

The FedEx Problem (Working Paper)

Amos Fiat* Kira Goldner† Anna R. Karlin‡ Elias Koutsoupias§

June 6, 2016

Abstract

“Remember that Time is Money”

— Benjamin Franklin in *Advice to a Young Tradesman* (1748)

Consider the following setting: a customer has a package and is willing to pay up to some value v to ship it, but needs it to be shipped by some deadline d . Given the joint prior distribution from which (v, d) pairs are drawn, we characterize the auction that yields optimal revenue, contributing to the limited understanding of optimal auctions beyond single-parameter settings. Our work requires a new way of combining and ironing revenue curves which illustrate why randomization is necessary to achieve optimal revenue. Finally, we strengthen the emerging understanding that duality is useful for both the design and analysis of optimal auctions in multi-parameter settings.

1 Introduction

Consider the pricing problem faced by FedEx. Each customer has a package to ship, a deadline d by which he needs his package to arrive, and a value v for a guarantee that the package will arrive by his deadline. FedEx can (and does) offer a number of different shipping options in order to extract more revenue from their customers. In this paper, we solve the optimal (revenue-maximizing) auction problem for the single-agent version of this problem. Our paper adds to the relatively short list of multi-parameter settings for which a closed-form solution is known.

This pricing problem is extremely natural and arises in numerous scenarios, whether it is Amazon.com providing shipping options, Internet Service Providers offering bandwidth plans, Bitcoin miners setting a policy for transaction fees, or a myriad of other settings in which customers have a sensitivity to time or some other feature of service. In these settings, a seller can price discriminate or otherwise segment his market by delaying service, or providing lower quality/cheaper versions of a product. It is important to understand how buyer deadline (or quality) constraints impact the design of auctions and what leverage they give to the seller to extract more revenue.

*Tel Aviv University. fiat@tau.ac.il.

†University of Washington. kgoldner@cs.washington.edu. Funded by the National Science Foundation under CCF grant 1420381.

‡University of Washington. karlin@cs.washington.edu. Funded by the National Science Foundation under CCF grant 1420381.

§University of Oxford. elias@cs.ox.ac.uk. Funded by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. 321171 (ALGAME).

This research was done in part while the authors were visiting the Simons Institute for Theoretical Computer Science.

We consider a model in which a seller provides n different options for service, and a customer is interested in buying an option that meets his quality demand of d . We use the running example of shipping packages by deadlines. A customer’s utility for getting his package shipped by day t at a price of p is $v - p$ if $t \leq d$ (i.e., it is received by his deadline) and $-p$ otherwise. A customer’s (v, d) pair is his private information. We study the Bayesian setting, where this pair (v, d) is drawn from a prior distribution known to the seller.

Related Work

The FedEx problem is a variant of price discrimination in which the customers are grouped by their deadline. Price discrimination offers different prices to users with the goal of improving revenue [5]. Alternatively one can view the FedEx problem as a multi-dimensional optimal auction problem. There are two ways to express the FedEx problem in this way. First, as a 2-dimensional (value \times deadline) problem of arbitrary joint distribution in which the second variable takes only integer values in a bounded interval. Alternatively, as a very special case of the n -dimensional unit-demand problem with correlated values (the customer buys a day among the n choices).

There is an extensive body of literature on optimal auction design. The seminal work of Myerson [33] has completely settled the case of selling a single item to multiple bidders and extends directly to the more general framework of single-parameter settings. The most complicated part of his solution is his handling of distributions that are not regular by “ironing” them, that is, by replacing the revenue curves by their upper concave envelope. Myerson’s ironing is done in quantile space. In this work, we also need to iron the revenue curves, but we need to do this in value space.

Extending Myerson’s solution to the multidimensional case has been one of the most important open problems in Microeconomics. For the case of unit-demand agents, a beautiful sequence of papers [15, 6, 16, 14, 1, 7] showed how to obtain approximately optimal auctions. For the case of finite type spaces, [9, 11, 10] are able to use linear and convex programming techniques to formulate and solve the optimal auction problem. This gives a black-box reduction from mechanism to algorithm design that yields a PTAS for revenue maximization in unit-demand settings. For the case of additive agents, additional recent breakthroughs [24, 29, 4, 38, 12, 13] have also resulted in approximately optimal mechanisms.

But if we insist on optimal auctions for continuous probability distributions, no general solution is known for the multidimensional case—even for the two-dimensional single-bidder case—and it is very possible that no such simple solution exists for the general case. One of the reasons that the multidimensional case is so complex is that optimal auctions are not necessarily deterministic [34, 37, 6, 26, 25, 31, 35, 19]. The optimal auction for the FedEx problem also turns out to be randomized with possibly exponentially many different price levels.

There are some relevant results that solve special cases of the two-parameter setting. One of the earliest works is by Laffont, Maskin, and Rochet [28], who solved a distant variant of the FedEx problem in which the utility of the bidder is expressed as a quadratic function of the two values (v, d) when the values are *uniformly* distributed in $[0, 1]$; unlike the FedEx problem the second parameter is drawn from a continuous probability distribution. Their solution was highly non-trivial, which was an early indication that the multidimensional optimal auction problem may not be easy. This work was followed by McAfee and McMillan [32] who generalized this example to a class of distributions that have the “single-crossing” property.

These initial results were followed by more general results. In particular, Haghpanah and

Hartline [23] considered the multidimensional problem of selling a product with multiple quality levels and gave sufficient conditions under which the optimal auction is to sell only the highest quality (in the FedEx problem, this corresponds to having a single price for every day). They also gave sufficient conditions under which selling the grand bundle to an agent with additive preferences is optimal. Their work generalizes results from Armstrong [3].

Our approach is based on a duality framework. Two such frameworks were proposed. The first framework by Daskalakis, Deckelbaum, and Tzamos [19, 20, 18] reduces the problem to optimal transport theory. With its use, the authors gave optimal mechanisms for a number of two-item settings and gave necessary and sufficient conditions under which grand bundle selling is optimal. The second framework is by Giannakopoulos and Koutsoupias [21, 22], which is based on expressing the problem as an optimization problem with linear partial differential inequalities, computing its dual, and using complementary slackness conditions to extract the optimal auction and prove its optimality. With its use, the authors gave the optimal auction for a large class of distributions for two items and for the uniform distribution for up to 6 items. Our solution of the FedEx problem follows the latter duality framework.

For much more on both exact and approximate optimal mechanism design, see [18, 17, 36, 27, 8]. For background on duality in infinite linear and convex programs, see e.g., [2, 30].

This Paper

Our main result is a characterization of the revenue-optimal auction for this setting. For each delivery option of 1 through n days, the mechanism specifies a distribution of prices. The customer, knowing the distributions, specifies a delivery option of i days, and then a price is drawn randomly from day i 's distribution and offered to the customer.

We formulate the optimal auction problem as a continuous infinite linear program, take its dual, and determine a sufficient set of conditions for optimality. We then show how to construct a sequence of “revenue-type” curves $\Gamma_{\geq i}(\cdot)$. Each such curve represents the optimal revenue on days i through n given a price that might be set on day i . That is, $\Gamma_{\geq i}(p)$ corresponds to the sum of (1) the revenue from selling day i delivery at a price of p to customers with a deadline of i and (2) the optimal revenue (when constrained by this choice on day i) from days $i + 1$ through n , where the constraint comes from the requirement that it is incentive compatible for a customer to report his true deadline rather than an earlier one. This curve also incorporates “ironing” so as to ensure incentive compatibility, and this ironing may lead to randomization over prices. We use these curves to construct a solution to the primal and dual linear programs that satisfies conditions sufficient for optimality.

Our result is one of relatively few exact and explicit closed-form generalizations of [33] to multi-parameter settings with arbitrary joint distributions, and contributes to recent breakthroughs in this space. Key take-aways are the following:

1. Our result requires a new way of combining and ironing revenue curves. In Myerson’s optimal auction for irregular distributions, ironing ensures incentive compatibility and gives an upper bound on the optimal revenue. Myerson shows that this upper bound is in fact achievable using randomization. Similarly, our combined and ironed curves yield upper bounds on the revenue, and we show that these upper bounds can be realized with lotteries. In Myerson’s setting, ironing is required to enforce incentive compatibility constraints among multiple bidders. In our setting, we need ironing even for one bidder because of the multiple parameters.
-

This may suggest that ironing is one of the biggest hurdles in extending Myerson’s results to more general settings.

2. The optimal auction we obtain is constructed inductively and, consequently, is relatively simple to describe. Once the distribution of prices has been determined for delivery by day i , using $\Gamma_{\geq i+1}(\cdot)$, we show how to define the distribution of prices for delivery by day $i + 1$. The latter involves randomizing over up to 2^i prices.
3. The duality approach gives a closed-form allocation rule. Naive attempts to solve this problem, even for the case where there are only two or three possible deadlines, leads to a massive case analysis depending on the priors. Even in cases where the optimal auction is deterministic, setting these prices is not straightforward. For example, for three possible deadlines, the optimal deterministic mechanism can require 1, 2, or 3 distinct prices, and determining how many prices to use and how to set them seems non-trivial.

Our duality approach, however, leads to a unified allocation rule with no case analysis at all. This paper strengthens the emerging understanding that duality is useful for determining the structure of the optimal auction in non-trivial settings and for analyzing the resulting auction.

2 Preliminaries

As discussed above, the type of a customer is a (value, deadline) pair. An auction takes as input a reported type $t = (v, d)$ and determines the shipping date in $\{1, \dots, n\}$ and the price. We denote by $a_d(v)$ the probability that the package is shipped by day d , when the agent reports (v, d) , and by $p_d(v)$ the corresponding expected price (the expectation is taken over the randomness in the mechanism).

Our goal is to design an optimal auction for this setting. By the revelation principle, we can restrict our attention to incentive-compatible mechanisms. Denote by $u(v', d' | v, d)$ the utility of the agent when his true type is (v, d) but he reports a type of (v', d') . That is,

$$u(v', d' | v, d) = \begin{cases} va_{d'}(v') - p_{d'}(v') & \text{if } d' \leq d \\ -p_{d'}(v') & \text{otherwise.} \end{cases}$$

The incentive compatibility requirement is that

$$u(v, d) := u(v, d | v, d) \geq u(v', d' | v, d) \quad \forall v', d'. \quad (1)$$

We also require individual rationality, i.e., $u(v, d) \geq 0$ for all (v, d) . Without loss of generality, $a_d(v)$ is the probability that the package is shipped *on day* d , since any incentive-compatible mechanism which ships a package early can be converted to one that always ships on the due date¹, while retaining incentive compatibility and without losing any revenue.

For each fixed d , this implies the standard (single parameter) constraints [33], namely

$$\forall d, a_d(v) \text{ is monotone weakly increasing and in } [0, 1]; \quad (2)$$

$$\forall d, p_d(v) = va_d(v) - \int_0^v a_d(x)dx \quad \text{and hence} \quad u(v, d) = \int_0^v a_d(x)dx. \quad (3)$$

¹At least that’s when the customer *thinks* it’s being shipped.

Clearly no agent would ever report $d' \geq d$, as this would result in negative utility. However, we do need to make sure that the agent has no incentive to under-report his deadline, and hence another IC constraint is that for all $2 \leq d \leq n$:

$$u(v, d-1|v, d) \leq u(v, d|v, d) \quad (4)$$

which is equivalent to

$$\int_0^v a_{d-1}(x)dx \leq \int_0^v a_d(x)dx \quad \forall d \text{ s.t. } 1 < d \leq n. \quad (5)$$

We sometimes refer to this as the inter-day IC constraint. Since $a_d(v)$ is the probability of allocation on day d , given report (v, d) , constraints (2), (3) and (5) are necessary and sufficient, by transitivity, to ensure that

$$u(v', d'|v, d) \leq u(v, d|v, d)$$

for *all* possible misreports (v', d') .

The prior

We assume that the agent's (value, deadline) comes from a known prior \mathcal{F} . Let q_i be the probability that the customer has a deadline $i \in \{1, \dots, n\}$, that is,

$$q_i = \Pr_{(v,d) \sim \mathcal{F}}[d = i]$$

and let $F_i(\cdot)$ be the marginal distribution function of values for bidders with deadline i , that is,

$$F_i(x) = \Pr_{(v,d) \sim \mathcal{F}}[v \leq x \mid d = i].$$

We assume that F_i is atomless and strictly increasing, with density function defined on $[0, H]$. Let $f_i(v)$ be the derivative of $F_i(v)$.

The objective

Let $\varphi_i(v) = v - \frac{1-F_i(v)}{f_i(v)}$ be the virtual value function for v drawn from distribution F_i . Applying the Myerson payment identity implies that the expected payment of a customer with deadline i is

$$\mathbb{E}_{v \sim F_i}[p_i(v)] = \mathbb{E}_{v \sim F_i}[\varphi_i(v)a_i(v)].$$

Thus, we wish to choose monotone allocation rules $a_i(v)$, for days $1 \leq i \leq n$, so as to maximize

$$\mathbb{E}_{(v,i) \sim \mathcal{F}}[p_i(v)] = \sum_{i=1}^n q_i \mathbb{E}_{v \sim F_i}[p_i(v)] = \sum_{i=1}^n q_i \mathbb{E}_{v \sim F_i}[\varphi_i(v)a_i(v)] = \sum_{i=1}^n q_i \int_0^H \varphi_i(v)f_i(v)a_i(v)dv,$$

subject to (2), (3) and (5).

A trivial case and discussion

Conditioned on the fact that a customer has a day d deadline, i.e., if we knew that his value for service was drawn from F_d , the optimal pricing would be trivial, since this is a single agent, single item auction. Thus, the optimal mechanism for such a customer is to set the price for service by day d to the reserve price r_d for his prior². Moreover, if it is the case that $r_d \geq r_{d+1}$ for each d , where r_d is the reserve price for the distribution F_d , the entire FedEx problem is trivial, since then we could set r_d as the price for service on day d , all IC constraints would be satisfied, and we would be optimizing pointwise for each conditional distribution.

In this paper, we do not make the assumption that the reserve prices are weakly decreasing with the deadline, let alone the stronger assumption that the distribution F_d stochastically dominates the distribution F_{d+1} for each d . There may be several reasons that these assumptions do not hold. For one, the prior \mathcal{F} captures the result of random draws from a population consisting of a mixture of different types. Obviously any particular individual with deadline d is at least as happy with day $d - 1$ service as with day d service, but two random individuals may have completely uncorrelated needs, so if one them is of type (v, d) , and the other is of type (v', d') , with $d' > d$, it is not necessarily the case that $v' \leq v$.

A second factor has to do with costs. It is likely that the cost that FedEx incurs for sending a package within d days is higher than the cost FedEx incurs for sending a package within $d' > d$ days, since in the latter case, for example, FedEx has more flexibility about which of many planes/trucks to put the package on, and even may be able to reduce the total number of plane/truck trips to a particular destination given this flexibility. More generally, in other applications of this problem, the cost of providing lower quality service is lower than the cost of providing higher quality service. Thus, even if reserve prices tend to decrease with d , all bets are off once we consider a customer's value for deadline d conditioned on that value being above the expected cost to FedEx of shipping a package by deadline d for each d .

In this paper, we are not explicitly modeling the costs that FedEx incurs, the optimization problems that it faces, the online nature of the problem, or any limits on FedEx's ability to ship packages. These are interesting problems for future research. The discussion in the preceding few paragraphs is here merely to explain why the problem remains interesting and relevant even when r_d is below r_{d+1} .

3 Warm-up: The case of $n = 2$

Suppose that the customer has a deadline of either one day or two days. By the taxation principle, the optimal mechanism is a menu, in this case a price p_i (possibly selected by some randomized procedure) for having the package shipped by (on) day i .

Let $R_i(v)$ be the *revenue curve* for day i , that is $R_i(v) := v \cdot (1 - F_i(v))$. Let $r_i := \operatorname{argmax}_v (R_i(v))$, the price at which expected revenue from a bidder with value drawn from F_i is maximized, and let $R_i^* := R_i(r_i)$ denote this maximum expected revenue. Since R_i^* is the optimal expected revenue from the agent [33], conditioned on having a deadline of i , $q_1 R_1^* + q_2 R_2^*$ is an upper bound on the optimal expected revenue for the two-day FedEx problem. If $r_1 \geq r_2$, then this optimum is indeed achievable by an IC mechanism: just set the day one shipping price p_1 to r_1 and the day two shipping price p_2 to r_2 .

² r_d is defined properly at the beginning of Section 3.

But what if $r_2 > r_1$? In this case, the inter-day IC constraint (5) is violated by this pricing (a customer with $d = 2$ will prefer to pretend his deadline is $d = 1$).

Attempt #1: One alternative is to consider the optimal single price mechanism (i.e., $p_1 = p_2 = p$). In this case, the optimal choice is clear:

$$p := \operatorname{argmax}_v [q_1 R_1(v) + q_2 R_2(v)], \quad (6)$$

i.e., set the price that maximizes the combined revenue from both days. There are cases where this is optimal, e.g., if both F_1 and F_2 are regular.³

Attempt #2: Another auction that retains incentive compatibility, and, in some cases, improves performance is to set the day one price p_1 to p and the day two price to

$$p_2 := \operatorname{argmax}_{v \leq p} R_2(v). \quad (7)$$

However, even if we fix $p_1 = p$, further optimization may be possible if F_2 is not regular.

Attempt #3: Consider the concave hull of $R_2(\cdot)$, i.e., the ironed revenue curve. If $R_2(v)$ is maximized at $r_2 > p$ and $R_2(\cdot)$ is ironed at p , then offering a lottery on day two with an expected price of p yields higher expected revenue than offering any deterministic day two price of p_2 . As we shall see, for this case, this solution is actually optimal. (See Figure 1.)

However, if $p > r_2$, (which is possible if F_1 and F_2 are not regular, even if $r_1 < r_2$), then we will see that the optimal day one price is indeed above r_2 , but not necessarily equal to p .

Attempt #4: If $p > r_2$, set the day one price to

$$p_1 := \operatorname{argmax}_{v \geq r_2} R_1(v).$$

This should make sense: if we're going to set a day one price above r_2 , we may as well set the day two price to r_2 , but in that case, the day two curve should not influence the pricing for day one (except to set a lower bound for it).

Admittedly, this sounds like a tedious case analysis, and extending this reasoning to three or more days gets much worse. Happily, though, there is a nice, and relatively simple way to put all the above elements together to describe the solution, and then, as we shall see in Section 5, prove its optimality via a clean duality proof.

A solution for $n = 2$

Define $\bar{R}(\cdot)$ to be the concave (ironed) revenue curve corresponding to revenue curve $R(\cdot)$ and let

$$R_{12}(v) := \begin{cases} q_1 R_1(v) + q_2 \bar{R}_2(v) & v \leq r_2 \\ q_1 R_1(v) + q_2 R_2(r_2) & v > r_2 \end{cases} \quad (8)$$

Note that because $\bar{R}_2(\cdot)$ is the least concave upper bound of $R_2(\cdot)$ and by definition of r_2 that $\bar{R}_2(r_2) = R_2(r_2)$. The optimal solution is to set

$$p_1 := \operatorname{argmax}_v R_{12}(v),$$

and then take

$$p_2 := r_2 \text{ if } r_2 \leq p_1 \text{ and } \mathbb{E}(p_2) := p_1 \text{ otherwise,}$$

³ A distribution F is regular if its virtual value function is increasing in v .




Figure 1: A two-day case: Suppose that the optimal thing to do on day one is to offer a price of p . In the upper left, we see the corresponding allocation curve $a_1(v)$. The bottom left graph shows the revenue curve $R_2(\cdot)$ for day two (the thin black curve) and the ironed version $\bar{R}_2(\cdot)$ (the thick blue concave curve). Optimizing for day two subject to the inter-day IC constraint $\int_0^v a_1(x)dx \leq \int_0^v a_2(x)dx$ suggests that the most revenue we can get from a deadline $d = 2$ customer is $\bar{R}_2(p)$ on day two, which can be done by offering the price of \underline{p} with probability $1/3$ and a price of \bar{p} with probability $2/3$ (since, in this example, $p = (1/3)\underline{p} + (2/3)\bar{p}$). This yields the pink allocation curve $a_2(v)$ shown in the upper right. The fact that these curves satisfy the inter-day IC constraint follows from the fact that the area of the two grey rectangles shown in the bottom right are equal.

where the randomized case is implemented via the lottery as in the example of Figure 1.

The key idea: $R_{12}(v)$ describes the best revenue we can get if we set a price of v on day 1. Since r_2 is the optimal day two price, if we are going to set a price above r_2 for day one, we may as well be optimal for day two. On the other hand, if the day one price is going to be below r_2 , we have to be careful about the inter-day IC constraint (5), and ironing the day two revenue curve may be necessary. This is precisely what the definition of $R_{12}(\cdot)$ in (8) does for us. The asymmetry between day one and day two, specifically the fact that the day one curve is never ironed, whereas the day two curve is, is a consequence of the inter-day IC constraint (5). We generalize this idea in the next section to solve the n -day problem.

4 An optimal allocation rule

4.1 Preliminaries

Our goal is to choose monotone allocation rules $a_i(v)$, for days $1 \leq i \leq n$, so as to maximize $\sum_{i=1}^n q_i \int_0^H \varphi_i(v) f_i(v) a_i(v) dv$.

For a distribution $f_i(\cdot)$ on $[0, H]$ with virtual value function $\varphi_i(\cdot) = v - \frac{1-F_i(v)}{f_i(v)}$, define $\gamma_i(v) := q_i \varphi_i(v) f_i(v)$. Then we aim to choose $a_i(v)$ to maximize $\sum_{i=1}^n \int_0^H \gamma_i(v) a_i(v) dv$.

Let $\Gamma_i(v) = \int_0^v \gamma_i(x) dx$. Observe that this function is the negative of the revenue curve, that is, $\Gamma_i(v) = -q_i R_i(v) = -q_i v(1 - F_i(v))$.⁴ Thus, $\Gamma_i(0) = \Gamma_i(H) = 0$ and $\Gamma_i(v) \leq 0$ for $v \in [0, H]$.

Definition 1. For any function Γ , define $\hat{\Gamma}(\cdot)$ to be the lower convex envelope⁵ of $\Gamma(\cdot)$. We say that $\hat{\Gamma}(\cdot)$ is ironed at v if $\hat{\Gamma}(v) \neq \Gamma(v)$.

Since $\hat{\Gamma}(\cdot)$ is convex, it is continuously differentiable except at countably many points and its derivative is monotone (weakly) increasing.

Definition 2. Let $\hat{\gamma}(\cdot)$ be the derivative of $\hat{\Gamma}(\cdot)$ and let $\gamma(\cdot)$ be the derivative of $\Gamma(\cdot)$.

Claim. The following facts are immediate from the definition of lower convex envelope (See Figure 2.):

- $\hat{\Gamma}(v) \leq \Gamma(v) \quad \forall v$.
- $\hat{\Gamma}(v_{min}) = \Gamma(v_{min})$ where $v_{min} = \operatorname{argmin}_v \Gamma(v)$. (This implies that there is no ironed interval that crosses over v_{min} .)
- $\hat{\gamma}(v)$ is an increasing function of v and hence its derivative $\hat{\gamma}'(v) \geq 0$ for all v .
- If $\hat{\Gamma}(v)$ is ironed in the interval $[\ell, h]$, then $\hat{\gamma}(v)$ is linear and $\hat{\gamma}'(v) = 0$ in (ℓ, h) .

We next define the sequence of functions that we will need for the construction:

⁴ $\Gamma_i(v) = q_i \int_0^v [x f_i(x) - (1 - F_i(x))] dx$. Integrating the first term by parts gives $\int_0^v x f_i(x) dx = v F_i(v) - \int_0^v F_i(x) dx$. Combining this with the second term yields $\Gamma_i(v) = -q_i v(1 - F_i(v))$.

⁵ The lower convex envelope of function $f(x)$ is the supremum over convex functions $g(\cdot)$ such that $g(x) \leq f(x)$ for all x . Notice that the lower convex envelope of $\Gamma(\cdot)$ is the negative of the ironed revenue curve $\bar{R}(v)$.




Figure 2: The black curve is $\Gamma_i(v)$, and its lower convex envelope $\hat{\Gamma}_i(v)$ is traced out by the thick light blue line. The curve is ironed in the interval $[\ell, h]$ (among others), so in that interval, $\hat{\Gamma}_i(v)$ is linear, and thus has second derivative equal to 0.

Definition 3. Let

$$\Gamma_{\geq n}(v) := \Gamma_n(v) \quad \text{and} \quad r_{\geq n} := \operatorname{argmin}_v \Gamma_{\geq n}(v).$$

Inductively, define, for $i := n - 1$ down to 1,

$$\Gamma_{\geq i}(v) := \begin{cases} \Gamma_i(v) + \hat{\Gamma}_{\geq i+1}(v) & v < r_{\geq i+1} \\ \Gamma_i(v) + \hat{\Gamma}_{\geq i+1}(r_{\geq i+1}) & v \geq r_{\geq i+1} \end{cases} \quad \text{and} \quad r_{\geq i} := \operatorname{argmin}_v \Gamma_{\geq i}(v).$$

The derivative of $\Gamma_{\geq i}(\cdot)$ is then

$$\gamma_{\geq i}(v) := \begin{cases} \gamma_i(v) + \hat{\gamma}_{\geq i+1}(v) & v < r_{\geq i+1} \\ \gamma_i(v) & v \geq r_{\geq i+1} \end{cases}.$$

Rewriting this yields

$$\gamma_{\geq i}(v) - \gamma_i(v) = \begin{cases} \hat{\gamma}_{\geq i+1}(v) & v < r_{\geq i+1} \\ 0 & v \geq r_{\geq i+1} \end{cases}. \quad (9)$$

4.2 The allocation rule

We define the allocation curves $a_i(\cdot)$ inductively. We will show later that they are optimal. Each allocation curve is piecewise constant. For day one, set

$$a_1(v) = \begin{cases} 0 & \text{if } v < r_{\geq 1}, \\ 1 & \text{otherwise} \end{cases}.$$

Suppose that a_{i-1} has been defined, for some $i < n$, with jumps at v_1, \dots, v_k , and values $0 = \beta_0 < \beta_1 \leq \beta_2 \leq \dots \leq \beta_k = 1$. That is,

$$a_{i-1}(v) = \begin{cases} 0 & \text{if } v < v_1, \\ \beta_j & v_j \leq v < v_{j+1} \quad 1 \leq j < k, \\ 1 & v_k \leq v \end{cases}.$$

Thus, we can write

$$a_{i-1}(v) = \sum_{j=1}^k (\beta_j - \beta_{j-1}) a_{i-1,j}(v)$$

where

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

Next we define $a_i(v)$.




Figure 3: This figure shows an example allocation curve $a_{i-1}(v)$ in purple, and illustrates some aspects of Definition 4. The curves $\Gamma_{\geq i}(v)$ and $\hat{\Gamma}_{\geq i}(v)$ are shown directly below the top figure. In this case, $r_{\geq i} \in [v_{j+1}, v_{j+2})$, so $j^* = j + 1$. The bottom figure shows how $a_{i,j}(v)$ is constructed from $a_{i-1,j}(v)$.

Definition 4. Let j^* be the largest j such that $v_j \leq r_{\geq i}$. For any $j \leq j^*$, consider two cases:

- $\hat{\Gamma}_{\geq i}(v_j) = \Gamma_{\geq i}(v_j)$, i.e. $\hat{\Gamma}_{\geq i}$ not ironed at v_j : In this case, define

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

- $\hat{\Gamma}_{\geq i}(v_j) \neq \Gamma_{\geq i}(v_j)$: In this case, let

- $\underline{v}_j :=$ the largest $v < v_j$ such that $\hat{\Gamma}_{\geq i}(v) = \Gamma_{\geq i}(v)$ i.e., not ironed, and
- $\bar{v}_j :=$ the smallest $v > v_j$ such that $\hat{\Gamma}_{\geq i}(v) = \Gamma_{\geq i}(v)$ i.e., not ironed.

Let $0 < \delta < 1$ such that

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

Then $\hat{\Gamma}_{\geq i}(\cdot)$ is linear between \underline{v}_j and \bar{v}_j :

$$\hat{\Gamma}_{\geq i}(v_j) = \delta \Gamma_{\geq i}(\underline{v}_j) + (1 - \delta) \Gamma_{\geq i}(\bar{v}_j).$$

Define

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

Finally, set $a_i(v)$ as follows:

$$a_i(v) = \begin{cases} \sum_{j=1}^{j^*} (\beta_j - \beta_{j-1}) a_{i,j}(v) & \text{if } v < r_{\geq i}, \\ 1 & v \geq r_{\geq i} \end{cases}. \quad (10)$$

Remark: In order to continue the induction and define $a_{i+1}(v)$ we need to rewrite $a_i(v)$ in terms of functions $a_{i,j}(v)$ that take only 0/1 values. This is straightforward.

Lemma 1. *The allocation curves $a_i(\cdot)$, for $1 \leq i \leq n$, are monotone increasing from 0 to 1 and satisfy the inter-day IC constraints (5). Moreover, each $a_i(\cdot)$ changes value only at points where $\hat{\Gamma}_{\geq i}(\cdot)$ is not ironed.*

Proof. That the allocation curves $a_i(\cdot)$ are weakly increasing, start out at 0, and end at 1 is immediate from the fact that they are convex combinations of the monotone allocation curves $a_{i,j}(\cdot)$. Also, by construction, each $a_i(\cdot)$ changes value only at points where $\hat{\Gamma}_{\geq i}(v)$ is not ironed.

So we have only left to verify that

$$\int_0^v a_{i-1}(x) dx \leq \int_0^v a_i(x) dx.$$

From the discussion above, for $v \leq r_{\geq i}$, we have

$$a_{i-1}(v) = \sum_{j=1}^{j^*} (\beta_j - \beta_{j-1}) a_{i-1,j}(v) \quad \text{and} \quad a_i(v) = \sum_{j=1}^{j^*} (\beta_j - \beta_{j-1}) a_{i,j}(v)$$

since $a_{i-1,j}(v) = 0$ for $v \leq r_{\geq i}$ and $j > j^*$. Thus, it suffices to show that for each $j \leq j^*$ and $v \leq r_{\geq i}$

$$\int_0^v a_{i-1,j}(x)dx \leq \int_0^v a_{i,j}(x)dx.$$

If $\Gamma_{\geq i}$ is not ironed at v_j , then this is an equality. Otherwise, for $v \leq v_j$, the left hand side is 0 and the right hand side is nonnegative. For $v_j \leq v \leq \bar{v}_j$ the left hand side is $(v - v_j)$, whereas the right hand side is $\delta(v - \underline{v}_j)$. Rearranging the inequality $v_j = \delta \underline{v}_j + (1 - \delta)\bar{v}_j \geq \delta \underline{v}_j + (1 - \delta)v$ implies that $v - v_j \leq \delta(v - \underline{v}_j)$. This completes the proof that (5) holds.

Notice that $\int_0^v a_{i-1,j}(x)dx = \int_0^v a_{i,j}(x)dx$ for $v < \underline{v}_j$ and $v > \bar{v}_j$, so $a_{i-1}(v) = a_i(v)$ unless $\Gamma_{\geq i}$ is ironed at v , or $v \geq r_{\geq i}$. We will use this fact in the proof of Claim 5.3 below. \square

5 Proof of optimality

In this section, we prove that the allocation rules and pricing of the previous section are optimal. To this end, we formulate our problem as an (infinite) linear program. We discussed the objective and constraints of the primal program in Section 2, and we have already shown above that our allocation rules are feasible for the primal program. We then construct a dual program, and a feasible dual solution for which we can prove strong duality and hence optimality of our solution.

5.1 The linear programming formulation

Recall the definitions from Section 2: The function $\gamma_i(v)$ is the derivative of $\Gamma_i(v) = \int_0^v q_i \varphi_i(x) f_i(x) dx$, where $\varphi_i(v) = v - \frac{1 - F_i(v)}{f_i(v)}$ is the day i virtual value function and q_i is the fraction of bidders with deadline i . Similarly $\hat{\gamma}_i(v)$ is the derivative of $\hat{\Gamma}_i(v)$. We use $[n]$ to denote the set of integers $\{1, \dots, n\}$.

The Primal

Variables: $a_i(v)$, for all $i \in [n]$, and all $v \in [0, H]$.

$$\text{Maximize } \sum_{i=1}^n \int_0^H a_i(v) \gamma_i(v) dv$$

Subject to

$$\begin{aligned} \int_0^v a_i(x) dx - \int_0^v a_{i+1}(x) dx &\leq 0 & \forall i \in [n-1] \quad \forall v \in [0, H] & \quad (\text{dual variables } \alpha_i(v)) \\ a_i(v) &\leq 1 & \forall i \in [n] \quad \forall v \in [0, H] & \quad (\text{dual variables } b_i(v)) \\ -a'_i(v) &\leq 0 & \forall i \in [n] \quad \forall v \in [0, H] & \quad (\text{dual variables } \lambda_i(v)) \\ a_i(v) &\geq 0 & \forall i \in [n] \quad \forall v \in [0, H]. & \end{aligned}$$

The Dual

Variables: $b_i(v), \lambda_i(v)$, for all $i \in [n]$, and all $v \in [0, H]$, $\alpha_i(x)$ for $i \in [n-1]$ and all $x \in [0, H]$.

$$\text{Minimize } \int_0^H [b_1(v) + \dots + b_n(v)] dv$$

Subject to

$$\begin{aligned} b_1(v) + \lambda'_1(v) + \int_v^H \alpha_1(x) dx &\geq \gamma_1(v) && \forall v \in [0, H] \text{ (primal var } a_1(v)) \\ b_i(v) + \lambda'_i(v) + \int_v^H \alpha_i(x) dx - \int_v^H \alpha_{i-1}(x) dx &\geq \gamma_i(v) && \forall v \in [0, H], i = 2, \dots, n-1 \text{ (primal var } a_i(v)) \\ b_n(v) + \lambda'_n(v) - \int_v^H \alpha_{n-1}(x) dx &\geq \gamma_n(v) && \forall v \in [0, H] \text{ (primal var } a_n(v)) \\ \lambda_i(H) &= 0 && \forall i \in [n] \\ \alpha_i(v) &\geq 0 && \forall v \in [0, H], i \in [n-1] \\ b_i(v), \lambda_i(v) &\geq 0 && \forall i \in [n] \forall v \in [0, H]. \end{aligned}$$

5.2 Conditions for strong duality:

As long as there are feasible primal and dual solutions satisfying the following conditions, strong duality holds. See Appendix A for a proof that these conditions are sufficient.

$$a_i(v) > 0 \Rightarrow \lambda_i(v) \text{ continuous at } v \quad i \in [n] \quad (11)$$

$$a_i(v) < 1 \Rightarrow b_i(v) = 0 \quad i \in [n] \quad (12)$$

$$a'_i(v) > 0 \Rightarrow \lambda_i(v) = 0 \quad i \in [n] \quad (13)$$

$$\int_0^v a_i(x) dx < \int_0^v a_{i+1}(x) dx \Rightarrow \alpha_i(v) = 0 \quad i \in [n-1] \quad (14)$$

$$b_i(v) + \lambda'_i(v) + \int_v^H \alpha_i(x) dx - \int_v^H \alpha_{i-1}(x) dx > \gamma_i(v) \Rightarrow a_i(v) = 0 \quad i = 2, \dots, n-1 \quad (15)$$

$$b_1(v) + \lambda'_1(v) + \int_v^H \alpha_1(x) dx > \gamma_1(v) \Rightarrow a_1(v) = 0 \quad (16)$$

$$b_n(v) + \lambda'_n(v) - \int_v^H \alpha_{n-1}(x) dx > \gamma_n(v) \Rightarrow a_n(v) = 0 \quad (17)$$

We allow $a'_i(v) \in \mathbb{R} \cup \{+\infty\}$, otherwise we could not even encode a single-price auction.⁶

5.3 The proof

Theorem 2. *The allocation curves presented in Subsection 4.2 are optimal, that is, obtain the maximum possible expected revenue.*

⁶ In particular, $a_i(v)$ may have (countably many) discontinuities, in which points $a'_i(v) = +\infty > 0$. However, in our proof of optimality $a'_i(v)$ appears only as a factor of the product $a'_i(v)\lambda_i(v)$. Every time $a'_i(v) = +\infty$, the corresponding dual value of $\lambda_i(v)$ is 0—by condition (13). See also Appendix A.

Proof. To prove the theorem, we verify that there is a setting of feasible dual variables for which all the conditions for strong duality hold. To this end, set the variables as follows:

$$\lambda_i(v) = \Gamma_{\geq i}(v) - \hat{\Gamma}_{\geq i}(v) \quad (18)$$

$$b_i(v) = \begin{cases} 0 & v < r_{\geq i} \\ \hat{\gamma}_{\geq i}(v) & v \geq r_{\geq i} \end{cases} \quad (19)$$

$$\alpha_i(v) = \begin{cases} \hat{\gamma}'_{\geq i+1}(v) & v < r_{\geq i+1} \\ 0 & v \geq r_{\geq i+1} \end{cases} \quad (20)$$

From Claim 4.1, it follows that $\lambda_i(v), \alpha_i(v) \geq 0$ for all v and i . Since $r_{\geq i}$ is the minimum of $\hat{\Gamma}_{\geq i}(\cdot)$, we have $\hat{\gamma}_{\geq i}(r_{\geq i}) = 0$. Moreover, since $\hat{\gamma}_{\geq i}(\cdot)$ is increasing, $b_i(v) \geq 0$ for all v and i .

Taking the derivative of (18), and using Equation (9), we obtain:

$$\gamma_i(v) - \lambda'_i(v) = \begin{cases} \hat{\gamma}_{\geq i}(v) - \hat{\gamma}_{\geq i+1}(v) & v < r_{\geq i+1} \\ \hat{\gamma}_{\geq i}(v) - 0 & v \geq r_{\geq i+1} \end{cases} \quad (21)$$

$$\gamma_n(v) - \lambda'_n(v) = \hat{\gamma}_n(v) \quad (22)$$

Also, using (20) and the fact that $\hat{\gamma}_{\geq i+1}(r_{\geq i+1}) = 0$, we get:

$$A_i(v) := \int_v^H \alpha_i(x) dx = \begin{cases} -\hat{\gamma}_{\geq i+1}(v) & v < r_{\geq i+1} \\ 0 & v \geq r_{\geq i+1} \end{cases} \quad (23)$$

Condition (11) from Section 5.2 holds since $\Gamma_{\geq i}(v)$ and $\hat{\Gamma}_{\geq i}(v)$ are both continuous functions. The proofs of all remaining conditions for strong duality from Section 5.2 can be found below. \square

Claim. *Condition (12): For all i and v , $a_i(v) < 1 \implies b_i(v) = 0$.*

Proof. If $a_i(v) < 1$, then $v < r_{\geq i}$, so by construction, $b_i(v) = 0$. \square

Claim. *Condition (13): For all i and v , $a'_i(v) > 0 \implies \lambda_i(v) = 0$.*

Proof. From Subsection 4.2, $a'_i(v) > 0$ only for unironed values of v , at which $\lambda_i(v) = 0$. \square

Claim. *Condition (14): For all i and v , $\int_0^v a_i(x) dx < \int_0^v a_{i+1}(x) dx \implies \alpha_i(v) = 0$.*

Proof. As discussed at the end of the proof of Lemma 1, $\int_0^v a_i(x) dx = \int_0^v a_{i+1}(x) dx$ unless $\Gamma_{\geq i+1}$ is ironed at v , or $v \geq r_{\geq i}$. In both of these cases $\alpha_i(v) = 0$ (by part 4 of Claim 4.1 and Definition 20, respectively). \square

Claim. *Conditions (15)- (17) and dual feasibility: For all i and v , $a_i(v) > 0 \implies$ the corresponding dual constraint is tight, and the dual constraints are always feasible.*

Proof. Rearrange the dual constraint $b_i(v) + A_i(v) - A_{i-1}(v) + \lambda'_i(v) \geq \gamma_i(v)$ to

$$b_i(v) - A_{i-1}(v) \geq \gamma_i(v) - \lambda'_i(v) - A_i(v).$$

Fact 1: For $i \in [n - 1]$, $\gamma_i(v) - \lambda'_i(v) - A_i(v) = \hat{\gamma}_{\geq i}(v)$ for all v . To see this use (21) and (23):

$$\gamma_i(v) - \lambda'_i(v) = \begin{cases} \hat{\gamma}_{\geq i}(v) - \hat{\gamma}_{\geq i+1}(v) & v < r_{\geq i+1} \\ \hat{\gamma}_{\geq i}(v) - 0 & v \geq r_{\geq i+1} \end{cases} \quad A_i(v) = \begin{cases} -\hat{\gamma}_{\geq i+1}(v) & v < r_{\geq i+1} \\ 0 & v \geq r_{\geq i+1} \end{cases}$$

Fact 2: For $i \in \{2, \dots, n\}$, $b_i(v) - A_{i-1}(v) = \hat{\gamma}_{\geq i}(v)$ for all v .

$$b_i(v) = \begin{cases} 0 & v < r_{\geq i} \\ \hat{\gamma}_{\geq i}(v) & v \geq r_{\geq i} \end{cases} \quad -A_{i-1}(v) = \begin{cases} \hat{\gamma}_{\geq i}(v) & v < r_{\geq i} \\ 0 & v \geq r_{\geq i} \end{cases}$$

Hence for $i \in \{2, \dots, n - 1\}$, $b_i(v) - A_{i-1}(v) = \gamma_i(v) - \lambda'_i(v) - A_i(v)$ for all v .

For $i = n$, since $\gamma_{\geq n} = \gamma_n$, and $\gamma_n(v) - \lambda'_n(v) = \hat{\gamma}_n(v)$. Combining this with Fact 2 above, we get that $b_n(v) - A_{n-1}(v) + \lambda'_n(v) = \gamma_n(v)$ for all v .

Finally, for $i = 1$, using Fact 1, for $v < r_{\geq 1}$, we get

$$b_1(v) = 0 \geq \hat{\gamma}_{\geq 1}(v) = \gamma_1(v) - \lambda'_1(v) - A_1(v)$$

which is true for $v < r_{\geq 1}$. For $v \geq r_{\geq 1}$, we get

$$b_1(v) = \gamma_{\geq 1}(v) = \gamma_1(v) - \lambda'_1(v) - A_1(v),$$

so the dual constraint is tight when $a_1(v) > 0$ as this starts at $r_{\geq 1}$. \square

The above claims prove that this dual solution satisfies feasibility and all complementary slackness and continuity conditions from Section 5.2 hold.

References

- [1] S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 512–521, Oct 2011.
- [2] Edward J Anderson and Peter Nash. *Linear programming in infinite-dimensional spaces: theory and applications*. John Wiley & Sons, 1987.
- [3] Mark Armstrong. Multiproduct Nonlinear Pricing. *Econometrica*, 64(1):51–75, 1996.
- [4] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of the 55th Annual Symposium on Foundations of Computer Science, FOCS '14*, 2014.
- [5] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The Limits of Price Discrimination . *American Economic Review*, 105(3):921–957, mar 2015.
- [6] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S Matthew Weinberg. Pricing randomized allocations. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 585–597. Society for Industrial and Applied Mathematics, 2010.




Figure 4: This figure illustrates what some of the dual variables might be for the case of two days when $r_{\geq 1} < r_2$. The upper figure plots the functions $\hat{\gamma}_{12}(v)$ and $\hat{\gamma}_2(v)$, and the lower figure shows $b_1(v)$ in dark grey, $b_2(v)$ in pink and $A_1(v) = \int_v^H \alpha_1(x) dx$ in green. Note that up to r_2 , the function $A_1(v) = -\hat{\gamma}_2(v)$.

- [7] Yang Cai and C. Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 522–531, Oct 2011.
 - [8] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. On optimal multidimensional mechanism design. *SIGecom Exch.*, 10(2):29–33, June 2011.
 - [9] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 130–139. IEEE Computer Society, 2012.
 - [10] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Reducing revenue to welfare maximization: Approximation algorithms and other generalizations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '13*, pages 578–595. SIAM, 2013.
 - [11] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Understanding incentives: Mechanism design becomes algorithm design. *CoRR*, abs/1305.4002, 2013.
 - [12] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the forty-eighth annual ACM symposium on Theory of computing*. ACM, 2016.
 - [13] Yang Cai and Zhiyi Huang. Simple and nearly optimal multi-item auctions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '13*, pages 564–577, 2013.
 - [14] S Chawla, JD Hartline, DL Malec, and B Sivan. Multi-parameter mechanism design and sequential posted pricing. *Proceedings of the forty- . . .*, 2010.
 - [15] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07*, pages 243–251, 2007.
 - [16] Shuchi Chawla, David Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297 – 317, 2015.
 - [17] Shuchi Chawla and Balasubramanian Sivan. Bayesian algorithmic mechanism design. *SIGecom Exch.*, 13(1):5–49, November 2014.
 - [18] Constantinos Daskalakis. Multi-item auctions defying intuition? *ACM SIGecom Exchanges*, 14(1):41–75, 2015.
 - [19] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the 14th ACM Conference on Electronic Commerce, EC '13*, pages 269–286, 2013.
 - [20] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 449–450. ACM, 2015.
-

- [21] Yiannis Giannakopoulos and Elias Koutsoupias. Duality and optimality of auctions for uniform distributions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, EC '14, pages 259–276, 2014.
 - [22] Yiannis Giannakopoulos and Elias Koutsoupias. Selling two goods optimally. In *ICALP (2)*, volume 9135 of *Lecture Notes in Computer Science*, pages 650–662. Springer, 2015.
 - [23] Nima Haghpanah and Jason Hartline. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 757–758. ACM, 2015. Updated version: <http://arxiv.org/abs/1404.1341>.
 - [24] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 656–656, 2012.
 - [25] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 565–566, 2013.
 - [26] Sergiu Hart and Philip J. Reny. Maximal revenue with multiple goods: Nonmonotonicity and other observations. Technical report, The Center for the Study of Rationality, Hebrew University, Jerusalem, 2012.
 - [27] Jason D. Hartline. Mechanism design and approximation. 2014.
 - [28] Jean-Jacques Laffont, Eric Maskin, and Jean-Charles Rochet. Optimal nonlinear pricing with two-dimensional characteristics. *Information, Incentives and Economic Mechanisms*, pages 256–266, 1987.
 - [29] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.
 - [30] David G Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, 1997.
 - [31] Alejandro M. Manelli and Daniel R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1 – 35, 2006.
 - [32] R. Preston McAfee and John McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335 – 354, 1988.
 - [33] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
 - [34] Gregory Pavlov. Optimal mechanism for selling two goods. *The BE Journal of Theoretical Economics*, 11(1), 2011.
 - [35] Marek Pycia. *Stochastic vs deterministic mechanisms in multidimensional screening*. PhD thesis, MIT, 2006. Chapter 3.
 - [36] Tim Roughgarden. Approximately optimal mechanism design: Motivation, examples, and lessons learned. *SIGecom Exch.*, 13(2):4–20, January 2015.
-

- [37] John Thanassoulis. Hagglng over substitutes. *Journal of Economic Theory*, 117(2):217 – 245, 2004.
- [38] Andrew Chi-Chih Yao. An n -to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109, 2015.

A Proof of strong duality

Theorem 3. *Let $a_i(\cdot)$, $b_i(\cdot)$, $\lambda_i(\cdot)$, $\alpha_i(\cdot)$ be functions feasible for the primal and dual, satisfying all the conditions from Section sec:CS. Then they are optimal.*

Proof. First, we prove weak duality. For any feasible primal and dual:

$$\int_0^H \sum_{i=1}^n b_i(v) dv \tag{24}$$

$$= \int_0^H \sum_{i=1}^n (1 \cdot b_i(v) + 0 \cdot [\lambda_i(v) + \alpha_i(v)]) dv. \tag{25}$$

Applying primal feasibility, we see that this quantity is

$$\geq \int_0^H \sum_{i=1}^n \left(a_i(v)b_i(v) - a'_i(v)\lambda_i(v) + \left[\int_0^v a_i(x) - a_{i+1}(x)dx \right] \alpha_i(v) \right) dv. \tag{26}$$

We rewrite this expression using the following.

- Applying integration by parts, using the facts that $\lambda_i(\cdot)$ is continuous (Condition (11)) and equal to 0 at any point that $a'_i(v) = \infty$,⁷ we get

$$- \int_0^H a'_i(v)\lambda_i(v) dv = -a_i(v)\lambda_i(v) \Big|_0^H + \int_0^H a_i(v)\lambda'_i(v) dv = \int_0^H a_i(v)\lambda'_i(v) dv,$$

since $a_i(0) = 0$ and $\lambda_i(H) = 0$.

- Second, interchanging the order of integration, we get

$$\int_0^H \int_0^v [a_i(x) - a_{i+1}(x)dx] \alpha_i(v) dv = \int_0^H \left(a_i(v) \int_v^H \alpha_i(x) dx - a_{i+1}(v) \int_v^H \alpha_i(x) dx \right) dv.$$

Combining these shows that (26) equals

$$\begin{aligned} & \int_0^H \left(\sum_{i=1}^n a_i(v) \left[b_i(v) + \lambda'_i(v) + \int_v^H \alpha_i(x) - \int_v^H \alpha_{i-1}(x) dx \right] \right) dv \\ & \geq \int_0^H \sum_{i=1}^H a_i(v)\gamma_i(v) dv \end{aligned} \tag{27}$$

⁷ $a'_i(v)$ can be ∞ at only countably many points.

where the last inequality is dual feasibility. (Note that $\alpha_0(\cdot) = \alpha_n(\cdot) = 0$.)

Comparing (24) and (27) yields weak duality, i.e., $\sum_i \int_0^H b_i(v) dv \geq \sum_i \int_0^H a_i(v) \gamma_i(v) dv$.

If the conditions (11)-(17) hold, we also have strong duality and hence optimality: To show that (25) = (26), observe that

- (12) $a_i(v) < 1$ implies that $b_i(v) = 0$;
- (13) $a'_i(v) > 0$ implies that $\lambda_i(v) = 0$.
- (14) $\int_0^v (a_{i+1}(x) - a_i(x)) dx > 0$ implies that $\alpha_i(v) = 0$ for $i = 1, \dots, n-1$.

Finally, (27) is an equality rather than an inequality because of conditions (15)-(17). □