}$$

$$= \sum_{j=0}^V f(j) \Pr(B = j \mid E_2)$$

□

□

We are now ready to prove our main “cleaning” result:

Theorem D.8 (Cleaning). *Consider any mechanism M_A for any FedEx instance with type space $[0, V] \times \{1, 2, \dots, n\}$. Let $\{x_i\}_{i=1}^{n-1}$ be a sequence of numbers in $[0, V]$. For any $k \geq 1$, there exists a mechanism M_B such that:*

- M_B mimics M_A on day 1.
- For all $1 \leq i \leq k-1$ and $v \leq x_i$, we have $b_{i+1}(v) \geq b_i(v)$.
- For all $1 \leq i \leq k$, we have $B_i \succeq A_i$ and for all $i > k$ we have $B_i = A_i$.
- For all $1 < i \leq k$, if the revenue curve R_i on day i of the bidder is increasing on $[0, x_{i-1}]$ such that $R_i(j) = \tilde{R}_i(j)$ for all $0 \leq j \leq x_{i-1}$, then $\sum_{j=0}^V R_i(j)b_i(j) \geq \sum_{j=0}^V R_i(j)a_i(j)$.

If it were not for the last two bullets, **Theorem D.8** would be trivial. We could have set M_B to be the optimal mechanism given the allocation of M_A on day 1. That the optimal mechanism is clean is an easy consequence of Fiat et al’s algorithm [FGKK16]. Intuitively, these bullets say that the curves A_i can be obtained from B_i by ‘splitting’ and ‘lowering’ prices appropriately and that these actions do not decrease the revenue generated. This structure ensures that we can convert results for the mechanism M_B into results for M_A .

This is exactly how we proceed. We take an arbitrary auction and clean it. We prove that a clean auction that generates high revenue must have a high menu complexity. We then translate this into a result for the original auction.

Proof of Theorem D.8. Proof by induction on k . If $k = 1$, we set $M_B = M_A$ and observe that all the conditions are satisfied. We assume the result holds for $k-1$ and prove it for $k > 1$. Let M_C be the mechanism promised by the induction hypothesis. Since M_C is feasible, we have

$$C_{k-1} \succeq C_k.$$

Applying **Lemma D.7** on $C_{k-1}, C_k = A_k$ and x_{k-1} , we get B_k such that

$$C_{k-1} \succeq B_k \succeq A_k, \tag{15}$$

and for all $v \leq x_{k-1}$, we have $b_k(v) \geq c_{k-1}(v)$. Furthermore, if f is a function that is increasing on $[0, x_{k-1}]$ such that $f(j) = \tilde{f}(j)$ for all $0 \leq j \leq x_{k-1}$, then $\sum_{j=0}^V f(j)b_k(j) \geq \sum_{j=0}^V f(j)a_k(j)$. Define the mechanism M_B to be M_C with the allocation C_k on day k replaced by B_k . **Equation 15** implies that M_B is feasible.

Observe that M_B satisfies all requirements.

□

D.5 Analysis of clean auctions

If f_i is the marginal density of bidder types on day i , then the revenue generated by the allocation density a_i is equal to $\sum_{x=0}^{v_{\max}} f_i(x)p_i(x)$ where $p_i(x)$ is the payment made by the bidder that reports type (x, i) . Due to Myerson's payment identity [Mye81], we know that $p_i(x) = \sum_{j=0}^x ja_i(j)$. Using this expression, the revenue generated by the allocation density a_i on day i is equal to

$$\begin{aligned} \sum_{x=0}^{v_{\max}} f_i(x)p_i(x) &= \sum_{x=0}^{v_{\max}} \sum_{j=0}^x ja_i(j)f_i(x) \\ &= \sum_{j=0}^{v_{\max}} ja_i(j)(1 - F(j-1)) = \sum_{j=0}^{v_{\max}} a_i(j)R_i(j). \end{aligned}$$

We will use this expression frequently. Lemma 2.1 says that given the allocation density a_i , the revenue generated by the optimal mechanism on days i through n is $\sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j)$. Of this revenue, an amount equal to $\sum_{j=0}^{v_{\max}} a_i(j)R_i(j)$ is generated on day i and the remainder $\sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j) - \sum_{j=0}^{v_{\max}} a_i(j)R_i(j)$ is generated on days $i+1$ and onwards. By 2.3, this is equal to $\sum_{j=0}^{v_{\max}} a_i(j)\tilde{R}_{\geq i+1}(\min(j, 3n+i+1))$.

D.5.1 Analyzing clean auctions

We described the optimal auction for our instance. It has the $i-1$ menu options $n+2, n+3, \dots, n+i-1$ and $3n+i-1$ on day $i-1$. On day i , the menu option $3n+i-1$ is split into $n+i$ and $3n+i$ while all the others are carried over. The first lemma we prove tries to formalize that this split is necessary. We intend to use this lemma in our analysis of clean auctions. By definition, a clean auction carries all prices in the range $[0, n+i-1]$ on day $i-1$ over to day i . We prove that if a price between $n+i$ and $3n+i-1/2$ is split without creating a new menu option of $n+i$, a significant revenue loss is incurred. This loss is the difference between what is actually obtained and what what could have been obtained.

If a price of $x < 3n+i$ is set on day $i-1$, then the maximum revenue that can be generated day i onwards is $\tilde{R}_{\geq i}(x)$. The revenue generated by an allocation a_i is at most $\sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j)$. We prove that if an allocation is feasible (this translates to $\sum_{j=0}^{v_{\max}} ja_i(j) \leq x$) but $a_i(n+i) = 0$, then the difference between the two terms above is large.

Lemma D.9. *Fix any $1 \leq i \leq n$ and $n+i \leq x \leq 3n+i-1/2$. For any allocation density a_i , if $a_i(n+i) = 0$ and $\sum_{j=0}^{v_{\max}} ja_i(j) \leq x$, it holds that*

$$\tilde{R}_{\geq i}(x) - \sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j) \geq \frac{n+1-i}{10(2n+1)}.$$

Proof. Consider the line joining the points $(n+i-1, \tilde{R}_{\geq i}(n+i-1))$ and $(3n+i, \tilde{R}_{\geq i}(3n+i))$. Any point on this line satisfies $y = \ell(x)$ for a linear function ℓ . Define the function $S(x) = \min(\ell(x), \tilde{R}_{\geq i}(x))$. The function S is the minimum of two concave functions and hence is concave. Further, $S(x) \geq R_{\geq i}(x)$ on all points $x \in \{0, 1, 2, \dots, 5n\} \setminus \{n+i\}$.¹⁴ Using these facts and Jensen's inequality,

$$\sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j) \leq \sum_{j=0}^{v_{\max}} a_i(j)S(j) \leq S\left(\sum_{j=0}^{v_{\max}} a_i(j)j\right).$$

¹⁴This can be verified by the interested reader. We omit the calculations.

We observe that the function S is increasing below $3n + i$. Thus,

$$\sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j) \leq S\left(\sum_{j=0}^{v_{\max}} a_i(j)j\right) \leq S(x).$$

By direct calculation, we have that

$$\begin{aligned} \sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j) &\leq S(x) \leq \ell(x) \\ &= \frac{3n+i-x}{2n+1}\tilde{R}_{\geq i}(n+i-1) + \frac{x+1-n-i}{2n+1}\tilde{R}_{\geq i}(3n+i) \\ &= \frac{3n+i-x}{2n+1}(n+1-i)S_{i-1} + \frac{x+1-n-i}{2n+1}(n+1-i)(n+S_i) \quad (\text{Equation 9}) \\ &= (n+1-i)\left(C(x-n-i) + S_i + \frac{3n+i-x}{2n+1}\left(\frac{1}{2} - \frac{1}{\lambda^{i-1}}\right)\right). \end{aligned}$$

We also have that

$$\tilde{R}_{\geq i}(x) = (n+1-i)(C(x-n-i) + S_i). \quad (\text{Equation 9})$$

Combining,

$$\begin{aligned} \sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j) &\leq \tilde{R}_{\geq i}(x) - \frac{n+1-i}{2n+1}(3n+i-x)\left(\frac{1}{\lambda^{i-1}} - \frac{1}{2}\right) \\ &\leq \tilde{R}_{\geq i}(x) - \frac{n+1-i}{2(2n+1)}\left(\frac{1}{\lambda^{i-1}} - \frac{1}{2}\right) \\ &< \tilde{R}_{\geq i}(x) - \frac{n+1-i}{10(2n+1)}. \end{aligned} \quad (\text{Claim D.1})$$

□

A corollary of the above result takes care of the case where $a_i(n+i) \neq 0$ but is small.

Corollary D.10. *Fix any $1 \leq i \leq n$ and $n+i \leq x \leq 3n+i-3/4$. For any allocation density a_i , if $a_i(n+i) < \frac{1}{100n}$ and $\sum_{j=0}^{v_{\max}} ja_i(j) \leq x$, it holds that*

$$\tilde{R}_{\geq i}(x) - \sum_{j=0}^{v_{\max}} a_i(j)R_{\geq i}(j) \geq \frac{n+1-i}{20(2n+1)}.$$

Proof. Construct the allocation

$$a'_i(x) = \begin{cases} 0, & \text{if } x = n+i \\ \frac{a_i(x)}{1-a_i(n+i)}, & \text{if } x \neq n+i. \end{cases}$$

Note that:

$$(1 - a_i(n + i)) \left(\sum_{j=0}^{v_{\max}} j a'_i(j) - (n + i) \right) = \sum_{j=0}^{v_{\max}} j a_i(j) - (n + i). \quad (16)$$

Set $y = \max \left(\sum_{j=0}^{v_{\max}} j a'_i(j), n + i \right) \geq n + i$. Note that $y \leq 3n + i - 1/2$ because:

$$\begin{aligned} \sum_{j=0}^{v_{\max}} j a'_i(j) &= \frac{1}{1 - a_i(n + i)} \left(\sum_{j=0}^{v_{\max}} j a_i(j) - (n + i) \right) + (n + i) \\ &\leq \frac{100n}{100n - 1} (2n - 3/4) + (n + i) \leq 3n + i - 1/2. \end{aligned}$$

Thus, [Lemma D.9](#) holds for y, a'_i and we have

$$\begin{aligned} &\sum_{j=0}^{v_{\max}} a_i(j) R_{\geq i}(j) - R_{\geq i}(n + i) \\ &= (1 - a_i(n + i)) \left(\sum_{j=0}^{v_{\max}} a'_i(j) R_{\geq i}(j) - R_{\geq i}(n + i) \right) \\ &\leq (1 - a_i(n + i)) \left(\tilde{R}_{\geq i}(y) - R_{\geq i}(n + i) - \frac{n + 1 - i}{10(2n + 1)} \right) \quad (\text{Lemma D.9}) \\ &= (1 - a_i(n + i)) \left((n + 1 - i)C(y - (n + i)) - \frac{n + 1 - i}{10(2n + 1)} \right) \quad (\text{Equation 9}) \\ &\leq (n + 1 - i)C(x - (n + i)) - (1 - a_i(n + i)) \frac{n + 1 - i}{10(2n + 1)} \quad (\text{Equation 16}) \\ &= \tilde{R}_{\geq i}(x) - R_{\geq i}(n + i) - (1 - a_i(n + i)) \frac{n + 1 - i}{10(2n + 1)}. \quad (\text{Equation 9}) \end{aligned}$$

Adding $R_{\geq i}(n + i)$ to both sides, we get the result. \square

[Corollary D.10](#) is our final result about the prices that split. These prices (in our instance) are in the range $[n + i, 3n + i]$ on day i . We next state results concerning prices outside this range. First, we deal with points $x > 3n + i$ such that $a_i(x) > 0$. The next pair of lemmas we prove upper bound the ‘allocation mass’ on this tail.

Lemma D.11. *For any $1 \leq i \leq n$ and $x > 3n + i$, $R_{\geq i}(3n + i) - R_{\geq i}(x) \geq n$.*

Proof. Proof by calculation.

$$\begin{aligned} R_{\geq i}(3n + i) - R_{\geq i}(x) &\geq \beta_i(3n + i) + (n - i)(C(2n - 1) + S_i) - (n - i)(2nC + S_i) \quad (\text{Equation 8}) \\ &= \beta_i(3n + i) - (n - i)C \\ &> \frac{3n + i}{2} - \frac{n - i}{2} \quad (\text{Claim D.3}) \\ &= n + i > n. \end{aligned}$$

\square

The next lemma shows that any approximately optimal mechanism for our instance must allocate very little mass to the tail end of the distribution.

Lemma D.12. *Consider any mechanism M for the FedEx instance described in Section D.1. If M generates at least a $1 - \frac{1}{1000n^2}$ fraction of the optimal revenue OPT , then on all days $1 \leq i \leq n$, the allocation curve A_i of the mechanism M satisfies,*

$$A_i(3n + i) \geq 1 - \frac{1}{200n}.$$

Proof. The statement can be verified for $i = 1$ and so we assume $i > 1$. Consider the allocation curve A_{i-1} on the $(i - 1)$ th deadline. Given A_{i-1} , the optimal revenue that can be generated on days i through n is $\sum_{j=0}^{v_{\max}} a_{i-1}(j) \tilde{R}_{\geq i}(\min(j, 3n + i))$. Since $A_{i-1} \succeq A_i$, we know by the third bullet in [Theorem D.6](#) that:

$$\sum_{j=0}^{v_{\max}} a_{i-1}(j) \tilde{R}_{\geq i}(\min(j, 3n + i)) \geq \sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i}(\min(j, 3n + i)).$$

The actual revenue generated on days i through n is at most $\sum_{j=0}^{v_{\max}} a_i(j) R_{\geq i}(j)$. Thus, allocating according to a_i on day i incurs a revenue loss of at least $\sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i}(\min(j, 3n + i)) - \sum_{j=0}^{v_{\max}} a_i(j) R_{\geq i}(j)$ on days i through n . This is at least:

$$\begin{aligned} \sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i}(\min(j, 3n + i)) - \sum_{j=0}^{v_{\max}} a_i(j) R_{\geq i}(j) & \\ & \geq \sum_{j=3n+i+1}^{v_{\max}} a_i(j) \left(\tilde{R}_{\geq i}(3n + i) - R_{\geq i}(j) \right) \quad (\text{As } \tilde{R}_{\geq i}(j) \geq R_{\geq i}(j)) \\ & \geq n \sum_{3n+i+1}^{v_{\max}} a_i(x) dx = n(1 - A_i(3n + i)). \quad (\text{Lemma D.11}) \end{aligned}$$

For a $1 - \frac{1}{1000n^2}$ approximate auction, the loss in revenue is upper bounded by $\frac{1}{200}$. Thus, $n(1 - A_i(3n + i)) \leq \frac{1}{200}$ implying $A_i(3n + i) \geq 1 - \frac{1}{200n}$. \square

Finally, we prove some technical results for allocating in the range $[0, n + i]$.

Lemma D.13. *For any clean mechanism M that generates revenue at least $\frac{99}{100}$ times OPT ,*

$$A_{n/3}(4n/3) \leq \frac{1}{2}.$$

Proof. The Myerson optimal revenue for day i is $\max_x R_i(x) = \beta_i(3n + i)$ by [Equation 7](#). We upper bound the revenue generated by M on days 1 through $n/3 - 1$ by $\sum_{i=1}^{n/3-1} \beta_i(3n + i)$. We have

$$\begin{aligned} \sum_{i=1}^{n/3-1} \beta_i(3n + i) &= \sum_{i=1}^{n/3} \frac{3n - i}{2} - \frac{n - i}{\lambda^i} + S_i && (\text{Equation 5}) \\ &\leq \sum_{i=1}^{n/3-1} \frac{3n - i}{2} - \frac{3(n - i)}{4} + n + i && (\text{D.1}) \\ &\leq \sum_{i=1}^{n/3-1} \frac{7n + 5i}{4} < \frac{3n^2}{4} - \frac{9n}{4}. \end{aligned}$$

The revenue generated on day $n/3$ onwards is upper bounded by $\sum_{j=0}^{v_{\max}} a_{n/3}(j)R_{\geq n/3}(j)$. Suppose for the sake of contradiction that $A_{n/3}(4n/3) > \frac{1}{2}$. If this is the case, then $\sum_{j=0}^{v_{\max}} a_{n/3}(j)R_{\geq n/3}(j)$ is upper-bounded by

$$\begin{aligned} \sum_{j=0}^{v_{\max}} a_{n/3}(j)R_{\geq n/3}(j) &\leq A_{n/3}(4n/3)R_{\geq n/3}(4n/3) + (1 - A_{n/3}(4n/3))R_{\geq n/3}(10n/3) \\ &\leq R_{\geq n/3}(10n/3) - A_{n/3}(4n/3) (R_{\geq n/3}(10n/3) - R_{\geq n/3}(4n/3)) \\ &\leq \frac{11n^2}{9} + \frac{11n}{6}. \end{aligned}$$

The sum of these two contributions is at most $\frac{3n^2}{4} + \frac{11n^2}{9} = \frac{71n^2}{36} \leq \frac{71}{72}\text{OPT}$, a contradiction. \square

Lemma D.14. *Let a_i and a_{i+1} be allocations that satisfy $\sum_{j=0}^x A_{i+1}(j) \geq \sum_{j=0}^x A_i(j)$ for all x . Further, let v_i, v_{i+1} be such that $A_i(v_i) = A_{i+1}(v_{i+1}) = \alpha$. Then,*

$$\sum_{j=0}^{v_{i+1}} ja_{i+1}(j) \leq \sum_{j=0}^{v_i} ja_i(j).$$

Proof. First suppose that $v_i < v_{i+1}$. We have:

$$\begin{aligned} \sum_{j=0}^{v_{i+1}} ja_{i+1}(j) - \sum_{j=0}^{v_i} ja_i(j) &= \alpha(v_{i+1} - v_i) - \sum_{j=0}^{v_{i+1}} A_{i+1}(j) + \sum_{j=0}^{v_i} A_i(j) \\ &\leq \alpha(v_{i+1} - v_i) - \sum_{j=v_i+1}^{v_{i+1}} A_i(j) \\ &\leq \alpha(v_{i+1} - v_i) - \sum_{j=v_i+1}^{v_{i+1}} \alpha = 0. \end{aligned}$$

If $v_i \geq v_{i+1}$. We have:

$$\begin{aligned} \sum_{j=0}^{v_{i+1}} ja_{i+1}(j) - \sum_{j=0}^{v_i} ja_i(j) &= \alpha(v_{i+1} - v_i) - \sum_{j=0}^{v_{i+1}} A_{i+1}(j) + \sum_{j=0}^{v_i} A_i(j) \\ &\leq \alpha(v_{i+1} - v_i) + \sum_{j=v_{i+1}+1}^{v_i} A_i(j) \\ &\leq \alpha(v_{i+1} - v_i) + \sum_{j=v_{i+1}+1}^{v_i} \alpha = 0. \end{aligned}$$

\square

We are now ready to prove a menu complexity lower bound for clean auctions.

Theorem D.15. *Consider any clean mechanism M for the FedEx instance described in [subsection D.1](#) that is $1 - \frac{1}{1000n^2}$ approximate. For any day i such that $n/4 \leq i < n/2$, the mechanism M satisfies $a_i(j) \geq \frac{1}{300n}$ for all $j \in [n+2, n+n/4]$.*

Proof. Consider any day $i < n/4$ and the pair a_i and a_{i+1} of allocation densities of the mechanism M on day i and day $i + 1$ respectively. Given a_i , the optimal revenue that can be generated on days $i + 1$ through n is $\sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n + i + 1))$. The actual revenue generated on days $i + 1$ through n is at most $\sum_{j=0}^{v_{\max}} a_{i+1}(j) R_{\geq i+1}(j)$. Thus, allocating according to a_{i+1} on day $i + 1$ incurs a revenue loss of at least $\sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n + i + 1)) - \sum_{j=0}^{v_{\max}} a_{i+1}(j) R_{\geq i+1}(j)$ on days $i + 1$ through n .

Let v^* be such that $A_i(3n + i) = A_{i+1}(v^*)$. We define two allocations b_i and b_{i+1} ¹⁵.

$$b_i(v) = \begin{cases} 0, & v \leq n + i \\ Z_1 a_i(v), & n + i < v \leq 3n + i \\ 0, & 3n + i < v, \end{cases} \quad b_{i+1}(v) = \begin{cases} Z_2(a_{i+1}(v) - a_i(v)), & v \leq n + i \\ Z_2 a_{i+1}(v), & n + i < v \leq v^* \\ 0, & v^* < v, \end{cases}$$

where Z_1, Z_2 are the normalization constants. Since the mechanism M is clean, we have $a_{i+1}(v) \geq a_i(v)$ for all $v \leq n + i$ by [D.1](#) and therefore b_i and b_{i+1} are valid allocation densities.

Claim D.14. $Z_1 = Z_2 \leq 3$.

Proof. We first prove that $Z_1 = Z_2$. This is because

$$\sum_{j=0}^{v_{\max}} \frac{b_i(j)}{Z_1} = A_i(3n + i) - A_i(n + i) = A_{i+1}(v^*) - A_i(n + i) = \sum_{j=0}^{v_{\max}} \frac{b_{i+1}(j)}{Z_2}.$$

We now prove that $Z_1 < 3$. Since the mechanism is clean and $i < n/4 < n/3$, any ‘allocation mass’ on points $[0, n + i]$ is pushed forward to all days. Thus, $A_i(n + i) \leq A_{n/3}(4n/3) \leq \frac{1}{2}$ by [Lemma D.13](#). Also, by [Lemma D.12](#), $A_i(3n + i) \geq 1 - \frac{1}{200n}$. Thus,

$$Z_1 = \frac{1}{A_i(3n + i) - A_i(n + i)} \leq \frac{1}{1/2 - \frac{1}{200n}} < 3.$$

□

For the remainder of this proof, we set Z denote the common value of Z_1 and Z_2 . We are now ready to lower bound the loss in revenue caused by allocating according to a_{i+1} . Since $\tilde{R}_{\geq i+1} = R_{\geq i+1}$ for all $j \leq n + i$, we get

$$\begin{aligned} & \sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n + i + 1)) - \sum_{j=0}^{v_{\max}} a_{i+1}(j) R_{\geq i+1}(j) \\ & \geq \frac{1}{Z} \left(\sum_{j=n+i+1}^{3n+i} b_i(j) \tilde{R}_{\geq i+1}(j) - \sum_{j=0}^{v^*} b_{i+1}(j) R_{\geq i+1}(j) \right) \\ & = \frac{1}{Z} \left(\tilde{R}_{\geq i+1} \left(\sum_{j=n+i+1}^{3n+i} j b_i(j) \right) - \sum_{j=0}^{v^*} b_{i+1}(j) R_{\geq i+1}(j) \right) \quad (\tilde{R}_{\geq i+1} \text{ is linear in } [n + i + 1, 3n + i + 1]) \end{aligned} \tag{17}$$

We next plan to upper bound $\sum_{j=0}^{v^*} b_{i+1}(j) R_{\geq i+1}(j)$ using [Corollary D.10](#). To this end, fix $y = \max\left(n + i + 1, \sum_{j=n+i+1}^{3n+i} j b_i(j)\right)$. Note that $n + i + 1 \leq y \leq 3n + i$. We first show that $y \geq \sum_{j=0}^{v^*} j b_{i+1}(j) \geq \sum_{j=0}^{v^*} j b_i(j)$. Adding $\sum_{j=0}^{n+i} Z a_i(v)$ on both sides, we get that this is

¹⁵In this definition we implicitly assume that $n + i \leq v^*$. This holds because $A_{i+1}(n + i) \leq \frac{1}{2}$ and $A_{i+1}(v^*) = A_i(3n + i) \geq 1 - \frac{1}{200n}$ by [Lemma D.12](#). Since the mechanism is clean and $i < n/4 < n/3$, any ‘allocation mass’ on day $i + 1$ on points $[0, n + i + 1]$ is pushed forward to all days. Thus, $A_{i+1}(n + i) \leq A_{i+1}(n + i + 1) \leq A_{n/3}(4n/3) \leq \frac{1}{2}$ by [Lemma D.13](#).

equivalent to showing that $\sum_{j=0}^{3n+i} Zj a_i(j) \geq \sum_{j=0}^{v^*} Zj a_{i+1}(j)$. This holds due to [Lemma D.14](#). Now suppose $b_{i+1}(n+i+1) < \frac{1}{100n}$. If this is the case, apply [Corollary D.10](#) to the last term in [Equation 17](#) to get

$$\begin{aligned} & \sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n+i+1)) - \sum_{j=0}^{v_{\max}} a_{i+1}(j) R_{\geq i+1}(j) \\ & \geq \frac{1}{Z} \left(\tilde{R}_{\geq i+1} \left(\sum_{j=n+i}^{3n+i} j b_i(j) \right) - R_{\geq i+1}(y) + \frac{n-i}{20(2n+1)} \right) \\ & \geq \frac{1}{3} \frac{n-i}{20(2n+1)} > \frac{1}{240}. \end{aligned}$$

where the last inequality is because $y \geq \sum_{j=0}^{3n+i} j b_i(j)$ and $\tilde{R}_{\geq i}$ is increasing below $3n+i$.

Since $\frac{1}{240} > \frac{\text{OPT}}{1000n^2}$, this is unaffordable. Thus, $b_{i+1}(n+i+1) \geq \frac{1}{100n}$ or $a_{i+1}(n+i+1) \geq \frac{1}{Z100n} > \frac{1}{300n}$. Since the mechanism is clean, this is pushed forward to all days implying the result. \square

D.5.2 Analyzing general auctions

[Theorem D.15](#) proves an $\Omega(n^2)$ -menu complexity lower bound for clean auctions. In this section, we wish to extend this lower bound to general auctions. Our proof will rely heavily on [Corollary D.5](#) which says that any auction can be viewed as the ‘muddled’ version of a clean auction. More precisely, the allocation curves A_i of any mechanism are stochastically dominated by clean curves B_i . We prove that removing a lot of menu options from B_i to get A_i while maintaining $B_i \succeq A_i$ results in a huge revenue loss day $i+1$ onwards. Since the revenue generated by B_j for $j \leq i$ is at least that generated by A_j (4th bullet in [Corollary D.5](#)), the huge loss of revenue on days $i+1$ onwards is tantamount to a huge loss overall.

Proof of [Theorem 5.1](#). Consider an arbitrary mechanism M_A and let A_i be the allocation curve of M_A on day i . [Corollary D.5](#) says that there exist allocation curves B_i of a clean mechanism M_B such that $B_i \succeq A_i$ for all i . Let $i \in (n/4, n/2]$ be such that the number of menu options on day i is at most $n/8$. Let $V = \{v_l\}_{l=1}^k$ be the vector of menu options on day i . Let S be the polygon approximation of $\tilde{R}_{\geq i+1}$ defined by the points in $V \cup [n+i, v_{\max}]$. We have that the revenue generated by M day $i+1$ onwards is at most

$$\sum_{j=0}^{v_{\max}} a_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n+i+1)) = \sum_{j=0}^{v_{\max}} a_i(j) S(\min(j, 3n+i+1)).$$

Since for all i , we have $B_i \succeq A_i$, and $S(\min(3n+i+1, x))$ is increasing and concave, it holds by the 3rd bullet in [Theorem D.6](#) that

$$\sum_{j=0}^{v_{\max}} a_i(j) S(\min(j, 3n+i+1)) \leq \sum_{j=0}^{v_{\max}} b_i(j) S(\min(j, 3n+i+1)).$$

Consider the mechanism that mimics M_B on days 1 through i and does optimally thereafter. Since it mimics M_B on days 1 through i , the revenue generated on these days is at least that of M_A (4th bullet in [Corollary D.5](#)). Day $i+1$ onwards, this mechanism generates a revenue of $\sum_{j=0}^{v_{\max}} b_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n+i+1))$. To prove our result it is sufficient to prove

that $\sum_{j=0}^{v_{\max}} b_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n+i+1)) - \sum_{j=0}^{v_{\max}} b_i(j) S(\min(j, 3n+i+1)) \geq \frac{1}{60000} \geq \frac{\text{OPT}}{200000n^2}$. Observe that since the functions S and $\tilde{R}_{\geq i+1}$ are the same for all values at least $n+i$, we have

$$\sum_{j=0}^{v_{\max}} b_i(j) \tilde{R}_{\geq i+1}(\min(j, 3n+i+1)) - \sum_{j=0}^{v_{\max}} b_i(j) S(\min(j, 3n+i+1)) = \sum_{j=0}^{n+i} b_i(j) \left(\tilde{R}_{\geq i+1}(j) - S(j) \right).$$

We know that M_B has revenue at least that of M_A and is clean. Thus, by [Theorem D.15](#) for M_B , we get that b_i has a mass of at least $1/300n$ on at least $n/4$ points in the range $[n, n+i]$. The function S_i in this range is defined by $V \cap [0, n+i]$ which has at most $n/8$ points. Thus, there are at least $n/8$ points in the support of b_i that are not in $V \cap [0, n+i]$. For any such point x ,

$$\begin{aligned} S(x) &\leq \frac{1}{2} (R_{\geq i+1}(x-1) + R_{\geq i+1}(x+1)) \\ &= \frac{1}{2} ((n-i)S_{x-n-1} + (n-i)S_{x-n+1}) \quad (\text{Equation 9}) \\ &= \frac{n-i}{2} \left(2S_{x-n} + \frac{1}{\lambda^{x-n}} - \frac{1}{\lambda^{x-n-1}} \right) \\ &= (n-i)S_{x-n} - \frac{n-i}{2} \frac{\lambda-1}{\lambda^{x-n}} \\ &= R_{\geq i+1}(x) - \frac{n-1}{8n} \frac{1}{\lambda^{x-n}} \quad (\text{Equation 9}) \\ &\leq R_{\geq i+1}(x) - \frac{3(n-i)}{32n} < \frac{3}{64}. \quad (\text{D.1}) \end{aligned} \tag{18}$$

Thus, the expression $\sum_{j=0}^{n+i} b_i(j) \left(\tilde{R}_{\geq i+1}(j) - S(j) \right)$ is at least the same sum over the $n/8$ points where [Equation 18](#) holds. This sum is thus, at least $\frac{n}{8} \cdot \frac{3}{64} \cdot \frac{1}{300n} > \frac{1}{60000}$. \square

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](https://arxiv.org/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 Pappas, 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: Non-monotonicity 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.
- [Mye81] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [Rot92] G. Rote. The convergence rate of the sandwich algorithm for approximating convex functions. *Computing*, 48(3-4):337–361, March 1992.
- [Rub16] Aviad Rubinfeld. Settling the complexity of computing approximate two-player nash equilibria. In *FOCS*, 2016.
- [RW15] Aviad Rubinfeld 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.