The optimal mechanism for selling to a budget constrained buyer: the general case

NIKHIL R. DEVANUR, Microsoft Research
S. MATTHEW WEINBERG, Princeton University

We consider a revenue-maximizing seller with a single item facing a single buyer with a private budget. The (value, budget) pair is drawn from an *arbitrary* and possibly correlated distribution. We characterize the optimal mechanism in such cases, and quantify the amount of price discrimination that might be present. For example, there could be up to $3 \cdot 2^{k-1} - 1$ distinct non-trivial menu options in the optimal mechanism for such a buyer with k distinct possible budgets (compared to k if the marginal distribution of values conditioned on each budget has decreasing marginal revenue [Che and Gale, 2000], or 2 if there is an arbitrary distribution and one possible budget [Chawla et al., 2011]).

Our approach makes use of the duality framework of Cai et al. [2016], and duality techniques related to the "FedEx Problem" of Fiat et al. [2016]. In contrast to [Fiat et al., 2016] and other prior work, we characterize the optimal primal/dual without nailing down an explicit closed form.

1 INTRODUCTION

The theory of optimal auction often equates willingness to pay with the ability to pay. While this leads to clean formulations and elegant solutions, there are many instances where this assumption is violated. For example, the valuation for an item could be based on future earnings but credit could be difficult to obtain due to imperfect capital markets (see Che and Gale [2000] for more discussion on this). This is especially true when the sums involved are huge, such as in spectrum auctions [Cramton, 1995]. Budget constraints have been a predominant feature in ad auctions, which has led to a revived interest in understanding how they impact auction design [Abrams, 2006, Borgs et al., 2005, Dobzinski et al., 2008, Goel et al., 2012]. Sometimes, budget constraints are imposed exogenously, e.g., in contests with some form of bidding [Gavious et al., 2002], and in auctions for players to form a team in a professional sports league [Venkateswaran, 2013]. Another reason for the existence of budget constraints is that a principal may employ an agency to bid on her behalf, and imposes a budget constraint to control the spending [Che and Gale, 1998].

In this paper we consider selling a single item to a single bidder, who has a quasi-linear utility function as well as a hard upper bound on her payment. In other words, if the buyer obtains the item and pays p, then her utility is $v - p$, as long as $p < b$, for some two numbers v and b. We call v her valuation and b her budget, and both of them are private information. We consider the problem in the Bayesian setting, where there is a joint probability distribution over types (v, b), which is known to the seller.

In such settings, it is known that the optimal mechanism necessitates selling *lotteries* (e.g. awarding the item with probability in $(0, 1)$). While there may be practical limitations to selling lotteries, there are many cases where lotteries are sold in disguise: an upgrade to a business class seat on an airplane is essentially a lottery since the upgrade is available only with a certain probability. A similar situation has been observed in a spot market for virtual machines, where they are evicted with some probability even when there is no excess demand [Kilcioglu and Maglaras, 2015a,b]. The same effect can be acheived via a degradation of service, depending on the context. E.g., if the item is time sensitive such as the latest fashion trend, then delaying the sale is an effective degradation of service. Equivalently, the item could be divisible, and x would be the fraction of the item sold.

Treating the item as divisible is a resonable approximation when there are a large number of copies of identical items, such as in the case of ad auctions.

Two variants of this problem have been considered [Che and Gale, 2000]: the optimal *unconditional* mechanism where a buyer can exaggerate her budget as long as the eventual payment is less than her actual budget, and the optimal *conditional* mechanism where the seller can prevent such exaggerations. This can be enforced by requiring a cash bond, for example, or by requiring the buyer to pay his full budget with some small probability. The optimal conditional mechanism can do a higher degree of *price discrimination* and as a result extract more revenue than the unconditional mechanism.[1] In this paper we focus on the optimal conditional mechanism (and henceforth, simply the optimal mechanism) .

The core of the matter is how to use the budget constraints to do price discrimination. Buyers can feasibly report a lower budget, so the ones with higher budgets extract higher information rents. The economic intuition gained from previous work on this subject is as follows. The following assumptions on the value distributions are commonly made: a distribution with a CDF F and density f is called *regular* if $v - \frac{1-F(v)}{f(v)}$ is non decreasing, and it satisfies *declining marginal revenues* (DMR) if the function $v(1 - F(v))$ is concave.

Fixed budget, regular or DMR distributions [Laffont and Robert, 1996]: If the buyer's value is drawn from a regular or DMR distribution and the budget is fixed and known to the seller, the optimal mechanism posts a price equal to the minimum of the budget and the Myerson reserve (at most one non-trivial menu option). There is no price discrimination in this case. This is not true when either the regularity or the DMR assumption is violated.

Fixed budget, arbitrary distributions [Chawla et al., 2011]: If the buyer's value is drawn from an arbitrary distribution and the budget is fixed and known to the seller, two non-trivial menu options are necessary and sufficient: receive the item with probability 1 and pay the budget, or receive the item with probability $q < 1$ and in this case pay a price $< qb$ (Example 3.8 shows that 2 options are necessary, as does Example 1 in [Chawla et al., 2011]). This shows that there could be a non-trivial price discrimation even when there is only one possible budget. The combination of arbitrary distributions as well as different budgets amplifies this aspect, and is one of the difficulties in handling such cases.

Private budget, DMR distributions [Che and Gale, 2000]: If there are k possible budgets $b_1 < \ldots < b_k$, private to the buyer, and the buyer's value conditioned on any given budget is drawn from a DMR distribution, then the optimal mechanism has the following format: the menu offered to buyers with budget b_1 has at most one non-trivial option. There is a cuttoff budget b_{i^*}, below which the budget binds and above which the menu contains an option to receive the item with probability 1. The menu presented to a buyer with budget b_j contains every option offered to the buyer with budget b_{j-1}, plus at most one additional option (meant for the highest valued types). For $j < i^*$, the additional option is to receive the item with a probability $q_j < 1$ and pay b_j. For $j \geq i^*$, the additional option offers the item with probability 1, at a possibly lower price than the same option for the lower budgeted buyers. (Higher budgeted types get larger discounts.) There are therefore at most k non-trivial menu options.

We show how to derive all of the above results using a common approach. The first two results appear in Section 3 as a warm-up, and the final one appears in Appendix D. In comparison to existing

[1] Che and Gale [2000] show that under cerain assumptions on the prior distribution, the optimal unconditional mechanism is essentially a price curve, which is a mapping from the allocation probability $x \in [0, 1]$ to a price $p(x)$. The optimal price curve is further guaranteed to be convex; the marginal cost per unit amount is non decreasing. The search space is therefore the set of all non decreasing convex functions p such that $p(0) = 0$. For the optimal conditional mechanism, the search space is a 2 dimensional function, since there can be a different price curve for each budget.

proofs of the above results, we believe our proof for the private budget, DMR case is considerably more structured.[2] Existing proofs for the single budet case are already quite straight-forward, but our approach shows how to properly view the single budget, regular/DMR setting as a special case of the single budget, arbitrary distribution setting where additional structure exists (rather than as a necessary assumption to get the proof going). Our principled approach allows us to extend this characterization all the way to the most general case. **Our main contribution in this paper is to characterize the optimal mechanism in the private budget setting without any assumption on the distributions.** The main economic intuition gained from our results is a quantitative understanding of the degree of price discrimination that might arise in the general case. For example:[3]

> If the buyer's value is drawn from an arbitrary distribution and there are possible budgets $b_1 < \ldots < b_k$, private to the buyer, the optimal mechanism has the following format: the menu offered to buyers with budget b_1 has at most two non-trivial options. There is a cuttoff budget b_i, below which the budget binds and above which the menu contains an option to receive the item with probability 1. Every option on the menu presented to a buyer with budget b_{j-1} *splits* into at most two options to be included on the menu presented to a buyer with budget b_j. In addition, there is at most one additional option at the top (i.e. higher allocation probability than all other options). If $j < i$, the additional option offers the item at probability $q < 1$ at price $b_j q$. If $j \geq i$, the additional option offers the item at probability 1 (at most $3 \cdot 2^{k-1} - 1$ non-trivial menu options) [Theorem 4.17].

1.1 Techniques

The predominant technique in characterizing optimal mechanisms for multi-dimensional types is to write it as a mathematical program and use its optimality conditions [Che and Gale, 2000, Malakhov and Vohra, 2004, 2009]. In particular, a very successul technique has been to guess the optimal mechanism as well as the optimal dual variables/Lagrange multipliers and verify that they satisfy these conditions [Daskalakis et al., 2013, 2015a, Fiat et al., 2016, Giannakopoulos and Koutsoupias, 2014, 2015, Haghpanah and Hartline, 2015, Laffont and Robert, 1996, Manelli and Vincent, 2006, Pai and Vohra, 2014]. Within these, there are several differences such as whether the types are discrete or continuous, giving rise to finite or continuous linear programs (LPs), whether to use "allocation variables" or "utility variables", and so on. We stick to discrete types and finite LPs that use allocation variables, althoug our results extend to continuous types (see [Cai et al., 2016] for a proof/discussion of why). Finite LPs are technically easier since strong duality holds without any additional work, while continuous versions have the advantage that one can do calculus more easily.

A key aspect of our approach is the idea of Lagrangifying only the incentive/budget constraints, as in the duality framework of Cai et al. [2016]. A particularly nice property of this is that the duals give a virtual valuation for each type. The optimal mechanism is then such that all the types with positive virtual values are allocated with probability 1, and all the types with negative virtual values are not allocated at all. For the types with a zero virtual value the allocation probability is

[2]We stop short of claiming that the complete proof is simpler - there are quite a few technical hurdles to overcome. But we believe that it is much simpler to get an understanding of both the approach and the special cases where these technical hurdles don't arise.

[3]Note that in the description below, we are targeting the optimal *ex-post IR* mechanism (charges at most budget when the item is awarded). Our approach, literally word-for-word with some obvious replacements also characterize the optimal *interim IR* mechanism (charges at most the budget). For the optimal interim IR mechanism, replace $b_j q$ in the text below with b_j.

typically in (0, 1); and we need to use additional structure of the dual to pin down these probabilities (specifically, complementary slackness). The optimal dual is such that the sign of the virtual value is monotone in the buyer's value.

A special case of our problem is the FedEx problem of Fiat et al. [2016]: each buyer has a value v for a service and a "deadline" d, where d is a time. The buyer gets a value v for a service as long as it is delivered before the deadine d. A buyer can be prevented from reporting a later deadline by making sure that he is always serviced exactly at his deadline. Thus, the only relevant IC constraints are when a buyer can misreport his value and an earlier deadline. This is exactly our problem, in the instances where none of the budgets are binding. The conditional mechanism can also prevent buyers from over-reporting their budget, so the IC constraints are exactly the same syntactically; semantically, deadlines are replaced by budgets. As mentioned earlier, we show that there is a threshold budget above which the budget constraint does not bind. In this regime, the problem becomes identical to the FedEx problem and we use essentially the same primal/dual solution from Fiat et al. [2016], translated to our LP. This translation shows that their dual solution indeed induces monotonically signed virtual values.

Finally, we emphasize one key technical departure from previous works. In prior works, an explicit optimal dual solution and explicit optimal primal solution are proposed, and then optimality/complementary slackness is proved. In our setting, this would be a complete nightmare - mostly due to the multiple budget constraints. Instead, we characterize what the optimal dual solution must look like, and show that there exists an optimal primal solution of the desired form which satisfies complementary slackness. All this is done without excessive algebraic calculation to nail down a closed form. Still, it is easy to see from our approach how one could go about finding the optimal primal/dual - the point is just that an exact calculation isn't necessary to characterize the solutions. We defer a more technical overview of our approach to later sections.

1.2 Conclusion and Future directions

Organization: In Section 2, we formally state our problem, notation, etc. We also write an LP formulation and go through a series of reformulations a la Cai et al. [2016], resulting in a program to search for the optimal dual. In Section 3, we instantiate this approach for the case of a single public budget. The results in this section are already known due to Chawla et al. [2011], Laffont and Robert [1996], but it will be instructive to see our approach address this special case. In Section 4, we prove our main result: a characterization of the optimal mechanism with arbitrary value distributions and private budgets. In the appendix, we provide a longer overview of other related work (Appendix A), derive Che and Gale's tighter characterization when all value distributions are DMR as a corollary (Appendix D), and provide several omitted proofs (Appendices B and C).

Future Work: Along with the FedEx Problem, our work is an example of problems that are "one-and-a-half-dimensional:" the space of optimal mechanisms is considerably richer than in single-dimensional settings, yet there is still enough structure to obtain meaningful characterizations. Another example of this is when there are multiple copies of an item, and the buyer's valuation grows linearly with the number of items subject to a capacity constraint. The type of a buyer is once again a pair, a per unit valuation and a capacity. Devanur et al. [2017] characterize the optimal mechanism when the marginal distributions satisfy DMR; extending this to arbitrary distributions is an interesting open question. Another enticing direction is understanding the optimal mechanism for *multiple* buyers in these settings. Non-trivial innovation is likely needed beyond existing techniques - see technical sections for more detail.[4] Additionally, it would be

[4]Essentially the challenge is that for a single bidder, one only needs virtual values that are monotonically signed (i.e. never go from positive to negative). But for multiple bidders, one additionally needs virtual values that are monotone

interesting to push the limits of these settings: for example, what about a single buyer with multiple of these "half-dimensions" (e.g., a budget and a deadline)?

2 PRELIMINARIES

We consider the problem faced by a single seller with a single item for sale to a single buyer. The buyer has a value v for receiving the item and a hard budget constraint b on her maximum ability to pay. Our techniques apply both to ex-post IR (where the buyer must pay at most b conditioned upon winning the item), and interim IR constraints (where the buyer must pay at most b, but pays before learning whether or not she receives the item). **Our characterization applies in both settings**, but we will give proofs only for ex-post IR. The same proof applies nearly verbatim for interim IR.

Buyer Types. The buyer's type, (v, b) is drawn jointly from an arbitrarily correlated distribution supported on $\{0, 1, \ldots, \bar{v}\} \times \{b_1, \ldots, b_k\}$, with $\{b_1, \ldots, b_k\} \subseteq \{0, 1, \ldots, \bar{v}\}$. The assumption that the support is finite and integral is for technical simplicity and without loss of generality, see [Cai et al., 2016] for a rigorous argument. Define $f_i(w) = \Pr[(v, b) = (w, b_i)]$, and $F_i(w) = \sum_{x \le w} f_i(w)$.

Mechanisms. We'll use the variables $\pi_i(v)$ to denote the probability that a bidder with type (v, b_i) receives the item, and $p_i(v)$ to denote the expected price they pay. We seek mechanisms that are:

- **Incentive Compatible**: $v \cdot \pi_i(v) - p_i(v) \ge v \cdot \pi_j(w) - p_i(w)$, for all $j \le i$, and all v, w.
- **Individually Rational**: $v \cdot \pi_i(v) - p_i(v) \ge 0$, for all i, v.
- **Feasible**: $\pi_i(v) \in [0, 1]$ for all i, v.
- **Budget-Respecting**: $p_i(v) \le b_i \cdot \pi_i(v)$.[5]

Note that in defining Incentive Compatibility above, we are assuming that the buyer may only lie by *under*-reporting her budget. This is standard in the literature on auction design with budget constraints, and can be enforced e.g. by asking the buyer to front her budget before any sale takes place. It is folklore knowledge that many of the above constraints are redundant (see e.g. [Fiat et al., 2016, Laffont and Robert, 1996]), and the revenue-optimal mechanism is the solution to the following LP. We write the LP for ex-post budget-respecting, the necessary change for interim budget-respecting is obvious.[6]

LP1:
- Maximize: $\sum_i \sum_v f_i(v) p_i(v)$ (expected revenue).
- Subject to:
  - $v \cdot \pi_i(v) - p_i(v) \ge v \cdot \pi_i(v-1) - p_i(v-1)$, for all i and $v \ge 1$ (Leftwards IC).
  - $v \cdot \pi_i(v) - p_i(v) \ge v \cdot \pi_i(v+1) - p_i(v+1)$, for all i and $v \le \bar{v}$ (Rightwards IC).
  - $v \cdot \pi_i(v) - p_i(v) \ge v \cdot \pi_{i-1}(v) - p_{i-1}(v)$, for all $i > 1$ and v (Downwards IC).
  - $\pi_1(0) = p_1(0) = 0$ (Individual Rationality).
  - $\pi_i(v) \le 1$, for all i, v (Feasibility).
  - $p_i(\bar{v}) \le b_i \cdot \pi_i(\bar{v})$ for all i (Budget-Respecting).

Proposition 2.1 (Folklore,[Fiat et al., 2016],[Laffont and Robert, 1996]). *Every solution to LP1 is an optimal mechanism.*

The above essentially just observes that many of the IC/budget constraints are redundant (implied by the others). Following the duality framework of Cai et al. [2016], we will also take the partial

non-decreasing when positive. The present results certainly fall short of providing clear guidance on how to achieve monotone non-decreasing virtual values.

[5]For interim IR, replace this with $p_i(v) \le b_i$ for all i. Adapting the proof to this case involves essentially just propagating this change through computation of virtual values, etc.

[6]This is the last time we will reference necessary changes for interim IR.

Lagrangian of LP1, with dual multipliers $\lambda_i(v, v-1)$ for the Leftwards IC constraints, $\lambda_i(v, v+1)$ for the Rightwards IC constraints, $\alpha_i(v)$ for the Downwards IC constraints, and γ_i for the budget constraints. For the reader unfamiliar with their framework, we'll go through their approach, omitting proofs (which can be found in [Cai et al., 2016]). By strong duality, the solution to LP1 corresponds to the primal variables in the solution to the following max-min program. We will use $\Lambda = (\vec{\lambda}, \vec{\alpha}, \vec{\gamma})$ to refer to a complete set of Lagrangian multipliers for the IC/budget constraints.

Max-Min1:

- Maximize: $\mathcal{L}_{\min}(\vec{\pi}, \vec{p})$, where $\mathcal{L}_{\min}(\vec{\pi}, \vec{p})$ is the value of the following LP1$(\vec{\pi}, \vec{p})$:
  - Minimize $\sum_{i,v} f_i(v)p_i(v) + \sum_i \gamma_i(b_i\pi_i(\bar{v}) - p_i(\bar{v}))$
    $+ \sum_{i,v} \big[\alpha_i(v)(v\pi_i(v) - p_i(v) - v\pi_{i-1}(v) + p_{i-1}(v))$
    $+ \lambda_i(v, v-1)(v \cdot \pi_i(v) - p_i(v) - v \cdot \pi_i(v-1) + p_i(v-1))$
    $+ \lambda_i(v, v+1)(v \cdot \pi_i(v) - p_i(v) - v \cdot \pi_i(v+1) + p_i(v+1))\big]$.
  - Subect to:
    - $\Lambda \geq 0$.[7]
- Subject to:
  - $\pi_1(0) = p_1(0) = 0$ (Individual Rationality).
  - $\pi_i(v) \leq 1$, for all i, v (Feasibility).

Again by strong duality, the optimal dual variables in the solution solve the following min-max program (the same as Max-Min1 with the min and max flipped):

Min-Max1:

- Minimize: $\mathcal{L}_{\max}(\Lambda)$, where $\mathcal{L}_{\max}(\Lambda)$ is the value of the following LP1(Λ):
  - Maximize $\sum_{i,v} f_i(v)p_i(v) + \sum_i \gamma_i(b_i\pi_i(\bar{v}) - p_i(\bar{v}))$
    $+ \sum_{i,v} \big[\alpha_i(v)(v\pi_i(v) - p_i(v) - v\pi_{i-1}(v) + p_{i-1}(v))$
    $+ \lambda_i(v, v-1)(v \cdot \pi_i(v) - p_i(v) - v \cdot \pi_i(v-1) + p_i(v-1))$
    $+ \lambda_i(v, v+1)(v \cdot \pi_i(v) - p_i(v) - v \cdot \pi_i(v+1) + p_i(v+1))\big]$.
  - Subect to:
    - $\pi_1(0) = p_1(0) = 0$ (Individual Rationality).
    - $\pi_i(v) \leq 1$, for all i, v (Feasibility).
- Subject to: $\Lambda \geq 0$.

Cai et al. [2016] characterize possible solutions to Min-Max1 as *flows*. Specifically, they show the following:

Proposition 2.2 ([Cai et al., 2016]). *$\mathcal{L}_{\max}(\Lambda) < \infty$ if and only if Λ forms a* flow. *That is, $\Lambda \geq 0$ and for all $(i, v) \neq (1, 0)$:*

$$f_i(v) + \lambda_i(v+1, v) + \lambda_i(v-1, v) + \alpha_{i+1}(v) = \lambda_i(v, v-1) + \lambda_i(v, v+1) + \alpha_i(v) + I(v = \bar{v}) \cdot \gamma_i.$$

Therefore, Min-Max1 clearly attains its minimum at a flow. They further define[8]

$$\Phi_i^{\Lambda}(v) = v - \frac{\lambda_i(v+1, v) - \lambda_i(v-1, v)}{f_i(v)}, \text{ for all } v < \bar{v},$$

$$\Phi_i^{\Lambda}(\bar{v}) = \bar{v} - \frac{-\lambda_i(\bar{v}-1, \bar{v}) + \gamma_i(\bar{v} - b)}{f_i(v)},$$

(for simplicity of notation in future sections, we will define $\Phi_1^{\Lambda}(0) = -\infty$,) and observe the following:

[7] By this we mean $\lambda_i(v, v') \geq 0$, $\alpha_i(v) \geq 0$, and $\gamma_i \geq 0$ for all i, v, v'.

[8] Technically, they did not consider budgets, but the derivation of $\Phi^{\Lambda}(\cdot)$ used here is exactly along their calculations.

Proposition 2.3 ([Cai et al., 2016]). *For any flow Λ, $\mathcal{L}_{\max}$ is equal to the value of LP2(Λ), where LP2(Λ) := Maximize $\sum_i \sum_v f_i(v) \cdot \pi_i(v) \cdot \Phi_i^{\Lambda}(v)$ (Expected Virtual Welfare), subject to:*

- $\pi_1(0) = p_1(0) = 0$ *(Individual Rationality).*
- $\pi_i(v) \leq 1$*, for all* i, v *(Feasibility).*

So $\Phi^{\Lambda}(\cdot)$ can be thought of as a *virtual valuation function*, and Min-Max1 can be rewritten as: Min-Max2:

- Minimize: $\mathcal{L}_{\max}(\Lambda)$.
- Subject to:
  - $\Lambda \geq 0$.
  - $f_i(v) + \lambda_i(v+1, v) + \lambda_i(v-1, v) + \alpha_{i+1}(v) = \lambda_i(v, v-1) + \lambda_i(v, v+1) + \alpha_i(v) + I(v = \bar{v}) \cdot \gamma_i,$ for all $(i, v) \neq (1, 0)$.

Theorem 2.4 (Strong Duality). *The following are equivalent:*

- Λ *solves Min-Max2 and* $(\vec{\pi}, \vec{p})$ *solves LP1.*
- $(\vec{\pi}, \vec{p})$ *solves LP2(Λ) and* $(\vec{\pi}, \vec{p})$ *and* Λ *satisfy complementary slackness.*[9]

Our approach will be the following: First, we'll characterize possible solutions for Min-Max2. Then, we'll characterize the mechanisms that can possibly maximize virtual welfare and satisfy complementary slackness (for these solutions). In the following sections, we'll reference Min-Max2, LP2(Λ), and $\mathcal{L}_{\max}(\Lambda)$. We conclude with some quick observations about LP2(Λ):

Observation 1 ([Cai et al., 2016]). *Every optimal solution to LP2(Λ) has $\pi_i(v) = 1$ if $\Phi_i^{\Lambda}(v) > 0$, and $\pi_i(v) = 0$ if $\Phi_i^{\Lambda}(v) < 0$ (if $\Phi_i^{\Lambda}(v) = 0$, then any $\pi_i(v) \in [0, 1]$ is possible). Therefore, $\mathcal{L}_{\max}(\Lambda) = \sum_i \sum_v \max\{f_i(v) \cdot \Phi_i^{\Lambda}(v), 0\}$.*

Observation 2 (Folklore). *If a bidder with value v and $v' \neq v$ are both indifferent between two (allocation, price) pairs, then the two (allocation, price) pairs must be the same.*

Definition 2.5 (Ironed inverval). We say that $[v, w]$ is an *ironed interval* within budget b_i for dual Λ if:

- $\lambda_i(x, x+1) > 0$ for all $x \in [v, w-1]$.
- $\lambda_i(x, x-1) > 0$ for all $x \in [v+1, w]$.
- $\lambda_i(v-1, v) = \lambda_i(w, w+1) = 0$.

By Observation 2, any mechanism satisfying complementary slackness with Λ awards the same (allocation, price) pair to an entire ironed interval.

3 WARMUP: PUBLIC BUDGET

In this section, we'll consider the case $k = 1$, which can be interpreted as a single buyer with a publicly known budget (b_1) and a value drawn from a distribution with density $f_1(\cdot)$. The outline of our approach is as follows. Due to space constraints, many proofs can be found in Appendix B. None of the proofs are overly technical, but the point of this warm-up is to give a high-level overview of the approach in a less technical setting, not to give a complete proof of known results.

1. In Section 3.1, we'll characterize possible solutions to Min-Max2. We'll do this by defining a sequence of elementary operations that can only *decrease* $\mathcal{L}_{\max}(\Lambda)$, and conclude that there exists a solution to Min-Max2 resulting from these operations.
2. In Section 3.2, we'll characterize what mechanisms can possibly satisfy Strong Duality with a dual from Section 3.1. Strong Duality guarantees an optimal mechanism of this format.

[9] That is, $\lambda_i(v, v-1) > 0 \Rightarrow v \cdot \pi_i(v) - p_i(v) = v \cdot \pi_i(v-1) - p_i(v-1)$, etc.

(3) In Section 3.3, we'll strengthen the characterization when $f_1(\cdot)$ is regular or DMR.[10]

3.1 Warmup: Characterizing the Optimal Dual

Here, we characterize possible solutions to Min-Max2. Note that as there is only one possible budget, there are no $\vec{\alpha}$ variables, and there is only one γ multiplier. So we will simplify notation and just discuss $\lambda(v, v')$ and γ (and drop the subscript of 1). Some of our intermediate lemmas are more general though, and will be stated with general subscripts when appropriate. The plan of attack is to show that there is always an optimal dual resulting in monotone non-decreasing virtual values. We begin with a lemma of Cai et al. [2016], which describes an "ironing" procedure that can be used to modify any flow-conserving dual.

Lemma 3.1 ([Cai et al., 2016]). *For any v, x, define $\lambda'(\cdot,\cdot)$ so that $\lambda'(w, w') = \lambda(w, w') + x$ whenever $\{w, w'\} = \{v, v-1\}$, and $\lambda'(w, w') = \lambda(w, w')$ otherwise. Then if $\Lambda = (\lambda, \gamma)$ is flow-conserving, $\Lambda' = (\lambda', \gamma)$ is also flow-conserving, and the following hold:*

- $\Phi^{\Lambda'}(w) = \Phi^{\Lambda}(w)$, *for all* $w \notin \{v, v-1\}$.
- $\Phi^{\Lambda'}(v) = \Phi^{\Lambda}(v) + x/f(v)$.
- $\Phi^{\Lambda'}(v-1) = \Phi^{\Lambda}(v-1) - x/f(v-1)$.
- $f(v) \cdot \Phi^{\Lambda'}(v) + f(v-1) \cdot \Phi^{\Lambda'}(v-1) = f(v) \cdot \Phi^{\Lambda}(v) + f(v-1) \cdot \Phi^{\Lambda}(v-1)$.

More specifically, Lemma 3.1 states that we can always add a cycle to any flow-conserving dual, and it will preserve the average virtual value among the two values. Now, we want to argue that this procedure can be used to iron out any non-monotonicities, and that this can only make the dual better (this and all other missing proofs in Appendix B).

Corollary 3.2. *There exists a solution to Min-Max2, Λ, such that the resulting $\Phi^{\Lambda}(\cdot)$ is monotone non-decreasing, and $\lambda(v-1, v) > 0 \Rightarrow \Phi^{\Lambda}(v) = \Phi^{\Lambda}(v-1)$. There also exists a solution to Min-Max2 such that the resulting $\Phi^{\Lambda}(\cdot)$ has $f(\cdot) \cdot \Phi^{\Lambda}(\cdot)$ monotone non-decreasing, and $\Lambda(v-1, v) > 0 \Rightarrow f(v) \cdot \Phi^{\Lambda}(v) = f(v-1) \cdot \Phi^{\Lambda}(v-1)$.*

In light of Corollary 3.2, we'll introduce the following terminology. Note that below, we are ironing so that $f_i(\cdot) \cdot \Phi_i^{\Lambda}(\cdot)$ is monotone non-decreasing, and *not* so that $\Phi_i^{\Lambda}(\cdot)$ is monotone non-decreasing.

Definition 3.3 (Proper ironing). We say that a dual solution Λ is *properly ironed* if for all i, $\Phi_i^{\Lambda}(\cdot) \cdot f_i(\cdot)$ is monotone non-decreasing and $\lambda_i(v-1, v) > 0 \Rightarrow f_i(v) \cdot \Phi_i^{\Lambda}(v) = f_i(v-1) \cdot \Phi_i^{\Lambda}(v-1)$.

Now, we turn to addressing what properties we can guarantee in an optimal dual related to the budget. We'll again name these conditions. Essentially one should interpret budget-feasible as meaning "it is possible to award $\bar{v}$ the item with non-zero probability in an IC mechanism that respects the budget constaint, while solving LP2(Λ) and satisfying complementary slackness."

Definition 3.4 (Budget-feasible dual). We say that Λ is *budget-feasible* at i if (for the warmup, we will just say budget-feasible):

- $\Phi_i^{\Lambda}(b_i) \geq 0$.
- Either $\gamma_i = 0$, or $\Phi_i^{\Lambda}(b_i - 1) \leq 0$.
- There exists an $x \geq b_i$ with $\lambda_i(x-1, x) = 0$.

[10] We say that a discrete distribution with integral support, PDF f and CDF F is *regular* if $v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing. We say that it is DMR if $vf(v) - (1 - F(v))$ is monotone non-decreasing.

Lemma 3.5. *If Λ is properly ironed, but* not *budget-feasible at i, then all feasible, IC, budget-respecting $(\vec{\pi}, \vec{p})$ that solve LP2(Λ) and satisfy complementary slackness have $\pi_i(v) = p_i(v) = 0$ for all v (if such a $(\vec{\pi}, \vec{p})$ even exists).*

Proof. Consider any mechanism with $\pi_i(\bar{v}) > 0$. First, consider the case where $\Phi_i^\Lambda(b_i) < 0$. Observe also that if $\lambda_i(b_i + 1, b_i) = 0$, then $\Phi_i^\Lambda(b_i) \geq b_i$. This can be seen immediately from the definition of Φ^Λ. So we may also conclude that $\lambda_i(b_i + 1, b_i) > 0$. Then to solve LP2($\Lambda$), we must have $\pi_i(b_i) = 0$. But to satisfy complementary slackness with $\lambda_i(b_i + 1, b_i)$ and also IC, we must have $p_i(\bar{v}) \geq (b_i + 1) \cdot \pi_i(\bar{v})$, so no mechanism can be feasible, IC, budget-respecting, solve LP2(Λ) and satisfy complementary slackness with $\pi_i(\bar{v}) > 0$.

The remaining properties characterize the ironed interval containing b_i in case the budget constraint binds, and we can actually prove that they hold no matter what. Suppose that $\gamma_i > 0$ and $\Phi_i^\Lambda(b_i - 1) > 0$. In this case, any solution to LP2(Λ) must have $\pi_i(b_i - 1) = 1$ and $p_i(b_i - 1) \leq b_i - 1$. $(i, \bar{v})$ will certainly prefer this to any option with $p_i(\bar{v}) = b_i$, so no IC solution to LP2(Λ) results in the i^{th} budget constraint binding and complementary slackness is violated.

Now if $\lambda_i(x-1, x) > 0$ for all $x \geq b_i$, by Observation 2 any mechanism satisfying complementary slackness necessarily awards the same (allocation, price) pair to all (i, x) with $x \geq b_i - 1$. If such a mechanism is additionally IC, the price must be at most $b_i - 1$, and therefore $p_i(\bar{v}) < b_i$.

As $\gamma_i > 0$, complementary slackness fails in either case. So there is no IC mechanism that satisfies complementary slackness with such a Λ at all.

Finally, briefly observe that if Λ is properly ironed and $\gamma_i = 0$, then $\Phi_i^\Lambda(\bar{v}) = \bar{v} > \Phi_i^\Lambda(\bar{v} - 1)$, so $\lambda_i(\bar{v} - 1, \bar{v}) = 0$ and the final bullet is also satisfied when $\gamma_i = 0$.

Just to summarize: we have shown that the second and third conditions above must hold for any Λ that can possibly have a feasible, IC, budget-respecting mechanism solve LP2(Λ) and satisfy complementary slackness. If the first condition fails, to hold, then all such mechanisms must have $\pi_i(v) = 0$ for all v. □

Corollary 3.6. *With a single budget, there exists an optimal solution to Min-Max2, Λ, such that:*

- *Λ is properly ironed.*
- *Λ is budget-feasible.*

Proof. We just need to observe that clearly the optimal mechanism is not to just award the item with probability 0 and make zero revenue. So by Corollary 3.2, there exists a solution to Min-Max2, Λ, that is properly ironed. By Lemma 3.5, it must also be budget-feasible. □

3.2 Warmup: Characterizing the Optimal Primal

In this section, we characterize what any mechanism that solves LP2(Λ) and satisfies complementary slackness must look like for a dual of the form guaranteed by Corollary 3.6.

Theorem 3.7 ([Chawla et al., 2011]). *The solution to LP1 corresponds to a menu with at most two options. The first offers the item with probability 1 at price b. The second offers the item with probability of $q < 1$ at price $p < qb$. The second option need not exist.*

Intuitively, when the budget doesn't bind, a posted price is optimal. Here is what's going on in the $\gamma > 0$ case: if there is no ironing around b, or if b is at the lower end of an ironed interval, then we could just set price b. Unfortunately, b might lie in the strict interior of an ironed interval, in which case such a mechanism wouldn't satisfy complementary slackness (beacuse the Rightwards IC constraint from $b - 1$ to b wouldn't be tight, yet $\lambda_i(b - 1, b) > 0$). Similarly, we can't just set price equal to the minimum v such that $\Phi^\Lambda(v) > 0$, as then the budget constraint wouldn't bind. So we may need to award the item to types in the ironed interval containing b with probability strictly between 0 and 1 in order to get the budget constraint to bind *and* respect all ironed intervals. But fortunately, Corollary 3.6 guarantees that no additional randomization is necessary.

3.3 Warmup: Regular Distributions

In this section, we prove a tighter characterization in the case that F is regular or DMR. We first prove that a tighter characterization is not possible in general with an example.

Example 3.8. Let $b = 11$, $f(20) = f(2) = 1/2$. Then the optimal deterministic mechanism simply sets a price of 11 on the item and achieves a revenue of 5.5. A strictly better mechanism offers the two options $(1, 11)$ and $(1/2, 1)$, and achieves a revenue of 6 (this is in fact optimal).

So Example 3.8 and Theorem 3.7 together say that two options are both necessary and sufficient to characterize the optimal mechanism for arbitrary F. But it turns out the additional option is not necessary in the special case that F is regular or DMR [Laffont and Robert, 1996]. The proof for DMR is more straight-forward, which we include in the body. The proof for regular is nearly identical after one additional technical lemma (which essentially shows that regular distributions are also DMR below the Myerson reserve). We will use the notation $\varphi(v) := v - \frac{1-F(v)}{f(v)}$ to denote Myerson's virtual valuation (which can be achieved by a flow in the discrete setting as well).

Proposition 3.9. *Let F be DMR or regular, and let r denote the maximum $v \leq b$ with $\Phi^\Lambda(v-1) \leq 0$. Then there exists an optimal solution to Min-Max2, Λ, such that:*

- *Λ is properly ironed.*
- *$\Phi^\Lambda(b) \geq 0$.*
- *$\lambda(r-1, r) = 0$.*
- *Either $\gamma = 0$ or $r = b$.*

Proof (for DMR, for regular see Appendix B). For a fixed γ, consider the following $\vec{\lambda}$: $\lambda(v, v-1) = 1 - \gamma - F(v-1)$. It is easy to verify that Λ is flow-conserving, and results in $\Phi^\Lambda(\bar{v}) = \bar{v} - \frac{(\bar{v}-b)\gamma}{f(\bar{v})}$, $\Phi^\Lambda(v) = v - \frac{1-F(v)-\gamma}{f(v)}$ for $v < \bar{v}$. Observe that as F is DMR, we have:

$$vf(v) - (1 - F(v)) \geq (v-1)f(v-1) - (1 - F(v-1)) \Rightarrow f(v)\Phi^\Lambda(v) \geq f(v-1)\Phi^\Lambda(v-1), \forall v < \bar{v}.$$

So except for at $\bar{v}$, $f(\cdot)\Phi^\Lambda(\cdot)$ is already monotone non-decreasing. It is easy to see the procedure described in Corollary 3.2 will result in a single ironed interval $[v^*, \bar{v}]$ at the top, and $\lambda(v-1, v) = 0$ for all $v \leq v^*$. Observe that we must have $v^* \geq b$, as otherwise complementary slackness cannot be satisfied: $v^* \neq \bar{v}$ implies $\gamma > 0$, and $v^* < b$ implies $p(\bar{v}) < b$ in any IC mechanism. For all $v \leq v^*$, we have $\lambda(v-1, v) = 0$. In particular, as $r \leq b \leq v^*$, we have $\lambda(r-1, r) = 0$.

Finally, if $\gamma > 0$, we must have $r = b$ due to complementary slackness, as no IC mechanism can simultaneously optimize LP2(Λ) and have $p(\bar{v}) = b\pi(\bar{v})$ - if $r < b$, then $\Phi^\Lambda(b-1) > 0$.

Characterizing the optimal dual for regular distributions requires one additional technical lemma (Lemma 3.10), which essentially says that regular distributions are also DMR below the Myerson reserve. □

Lemma 3.10. *Let F be regular, $v \leq w$, and $\varphi(v) \leq 0$. Then $vf(v) - (1 - F(v)) \leq wf(w) - (1 - F(w))$.*

Proposition 3.11. *For any (Λ) of the form guaranteed by Proposition 3.9, the mechanism that sets price r (formally, $(\pi(v), p(v)) = (1, r)$ for all $v \geq r$, $(\pi(v), p(v)) = (0, 0)$ for all $v < r$) optimizes LP2(Λ) and satisfies complementary slackness.*

Proof. All IC constraints are tight, except for the one between $r-1$ and r. However, we are guaranteed that $\lambda(r-1, r) = 0$. Moreover, we are also guaranteed that whenever $\gamma > 0$, $r = b$ and therefore $p(\bar{v}) = b$. So complementary slackness is satisfied.

It's also easy to see that the proposed mechanism solves LP2(Λ): it awards the item to every type with non-negative virtual value. □

The last remaining observation is that we can characterize exactly what r is:

Theorem 3.12 ([Laffont and Robert, 1996]). *Let F be regular or DMR, and let x be the minimum v such that $\varphi(v) \geq 0$. Then the optimal mechanism sets reserve* $\min\{b, x\}$.

4 MAIN RESULT: PRIVATE BUDGET

In this section, we consider the general case. The outline of our approach follows that of the warm-up, but each part will be more technically challenging. Many proofs of technical lemmas/observations/etc. are omitted and can be found in Appendix C.

(1) In Section 4.1, we'll characterize possible solutions to Min-Max2. Similarly to the warm-up, we'll define a sequence of elementary operations that can only decrease $\mathcal{L}_{\max}(\Lambda)$ and conclude that there exists an optimal dual resulting from these operations. With multiple budgets, we will have to define additional operations beyond just the ironing considered in the warm-up to guarantee a "nice" dual.
(2) In Section 4.2, we'll characterize what mechanisms (primal solutions) can possibly satisfy Strong Duality with the duals found in Section 4.1. There are two parts to our approach here. The first half looks a lot like the solution to the FedEx Problem [Fiat et al., 2016]: once we set a menu for budget b_1, we describe how to generate menus for the larger budgets that respects/makes tight all necessary IC constraints (but may not respect all budget constraints). The second half is proving that there exists a menu for b_1 such that the mechanism resulting from this procedure respects/makes tight all necessary budget constraints as well.
(3) We defer to Appendix D, a stronger characterization in the case that each $f_i(\cdot)$ is DMR [Che and Gale, 2000] - the proof is quite similar to that of Theorem 3.12.[11]

4.1 The General Case: Characterizing the Optimal Dual

Here, we characterize the possible solutions to Min-Max2 in the general case. Things will be more technical here than in Section 3.1 due to the multiple budgets (and in particular, the existance of the $\vec{\alpha}$ variables). We'll describe a set of modifications that can be performed on any potential dual solution that can only decrease $\mathcal{L}_{\max}(\Lambda)$. At a high level, the modifications are the following:

(1) Properly ironing Λ can only help.
(2) If (i, v) isn't at the bottom of an ironed interval, then "splitting" $\alpha_i(v)$ between $\alpha_i(v+1)$ and $\alpha_i(v-1)$ can only help.
(3) Subject to (1) above, if $\Phi_i^{\Lambda}(v-1) < 0$, then "boosting" $\alpha_i(v)$ can only help.
(4) Subject to (1), if $\Phi_i^{\Lambda}(v-1) > 0$ and $\alpha_i(v) > 0$, then "re-routing" flow from $\alpha_i(v)$ to $\alpha_i(v-1)$ can only help.

Below, we'll make these claims formal and conclude the subsection with a characterization of potentially optimal solutions to Min-Max2.

[11] We leave as an open problem whether the strengthened characterization extends to regular distributions as well.

4.1.1 Characterizing The Optimal Dual: Proper Ironing can Only Help. Here, we prove the analogy of Corollary 3.6. The majority of the proof is identical and omitted. The only difference is that we can't necessarily guarantee that Λ is budget-feasible at i for all i, because it may simply be that $\Phi_i^\Lambda(\bar{v}) = 0$ and the item is never awarded (in which case the mechanism need not charge any budget-respecting price).

Corollary 4.1. *There exists an optimal solution Λ to Min-Max2 such that:*

- *Λ is properly ironed.*
- *For all i, either Λ is budget-feasible at i, or $\Phi_i^\Lambda(\bar{v}) = 0$.*

Proof. That Λ may be properly ironed follows exactly the proof of Corollary 3.6. For the second bullet, we observe that by exactly the same proof as Corollary 3.6, if we are to have $\pi_i(\bar{v}) = 1$, then we necessarily have Λ budget-feasible at i. If $\Phi_i^\Lambda(\bar{v}) > 0$, then we necessarily have $\pi_i(\bar{v}) = 1$ in an solution to LP2(Λ), and we can invoke this argument. So the only way that Λ can be optimal yet not budget-feasible at i is if $\Phi_i^\Lambda(\bar{v}) = 0$. □

4.1.2 Characterizing the Optimal Dual: How to Set $\vec{\alpha}$. In this section, we'll see what $\vec{\alpha}$ must look like in an optimal solution. Let's first consider three possible modifications to Λ. We'll call the first *boosting*, the second *re-routing*, and the third *splitting*. All modifications will preserve that Λ is flow-conserving.

Definition 4.2. Consider any (i, v). We say that we are *boosting* $\alpha_i(v)$ by c if we increase $\alpha_i(v)$ by c, decrease $\lambda_i(w, w-1)$ by c for all $w \leq v$, and increase $\lambda_{i-1}(w, w-1)$ by c for all $w \leq v$.

Observation 3. *Boosting $\alpha_i(v)$ by c increases $f_i(w) \cdot \Phi_i^\Lambda(w)$ by c, and decreases $f_{i-1}(w) \cdot \Phi_{i-1}^\Lambda(w)$ by c for all $w < v$.*

Definition 4.3. Consider any (i, v) such that $\alpha_i(v) > 0$. We say that we are *re-routing* $\alpha_i(v)$ by c if we decrease $\alpha_i(v)$ by c, and increase $\lambda_i(v, v-1)$, $\alpha_i(v-1)$, and $\lambda_{i-1}(v-1, v)$ by c.

Observation 4. *Re-routing $\alpha_i(v)$ by c decreases $f_i(v-1) \cdot \Phi_i^\Lambda(v-1)$ by c and increases $f_{i-1}(v) \cdot \Phi_{i-1}^\Lambda(v)$ by c.*

Definition 4.4. Consider any (i, v), $v \in (0, 1)$ such that $\alpha_i(v) > 0$, $\lambda_i(v-1, v) > 0$, and $\lambda_i(v+1, v) > 0$. We say that we are *splitting* $\alpha_i(v)$ by c if we:

- Increase $\alpha_i(v+1)$ and $\lambda_{i-1}(v+1, v)$ by $c/2$, and decrease $\lambda_i(v+1, v)$ and $\alpha_i(v)$ by $c/2$.
- Increase $\alpha_i(v-1)$ and $\lambda_{i-1}(v-1, v)$ by $c/2$, and decrease $\lambda_i(v-1, v)$ and $\alpha_i(v)$ by $c/2$.

If $\alpha_i(\bar{v}) > 0$, $\lambda_i(\bar{v}-1, \bar{v}) > 0$, and $\gamma_i > 0$, we say that we are splitting $\alpha_i(\bar{v})$ by c if we:

- increase $\alpha_i(\bar{v}-1)$ and $\lambda_{i-1}(\bar{v}-1, \bar{v})$ by c, and decrease $\lambda_i(\bar{v}-1, \bar{v})$ and $\alpha_i(\bar{v})$ by c.
- increase $\alpha_i(\bar{v})$ and γ_{i-1} by $\frac{c}{\bar{v}-b_i}$, and decrease γ_i by $\frac{c}{\bar{v}-b_i}$.

Observation 5. *For $v < \bar{v}$, splitting $\alpha_i(v)$ by c doesn't change $\Phi_i^\Lambda(\cdot)$ anywhere.*

Observation 6. *Splitting $\alpha_i(\bar{v})$ by c decreases $\Phi_{i-1}^\Lambda(\bar{v})$ by $c\left(\frac{\bar{v}-b_{i-1}}{\bar{v}-b_i} - 1\right)$.*

It's obvious that we can always split any $\alpha_i(v)$ without harming the quality of Λ (as all virtual values remain unchanged or decrease). We'll now prove that boosting and re-routing can only help as well, as long as Λ is properly ironed.

Proposition 4.5. *Let Λ be any properly ironed dual, and let $f_i(v-1) \cdot \Phi_i^\Lambda(v-1) = -\epsilon < 0$ for some $i > 1$, $v > 0$. If Λ' denotes Λ after boosting $\alpha_i(v)$ by ϵ, then $\mathcal{L}_{\max}(\Lambda') \leq \mathcal{L}_{\max}(\Lambda)$.*

Proof. Because $f_i(\cdot) \cdot \Phi_i^\Lambda(\cdot)$ is presently monotone non-decreasing, $f_i(w) \cdot \Phi_i^\Lambda(w) \le -\epsilon$ for all $w \le v-1$. The solution to LP2(Λ) therefore necessarily has $\pi_i(w) = 0$ for all $w \le v$. Intuitively what we are doing is this: boosting $\alpha_i(v)$ by ϵ causes $f_i(w) \cdot \Phi_i^\Lambda(w)$ to increase for all $w < v$. A priori, this could be bad: increasing virtual values could cause $\mathcal{L}_{\max}$ to increase. But because all considered virtual values are already negative, they don't contribute to $\mathcal{L}_{\max}$ anyway, as $\pi_i(w)$ is 0 regardless. So in a sense we can increase $\Phi_i^\Lambda(v-1)$ to 0 "for free." This certainly doesn't hurt, but maybe it helps since it allows us to decrease $\Phi_{i-1}^\Lambda(w)$ for all $w < v$, which can only help.

A little more formally, by Observation 1, $\mathcal{L}_{\max}(\Lambda) - \mathcal{L}_{\max}(\Lambda') = \sum_{w<v} \max\{0, f_i(w) \cdot \Phi_i^\Lambda(w)\} - \max\{0, f_i(w) \cdot \Phi_i^\Lambda(w) + \epsilon\} + \max\{0, f_{i-1}(w) \cdot \Phi_{i-1}^\Lambda(w)\} - \max\{0, f_{i-1}(w) \cdot \Phi_{i-1}^\Lambda(w) - \epsilon\}$. As $f_i(w) \cdot \Phi_i^\Lambda(w) \le -\epsilon$ for all $w < v$, the first two terms are always both zero, and the difference of the last two terms is clearly non-negative. So $\mathcal{L}_{\max}(\Lambda') \le \mathcal{L}_{\max}(\Lambda)$, and we have only improved. □

Proposition 4.6. *Let Λ be any properly ironed dual and suppose that* $\min\{\alpha_i(v), f_i(v-1) \cdot \Phi_i^\Lambda(v-1)\} = \epsilon > 0$ *for some $i > 1$, $v > 0$. If Λ' denotes Λ after re-routing $\alpha_i(v)$ by ϵ, then* $\mathcal{L}_{\max}(\Lambda') \le \mathcal{L}_{\max}(\Lambda)$.

The proof is very similar to that of Proposition 4.5 and can be found in Appendix C.

4.1.3 Characterizing the Optimal Dual: Final Steps. Here, we combine the propositions from the previous sections with one extra step to prove our main theorem characterizing the dual. We already know that we may take the optimal solution to be properly ironed and for all i to either be budget-feasible or have $\Phi_i^\Lambda(\bar{v}) = 0$. Throughout this section, we'll use v_i^+ to denote the minimum v such that $\Phi_i^\Lambda(v) > 0$, if one exists.

We can derive some additional properties with little extra work using the results of Section 4.1.2. We first separate out one extra property we'd like that is a bit more technical to prove. Intuitively, we want to define two budgets to be *linked* if for every $(\vec{\pi}_{i-1}, \vec{p}_{i-1})$, there is at most one $(\vec{\pi}_i, \vec{p}_i)$ such that $(\vec{\pi}, \vec{p})$ can possibly satisfy complementary slackness and solve LP2(Λ). We provide formal conditions below, and prove that they have the desired semantic meaning.

Definition 4.7. We say that Λ *links* budgets i and $i-1$ if for every ironed interval within budget i, $[w, v]$, either $\alpha_i(w) > 0$ or $\Phi_i^\Lambda(w-1) > 0$. Formally, $\lambda_i(w-1, w) = 0 \Rightarrow \alpha_i(w) > 0 \vee \Phi_i^\Lambda(w-1) > 0$.

There is one exception to the rule: if $\Phi_i^\Lambda(\bar{v}) = 0$ and $[b_i, \bar{v}]$ is an ironed interval within budget i, then Λ *does not* link budgets i and $i-1$, even if the above conditions hold.

Lemma 4.8. *If Λ links budgets i and $i-1$, and $(\vec{\pi}, \vec{p})$ and $(\vec{\pi}', \vec{p}')$ both solve LP2(Λ) and satisfy complementary slackness, and $(\vec{\pi}_{i-1}, \vec{p}_{i-1}) = (\vec{\pi}'_{i-1}, \vec{p}'_{i-1})$, then $(\vec{\pi}_i, \vec{p}_i) = (\vec{\pi}'_i, \vec{p}'_i)$ as well.*

Proof. We basically observe that the (allocation,price) pair for each ironed interval has two variables and two non-degenerate linear constraints it must satisfy. Specifically, for any ironed interval $[w, v]$, with $v < \bar{v}$ and $\Phi_i^\Lambda(w) \le 0$, we know that $(i, v+1)$ is indifferent between their own (allocation, price) and that of $(i-1, v+1)$ by complementary slackness (as $\alpha_i(v+1) > 0$ by linkage). We also know that $(i, v+1)$ is indifferent between their own (allocation, price) and that of (i, v) (again by complementary slackness). So $(i, v+1)$ is indifferent between the (allocation, price) awarded to the ironed interval and that awarded to $(i-1, v+1)$. Similarly by linkage, we know that $\alpha_i(w) > 0$, and therefore (i, w) is indifferent between the (allocation, price) awarded to the ironed interval and that awarded to $(i-1, w)$. So this results in two non-degenerate linear equations for the (allocation, price) awarded to the ironed interval, and its solution is unique.

Now consider the minimum w such that $\Phi_i^\Lambda(w) > 0$, if one exists. The reasoning above says that $(\pi_i(w-1), p_i(w-1))$ is already determined. So as we must have $\pi_i(w) = 1$ to solve LP2(Λ), and (i, w) indifferent between her (allocation, price) and that of $(i, w-1)$, this is again two variables and two constraints. IC then implies that all $v > w$ receive the same (allocation, price) as well, so the entire menu is determined.

If instead $\Phi_i^\Lambda(\bar{v}) = 0$, then we clearly have $\gamma_i > 0$. In this case, all (allocation, price) pairs for all ironed intervals except for $[v_i^+, \bar{v}]$ have been determined. The (allocation, price) awarded to this interval is again governed by two constraints: that $p_i(\bar{v}) = b_i \cdot \pi_i(\bar{v})$, and that (i, v_i^+) is indifferent between the (allocation, price) awarded to this ironed interval, and that awarded to $(i, v_i^+ - 1)$. As long as $v_i^+ \neq b_i$, these two constraints are non-degenerate and have a unique solution. If $v_i^+ = b_i$, then the constraints are degenerate - either no solution exists or there could be infintiely many solutions. □

Corollary 4.9. *Suppose that Λ links budgets i and $i-1$, and $(\vec{\pi}, \vec{p})$ solves LP2(Λ) and satisfies complementary slackness with $(\vec{\pi}_{i-1}, \vec{p}_{i-1}) = (\vec{0}, \vec{0})$.*

- *If $\Phi_i^\Lambda(\bar{v}) = 0$ then $(\vec{\pi}_i, \vec{p}_i) = (\vec{0}, \vec{0})$ as well.*
- *If $\Phi_i^\Lambda(\bar{v}) > 0$,then*
  - *$\pi_i(v) = p_i(v) = 0$ for all $v < v_i^+$.*
  - *$(\pi_i(v), p_i(v)) = (1, v_i^+)$ for all $v \geq v_i^+$.*

Theorem 4.10. *There exists an optimal solution Λ for Min-Max2 with the following properties:*

- *Λ is properly ironed.*
- *For all i, either Λ is budget-feasible at i or $\Phi_i^\Lambda(\bar{v}) = 0$.*
- *$\Phi_i^\Lambda(v) \geq 0$ for all v and $i > 1$.*
- *$\alpha_i(v) > 0 \Rightarrow (\Phi_i^\Lambda(v-1) = 0$ AND $\lambda_i(v-1, v) = 0)$.*
- *$\gamma_i > 0 \Rightarrow \Phi_{i-1}^\Lambda(\bar{v}) \leq 0$.*

Bullet points two and five necessarily hold in any solution to Min-Max2. Whenever bullet points one, three, or four don't hold, there exist elementary operations (proper ironing, boosting, rereouting, splitting) that can only improve the dual. So these operations can be repeatedly performed until all bullet points hold.

4.2 The General Case: Characterizing the Optimal Primal

In this section, we show that there exists a mechanism of a certain format that solves LP2(Λ) and satisfies comlementary slackness for a dual of the form guaranteed by Theorem 4.10. We begin by proposing one method for generating a menu for types with budget b_i from the menu for types with budget b_{i-1}. It is essentially assuming that budgets i and $i-1$ are *linked*, and propagating the equalities.

Generate-Menu($i, \vec{\pi}_{i-1}, \vec{p}_{i-1}$):

- If $\Phi_i^\Lambda(\bar{v}) > 0$, let v_i^+ be the miminum v such that $\Phi_i^\Lambda(v) > 0$.
- Else, let v_i^+ be the maximum v such that $\lambda_i(v-1, v) = 0$ (bottom of the ironed interval in budget i containing $\bar{v} : v_i^+$ might be 0).
  - Consider any ironed interval $[w, v]$ with $w < v_i^+$.
    * If $(i-1, w)$ and $(i-1, v)$ receive the same (allocation, price) under $(\vec{\pi}_{i-1}, \vec{p}_{i-1})$, then copy this option for the entire ironed interval $[w, v]$ for budget i. Formally: if $(\pi_{i-1}(w), p_{i-1}(w)) = (\pi_{i-1}(v), p_{i-1}(v))$, then set $(\pi_i(x), p_i(x)) = (\pi_{i-1}(x), p_{i-1}(x))$ for all $x \in [w, v]$.
    * It $(i-1, w)$ and $(i-1, v)$ receive different (allocation, price) under $(\vec{\pi}_{i-1}, \vec{p}_{i-1})$, set $(\pi_i(x), p_i(x))$ (for all $x \in [w, v]$) so that (i, w) is indifferent between $(\pi_i(x), p_i(x))$ and $(\pi_{i-1}(w), p_{i-1}(w))$, *and* $(i, v+1)$ is indifferent between $(\pi_i(x), p_i(x))$ and $(\pi_{i-1}(v), p_{i-1}(v))$. Note that there are two variables and two linear constraints, so such a setting of

$(\pi_i(x), p_i(x))$ always exists (although it is not guaranteed that $\pi_i(x) \in [0,1], p_i(x) \geq 0$). Also observe that if $(\pi_{i-1}(v), p_{i-1}(v)) = (\pi_{i-1}(w), p_{i-1}(w))$, then the above is just a special case.[12]

Formally: if $\pi_{i-1}(w) < \pi_{i-1}(v)$, then set $\pi_i(x) = \frac{(v+1)\pi_{i-1}(v+1) - w\pi_{i-1}(w) - p_{i-1}(v+1) + p_{i-1}(w)}{v+1-w}$, and $p_i(x) = (v+1)(\pi_i(v) - \pi_{i-1}(v+1)) + p_{i-1}(v+1) = w(\pi_i(w) - \pi_{i-1}(w)) + p_{i-1}(w)$ for all $x \in [w, v]$.

- If $\Phi_i^\Lambda(\bar{v}) > 0$, set $(\pi_i(x), p_i(x)) = (1, p_{i-1}(v_i^+) + v_i^+(1 - \pi_{i-1}(v_i^+)))$ for all $x \in [v_i^+, \bar{v}]$. This price guarantees that the IC constraint from (i, v_i^+) to $(i-1, v_i^+)$ binds, as does the IC constraint from (i, v_i^+) to $(i, v_i^+ - 1)$ (by how $(\pi_i(v_i^+ - 1), p_i(v_i^+ - 1)$ is determined above).
- Else, set $(\pi_i(x), p_i(x) = (\frac{v_i^+ \pi_{i-1}(v_i^+) - p_{i-1}(v_i^+)}{v_i^+ - b_i}, b_i \cdot \pi_i(x))$ for all $x \geq v_i^+$. This price guarantees that the IC constraints from (i, v_i^+) to $(i-1, v_i^+)$ and (i, v_i^+) to $(i, v_i^+ - 1)$ bind, as does the i^{th} budget constraint bind (but not necessarily that $\pi_i(\bar{v}) \in [0,1], p_i(\bar{v}) \geq 0$).

Note that Generate-Menu$(i, \vec{\pi}_{i-1}, \vec{p}_{i-1})$ might sometimes output probabilities outside of $[0,1]$, negative prices, or have undefined behavior depending on the input. But we claim that when properly seeded with a menu for types with budget b_1, and with a dual of the form guaranteed by Theorem 4.10, the resulting menus have nice properties. Let's begin by proving that complementary slackness is at least satisfied for all IC constraints. Many of the following are essentially what is shown in [Fiat et al., 2016], but in different language.

We begin by proving several small claims to show that when $(\vec{\pi}_i, \vec{p}_i)$ is sufficiently reasonable, Generate-Menu$(i+1, \vec{\pi}_i, \vec{p}_i)$ is somewhat reasonable as well. We break these up into smaller subsections to make the big picture clearer.

4.2.1 Generate Menu: When is the Output Feasible? In this section, we'll prove some small claims to address when we can claim that the output of Generate-Menu has all allocation probabilities in $[0,1]$ and all prices ≥ 0. Most of these claims are already known due to Fiat et al. [2016], and are just presented here in a different language.

Proposition 4.11 (First two bullet points are in [Fiat et al., 2016]). *Let $(\vec{\pi}_i, \vec{p}_i)$ satisfy all leftwards/rightwards IC constraints, and $\pi_i(v) \in [0,1], p_i(v) \geq 0$ for all v, and $(\vec{\pi}_{i+1}, p_{i+1}) =$ Generate-Menu$(i+1, \vec{\pi}_i, \vec{p}_i)$. Then*

- *$\pi_{i+1}(v) \in [0,1]$ and $p_{i+1}(v) \geq 0$ for all $v < v_i^+$.*
- *Further, if $\Phi_{i+1}^\Lambda(\bar{v}) > 0$, then $\pi_{i+1}(\bar{v}) \in [0,1]$ and $p_{i+1}(\bar{v}) \geq 0$.*
- *If Λ is budget-feasible at i, then $\pi_{i+1}(v), p_{i+1}(v) \geq 0$ for all v.*
- *If $v_i^+ \cdot \pi_i(v_i^+) - p_i(v_i^+) \leq v_i^+ - b_{i+1}$, then $\pi_{i+1}(v) \in [0,1]$ and $p_{i+1}(v) \geq 0$ for all v.*

We provide a complete proof of all bullets separated as small lemmas in Appendix C. Below is a useful observation specifically about Generate-Menu when given the trivial menu as input.

Observation 7. *If Λ is budget-feasible at i, then Generate-Menu$(i, \vec{0}, \vec{0})$ outputs a feasible menu.*

So at this point, here is what we know: if we input a feasible menu for budget b_{i-1} into Generate-Menu, and $\Phi_i^\Lambda(\bar{v}) > 0$, then we will definitely get out a feasible menu for budget b_i. If $\Phi_i^\Lambda(\bar{v}) = 0$, then we definitely get a feasible menu for all ironed intervals except the highest. For the highest ironed interval, it's possible that the options on the input menu were "too cheap" to possibly make budget b_i bind with $\pi_i(v_i^+) \in [0,1]$. So the only possible issue with feasibility is **maybe $\pi_i(\bar{v}) \notin [0,1]$ when $\Phi_i^\Lambda(\bar{v}) = 0$.**

[12] This requires observing that in this case we necessarily have $p_{i-1}(v+1) = (v+1)(\pi_{i-1}(v+1) - \pi_{i-1}(w)) + p_{i-1}(w)$.

4.2.2 Generate-Menu: Where does Complementary Slackness Hold? In this section, we'll figure out which constraints definitely satisfy complementary slackness with an output of Generate-Menu, and which we need to keep an eye out for. To ease notation, we will say that a constraint $\vec{a} \cdot \vec{x} \leq b$ with dual variable β is *strongly satisfied* by $\vec{y}$ if $\beta = 0$ and $\vec{a} \cdot \vec{y} \leq b$, or $\beta > 0$ and $\vec{a} \cdot \vec{y} = b$.

Proposition 4.12 ([Fiat et al., 2016]). *If $(\vec{\pi}_1, \vec{p}_1)$ strongly satisfies all leftwards/rightwards IC constraints within budget b_1, then the full $(\vec{\pi}, \vec{p})$ obtained by iterating $(\vec{\pi}_i, \vec{p}_i) =$ Generate-Menu$(i, \vec{\pi}_{i-1}, \vec{p}_{i-1})$ strongly satisfies all leftwards/rightwards/downwards IC constraints.*

So at least we know that all IC constraints are strongly satisfied. Now, let's confirm that Generate-Menu indeed results in a menu that solves LP2(Λ), assuming it outputs a feasible mechanism.

Corollary 4.13. *Let $\pi_1(v) = 1$ if $\Phi_1^\Lambda(v) > 0$ and $\pi_1(v) = 0$ if $\Phi_1^\Lambda(v) < 0$. Let also $(\vec{\pi}_1, \vec{p}_1)$ strongly satisfy all leftwards/rightwards IC constraints, and define $(\vec{\pi}_{i+1}, p_{i+1}) =$ Generate-Menu$(i+1, \vec{\pi}_i, \vec{p}_i)$ for all i. Then if $\pi_i(v) \in [0,1], p_i(v) \geq 0$ for all v, $(\vec{\pi}, \vec{p})$ solves LP2(Λ).*

Let's now parse what possible problems might arise when iterating Generate-Menu from an arbitrary menu for budget b_1 (that at least strongly satisfies the leftwards/rightwards IC constraints within budget b_1, and the budget constraint).

- Definitely, all IC constraints are strongly satisfied.
- If the resulting mechanism is feasible, it definitely solves LP2(Λ) (but we might have $\pi_j(\bar{v}) > 1$ when $\Phi_j^\Lambda(\bar{v}) = 0$).
- If i denotes the minimum i such that $\Phi_i^\Lambda(\bar{v}) > 0$, then all budget constraints for $j < i$ are definitely strongly satisfied (this is by definition in Generate-Menu).
- For the above, if we are able to strongly satisfy the i^{th} budget constraint, then definitely we will strongly satisfy all budget constraints, as the most expensive menu option awarded to buyers with budgets $\geq b_i$ will cost $\leq b_i$.

So there are only two things we need to watch out for when carefully choosing seeds: that we choose one in a way so that $\pi_j(\bar{v}) \leq 1$ always, and also that that $p_i(\bar{v}) = b_i$ if $\gamma_i > 0$ and $\Phi_i^\Lambda(\bar{v}) > 0$.

4.2.3 Generate-Menu: Setting Appropriate Seeds. This is perhaps the most technical part of the proof, and will require the concept of "linked" budgets to get all the way to the main result. In Appendix C, we'll build up to the main result stated below in a couple steps.

- Observe that if $\Phi_1^\Lambda(\bar{v}) > 0$, then we are basically just the FedEx Problem, except that maybe budget 1 binds. So we need to be a little careful how we set the menu for budget 1, but otherwise the approach looks very similar to that of Fiat et al. [2016] (Proposition C.4).
- Observe that if we seed Generate-Menu with the trivial menu for budget 1, this will propagate through all j such that $\Phi_j^\Lambda(\bar{v}) = 0$. Generate-Menu will introduce a single non-trivial menu option at the first i with $\Phi_i^\Lambda(\bar{v}) > 0$ with allocation probability 1, but possibly price $> b_i$. So all feasibility concerns are addressed, but budget constraints might be violated, and it can only be violated by charging too much (Proposition C.5).
- Observe that if we instead seeded Generate-Menu(i) with a menu for $i-1$ that offered the option $(1, b_{i-1})$, then the menu output for budget i would *definitely* charge a price $\leq b_{i-1} < b_i$. Since Generate-Menu is a continuous function, the intermediate value theorem says that there is some $q \in [0,1]$ such that if a buyer with bidder $i-1$ has the option to purchase $(q, q \cdot b_{i-1})$, then Generate-Menu(i) will output a menu with the option $(1, b_i)$ as desired.
- If Λ happens to be budget feasible everywhere, then the above can be achieved by carefully tuning the menu for budget 1 (to contain the option $(q, q \cdot b_1)$ for some $q \in [0,1]$). But

if Λ isn't budget feasible at 1 (or at any $j > 1$), this approach fails because it proposes a non-trivial menu for a budget where Λ isn't budget feasible (Proposition C.6).

- If Λ isn't budget feasible everywhere, a more complicated approach can be used to find an appropriate "seed budget" to be the first one that is seeded with a non-trivial budget. This is somewhat complex because if budgets i and $i-1$ are linked, we can't just set the menu for budget i arbitrarily once we've decided that the menu for budget $i-1$ is trivial (exactly the definition of being linked). This is discussed (without proofs) below.

Definition 4.14 (Seed budget). Define the *seed budget*, s^*, (which will be the first non-trivial menu used to seed Generate-Menu) in the following way:

- First, set s^* to be the minimum i such that $\Phi_i^{\Lambda}(\bar{v}) > 0$.
- Iterate the following until termination: If $\gamma_{s^*} > 0$ *and* Λ links s^* to s^*-1, then update $s^* := s^* - 1$.

LEMMA 4.15. *Λ is budget-feasible for all $i \geq s^*$.*

PROPOSITION 4.16. *Let s^* be the seed budget for Λ. Let also $[v^+, \bar{v}]$ denote the ironed interval within budget s^* containing $\bar{v}$, and $[v^-, x]$ the ironed interval containing b_{s^*}. There exists a $q \in [0,1]$ such that the following mechanism solves LP2(Λ) and satisfies complementary slackness.*

- $(\pi_i(v), p_i(v)) = (0,0)$ *for all v and $i < s^*$.*
- $(\pi_{s^*}(v), p_{s^*}(v)) = (0,0)$ *for all $v < v^-$.*
- $(\pi_{s^*}(v), p_{s^*}(v)) = (q, b_{s^*}q)$ *for all $v \geq v^+$, and* $(\pi_{s^*}(v), p_{s^*}(v)) = (q\frac{v^+-b_{s^*}}{v^+-v^-}, v^- q \cdot \frac{v^+-b_{s^*}}{v^+-v^-})$ *for all $v \in [v^-, v^+)$.*
- $(\vec{\pi}_i, \vec{p}_i) =$ *Generate-Menu*$(i, \vec{\pi}_{i-1}, \vec{p}_{i-1})$ *for all $i > s^*$.*

THEOREM 4.17 (PROCESSES FORMAT OF PROPOSITION 4.16). *The optimal mechanism takes the following form, where ℓ_i denote the number of non-trivial menu options for types with budget b_i.*

- *There exists an i such that $b_j \cdot \pi_j(\bar{v}) = p_j(\bar{v})$ for all $j \leq i$, and there exists an a such that $\pi_j(\bar{v}) = 1$ for all $j \geq a$. There exists a c such that $\pi_j(v) = p_j(v) = 0$ for all v and $j < c$.*
- *Above, $c \leq a$ and $a \in \{i, i+1\}$. So there is at most one j with $b_j \cdot \pi_j(\bar{v}) = p_j(\bar{v})$ and $\pi_j(\bar{v}) = 1$.*
- *$\ell_c \leq 2$, and $\ell_j \leq 2\ell_{j-1} + 1$ for all j.*
- *The menu for types with budget b_j is obtained by Generate-Menu with input equal to the menu for types with budget b_{j-1} for all $j > c$.*

REFERENCES

Zoë Abrams. 2006. Revenue Maximization when Bidders Have Budgets. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm (SODA '06)*. 1074–1082.

Jean-Pierre Benoit and Vijay Krishna. 2001. Multiple-object auctions with budget constrained bidders. *The Review of Economic Studies* 68, 1 (2001), 155–179.

Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. 2010. Budget constrained auctions with heterogeneous items. In *STOC*. 379–388.

Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. 2005. Multi-unit auctions with budget-constrained bidders. In *ACM Conference on Electronic Commerce*. 44–51.

Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. 2016. A duality based unified approach to Bayesian mechanism design. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 926–939.

Shuchi Chawla, David L. Malec, and Azarakhsh Malekian. 2011. Bayesian mechanism design for budget-constrained agents. In *ACM Conference on Electronic Commerce*. 253–262.

Yeon-Koo Che and Ian Gale. 1998. Standard Auctions with Financially Constrained Bidders. *The Review of Economic Studies* 65, 1 (1998), 1–21.

Yeon-Koo Che and Ian Gale. 2000. The Optimal Mechanism for Selling to a Budget-Constrained Buyer. *Journal of Economic theory* 92, 2 (2000), 198–233.

Yeon-Koo Che, Ian Gale, and Jinwoo Kim. 2012. Assigning Resources to Budget-Constrained Agents. *Review of Economic Studies* 80, 1 (2012), 73–107.

Peter C Cramton. 1995. Money out of thin air: The nationwide narrowband PCS auction. *Journal of Economics & Management Strategy* 4, 2 (1995), 267–343.

Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2013. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 269–286.

Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2015a. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 449–450.

Constantinos Daskalakis, Nikhil Devanur, and S Matthew Weinberg. 2015b. Revenue maximization and ex-post budget constraints. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 433–447.

Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. 2013. Prior-free Auctions for Budgeted Agents. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC '13)*. 287–304.

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

Shahar Dobzinski, Ron Lavi, and Noam Nisan. 2008. Multi-unit Auctions with Budget Limits. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*. IEEE Computer Society, Washington, DC, USA, 260–269. DOI:http://dx.doi.org/10.1109/FOCS.2008.39

Amos Fiat, Kira Goldner, Anna R Karlin, and Elias Koutsoupias. 2016. The FedEx problem. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 21–22.

Arieh Gavious, Benny Moldovanu, and Aner Sela. 2002. Bid costs and endogenous bid caps. *RAND Journal of Economics* (2002), 709–722.

Yiannis Giannakopoulos and Elias Koutsoupias. 2014. Duality and optimality of auctions for uniform distributions. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 259–276.

Yiannis Giannakopoulos and Elias Koutsoupias. 2015. Selling two goods optimally. In *International Colloquium on Automata, Languages, and Programming*. Springer, 650–662.

Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. 2012. Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the 44th symposium on Theory of Computing (STOC '12)*. Association for Computing Machinery, 107–122.

Nima Haghpanah and Jason Hartline. 2015. Reverse Mechanism Design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*.

Zhiyi Huang, Nikhil R Devanur, and David Malec. 2012. Sequential auctions of identical items with budget-constrained bidders. *arXiv preprint arXiv:1209.1698* (2012).

Cinar Kilcioglu and Costis Maglaras. 2015a. AWS EC2 Spot vs On-Demand Analysis. (2015). http://www.columbia.edu/~ck2560/awsplot.html

Cinar Kilcioglu and Costis Maglaras. 2015b. Revenue Maximization for Cloud Computing Services. *SIGMETRICS Performance Evaluation Review* 43, 3 (2015), 76.

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

Alexey Malakhov and Rakesh V Vohra. 2004. *Single and multi-dimensional optimal auctions: a network approach*. Technical Report. Discussion paper//Center for Mathematical Studies in Economics and Management Science.

Alexey Malakhov and Rakesh V Vohra. 2009. An optimal auction for capacity constrained bidders: a network perspective. *Economic Theory* 39, 1 (2009), 113–128.

Alejandro M Manelli and Daniel R Vincent. 2006. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory* 127, 1 (2006), 1–35.

Eric S Maskin. 2000. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review* 44, 4 (2000), 667–681.

Roger Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (February 1981), 58–73.

Mallesh M. Pai and Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150, C (2014), 383–425.

Carolyn Pitchik. 2009. Budget-constrained sequential auctions with incomplete information. *Games and Economic behavior* 66, 2 (2009), 928–949.

Carolyn Pitchik and Andrew Schotter. 1988. Perfect equilibria in budget-constrained sequential auctions: An experimental study. *The Rand Journal of Economics* (1988), 363–388.

Anand Venkateswaran. 2013. IPL auction Day 1 - As it happened - The Hindu. (2013). http://www.thehindu.com/sport/cricket/ipl-auction-day-1-as-it-happened/article1072914.ece Accessed: 2013-10-30.

A OTHER RELATED WORK

Laffont and Robert [1996] consider the budget constrained auction design problem with *multiple* bidders. They characterize the optimal auction under the assumption that all the bidders have the same budget which is known to the seller, and that their values are drawn i.i.d. from a *regular* distribution. While regularity implies "no pooling" in the absence of budgets [Myerson, 1981], Laffont and Robert [1996] show that there maybe pooling at the top when the budget constraint is binding; however, regularity still ensures that there is no additional pooling at the bottom. Pai and Vohra [2014] consider the case where the budgets are different and private, under the assumption that the distribution of budgets and values are *independent* of each other, and that the marginal distribution over values satisfies the *monotone hazard rate* (MHR) and the *decreasing density* conditions.[13] Their techniques could also be applied to the *social welfare* maximizing auction, under similar assumptions. This was preceded by Maskin [2000] who characterized the social welfare maximizing auction, with 3 bidders with a common budget constraint. There are several open questions here, about how to obtain improvements that are analogous to our result over Che and Gale [2000]: characterize the optimal multi bidder auction with no assumption about the value distributions.

The computer science community has addressed the difficulty of characterizing optimal mechanisms by designing apprximations. An α-approximate mechanism is one whose revenue is always guaranteed to be within an α multiplicative factor of the revenue of the optimal mechanism. Chawla et al. [2011] design constant factor approximations for several cases: single dimensional buyer types with *public* budgets and a seller facing *downward closed* service constraints (this includes single item auctions), single dimensional buyer types drawn from an *MHR* distribution with *private* budgets and a seller facing *matroid* service constraints, and multidimensional *unit-demand* buyer types with *public* budgets and a seller facing *matroid* service constraints. They also have similar guarantees for welfare maximization. Bhattacharya et al. [2010] give a constant factor approximation with multiple buyers and multiple items, when buyers valuations are *additive*, and are drawn from *MHR* distributions, under the *public* budgets assumption. Daskalakis et al. [2015b] improved this to a 3-approximation, for *all* distributions, and *private* budgets. In the much harder *prior-free* (aka *worst case*) setting, there have been several approximation results on both revenue and welfare maximization [Abrams, 2006, Borgs et al., 2005, Devanur et al., 2013], as well as on the design of *Pareto optimal* auctions [Dobzinski et al., 2008, Goel et al., 2012].

Others have studied how standard auctions are affected by the presence of budget constraints. Che and Gale [1998] rank different auction formats and show that all-pay auctions do best followed by first price auctions and then second price auctions. In general, all-pay auctions are a good idea in the presence of budget constraints since the payments for any one bidder are the lowest and hence the budget constraints are the least binding. Benoit and Krishna [2001], Huang et al. [2012], Pitchik [2009], Pitchik and Schotter [1988] all study *sequential* auctions with budget constraints and characterize sub game perfect Bayes Nash equilibria. This has been difficult to do beyond very simple settings. Finally, Che et al. [2012] study welfare maximizing mechanisms when the buyers are allowed to trade goods after the auction. They show that random assignments/subsidies and allowing ex-post trade leads to higher social welfare.

B OMITTED PROOFS FROM SECTION 3

[13] A distribution with a cdf F and density f satisfies the monotone hazard rate condition if $\frac{1-F(v)}{f(v)}$ is non increasing. It satisfies the decreasing density condition if $f(v)$ is decreasing.

Proof of Corollary 3.2. Consider any Λ such that $\Phi^\Lambda(v) < \Phi^\Lambda(v-1)$, and consider x such that $\Phi^\Lambda(v) + x/f(v) = \Phi^\Lambda(v-1) - x/f(v-1)$ (such an x clearly exists as $\Phi^\Lambda(v) < \Phi^\Lambda(v-1)$). By Lemma 3.1, adding a cycle of weight x between v and $v-1$ results in $\Phi^{\Lambda'}(v) = \Phi^{\Lambda'}(v-1)$, while maintaining $f(v)\cdot\Phi^{\Lambda'}(v) + f(v-1)\cdot\Phi^{\Lambda'}(v-1) = f(v)\cdot\Phi^\Lambda(v) + f(v-1)\cdot\Phi^\Lambda(v-1)$. Similarly, adding a cycle of weight y between v and $v-1$, where $y = (\Phi^\Lambda(v-1)f(v-1) - \Phi^\Lambda(v)f(v))/2$ results in $\Phi^{\Lambda'}(v)f(v) = \Phi^{\Lambda'}(v-1)f(v-1)$, while maintaining $f(v)\cdot\Phi^{\Lambda'}(v) + f(v-1)\cdot\Phi^{\Lambda'}(v-1) = f(v)\cdot\Phi^\Lambda(v) + f(v-1)\cdot\Phi^\Lambda(v-1)$. Let's now compare the optimal solution to LP2(Λ) to that of LP2(Λ') in either case.

By Observation 1, $\mathcal{L}_{\max}(\Lambda) - \mathcal{L}_{\max}(\Lambda')$ is exactly $\max\{0, f(v)\Phi^\Lambda(v)\} + \max\{0, f(v-1)\phi^\Lambda(v-1)\} - \max\{0, f(v)\Phi^\Lambda(v) + f(v-1)\Phi^\Lambda(v-1)\}$, which is clearly non-negative. Here, we are basically just observing the sum of maxes is at least as large as the max of sums. Therefore, $\mathcal{L}_{\max}(\Lambda) \geq \mathcal{L}_{\max}(\Lambda')$, and if Λ solves Min-Max2, then so does Λ'.

The above argument can be iterated to eventually yield a Λ where $\Phi^\Lambda(\cdot)$ is monotone non-decreasing.

To see the second claim, observe that whenever $\lambda(v-1, v) = x > 0$, there is a cycle of weight x between $v-1$ and v. So the weight of this cycle can be decreased by any $y < x$, similarly preserving $\Phi^\Lambda(v)\cdot f(v) + \Phi^\Lambda(v-1)\cdot f(v-1)$, except it will *increase* $\Phi^\Lambda(v-1)$ and *decrease* $\Phi^\Lambda(v)$. If $\Phi^\Lambda(v) > \Phi^\Lambda(v-1)$, then the exact same reasoning above shows that decreasing the weight on the cycle can only decrease $\mathcal{L}_{\max}(\Lambda)$. The argument can again be iterated to eventually yield a Λ where $\Phi^\Lambda(\cdot)$ that is both monotone non-decreasing, and satisfies $\lambda(v-1, v) > 0 \Rightarrow \Phi^\Lambda(v) = \Phi^\Lambda(v-1)$, and an identical argument completes the proof for $f(\cdot)\cdot\Phi^\Lambda(\cdot)$ as well. □

Proof of Theorem 3.7. Consider any dual of the form promised by Corollary 3.6. First consider the case that $\gamma = 0$. In this case, let x denote the minimum value with $\Phi^\Lambda(x) \geq 0$, which is guaranteed to be at most b. Now, consider the mechanism that simply sets price x on the item (formally, $(\pi(v), p(v)) = (1, x)$ for $v \geq x$ and $(\pi(v), p(v)) = (0, 0)$ for $v < x$). This is certainly IC, and certainly solves LP2(Λ). Moreover, all IC constraints between any two types $\geq x$ are tight (they get the same allocation/price), as are any IC constraints between any two types $< x$. The IC constraint between x and $x-1$ is also tight, and the IC constraint between $x-1$ and x is not. However, Corollary 3.6 promises that $\lambda(x-1, x) = 0$, as we necessarily have $\Phi^\Lambda(x) > \Phi^\Lambda(x-1)$, and therefore complementary slackness is satisfied. So in this case, there is a mechanism that solves LP2(Λ) and satisfies complementary slackness, and it is therefore optimal.

Next, consider the case that $\gamma > 0$ and $\Phi^\Lambda(b-1) \leq 0$. Let y denote the minimum v such that $\Phi^\Lambda(v) \geq 0$ ($y \leq b$ by Corollary 3.6). Then by Corollary 3.6, $\lambda(y, y-1) = 0$. Let also x denote the maximum v such that $\lambda(v, v-1) = 0$. By Corollary 3.6, $x \geq b$. As $y \leq b$, and $x \geq b$, there exists a $q \in [0, 1]$ such that $b = qy + (1-q)x$. Consider the mechanism that offers the two options $(1, b)$ and (q, qy) (formally, $(\pi(v), p(v)) = (1, b)$ for $v \geq x$, $(\pi(v), p(v)) = (q, qy)$ for $v \in [y, x)$, $(\pi(v), p(v)) = (0, 0)$ for $v < y$). This mechanism is clearly IC, solves LP2(Λ), and has $p(\bar{v}) = b$. The only two IC constraints that aren't tight are those between $x-1$ and x, and $y-1$ and y, but we are guaranteed to have $\lambda(x-1, x) = \lambda(y-1, y) = 0$ by Corollary 3.6, and complementary slackness holds.

So in either case, there exists a mechanism of the desired form that solves LP2(Λ) and satisfies complementary slackness. By Strong Duality, such a mechanism solves LP1 (and is therefore optimal). □

Proof of Lemma 3.10. Let's first prove the lemma for $w = v + 1$, and then for all $w \geq v + 1$ at the end.

$$\begin{aligned}(v+1)f(v+1) - 1 + F(v+1) - vf(v) + 1 - F(v) \\ = (v+1)f(v+1) - vf(v) + f(v+1) \\ = v(f(v+1) - f(v)) + 2f(v+1).\end{aligned}$$

So maybe $f(v+1) \geq f(v)$, in which case $v(f(v+1) - f(v)) + 2f(v+1)$ is clearly positive, and we have $vf(v) - (1 - F(v)) \leq (v+1)f(v+1) - (1 - F(v+1))$ as desired. In case $f(v+1) < f(v)$, consider the following two inequalities which we know to be true:

$$\begin{aligned}0 &\geq \varphi(v)\\ \varphi(v+1) &\geq \varphi(v).\end{aligned}$$

The second equation is due to F being regular. Now consider multiplying the first equation by $f(v) - f(v+1)$, which is positive by assumption, and the second by $f(v+1)$. Then we get:

$$f(v+1)\varphi(v+1) \geq f(v)\varphi(v),$$

as desired. To see that this holds for any $w \geq v + 1$, observe that we can chain the above from v up thru any w such that $\varphi(w) \leq 0$. For any w such that $\varphi(w) > 0$, we clearly have $f(w)\varphi(w) > 0 \geq f(v)\varphi(v)$. □

Corollary B.1. *Let F be regular, and let r denote the maximum $v \leq b$ with $\Phi^\Lambda(v-1) \leq 0$. Then there exists an optimal solution to Min-Max2, Λ, such that:*

- *Λ is properly ironed.*
- *$\Phi^\Lambda(b) \geq 0$ (so r exists).*
- *$\lambda(r-1, r) = 0$.*
- *Either $\gamma = 0$ or $r = b$.*

Proof. Again, consider a fixed γ and the following $\vec{\lambda}$: $\lambda(v, v-1) = 1 - \gamma - F(v-1)$. Recall that (Λ) is flow-conserving, and results in $\Phi^\Lambda(\bar{v}) = \bar{v} - \frac{(\bar{v}-b)\gamma}{f(\bar{v})}$, $\Phi^\Lambda(v) = v - \frac{1-F(v)-\gamma}{f(v)}$ for $v < \bar{v}$. Observe that as F is regular, we have (by Lemma 3.10):

$$vf(v) - (1 - F(v)) \leq wf(w) - (1 - F(w)) \Rightarrow f(v)\Phi^\Lambda(v) \geq f(w)\Phi^\Lambda(w), \forall v < w, \varphi(v) \leq 0.$$

So for all v such that $\varphi(v) \leq 0$, $f(\cdot)\Phi^\Lambda(\cdot)$ is already monotone non-decreasing. It is easy to see the procedure described in Corollary 3.2 will result in potentially multiple ironed intervals of the following form. There will be a single ironed interval $[v^*, \bar{v}]$ at the top, and potentially other ironed intervals $[v, w]$ with $\varphi(v) > 0$ (and therefore also $\Phi^\Lambda(v) > 0$). But it is not possible to have an ironed interval $[v, w]$ with $\varphi(v) \leq 0$ and $w < \bar{v}$ by the work above, because we would have already had $\Phi^\Lambda(v) \leq \Phi^\Lambda(v')$ for all $v' \in [v, w]$. As in Proposition 3.9, we must have $v^* \geq b$, as otherwise complementary slackness cannot be satisfied - $v^* \neq \bar{v}$ implies $\gamma > 0$, and $v^* < b$ implies $p(\bar{v}) < b$ in any IC mechanism. Now, let x be the minimum v such that $\varphi(v) \geq 0$. If $b \leq x$, then we necessarily have $\lambda(r-1, r) = 0$, as $r - 1$ cannot be in an ironed interval (it cannot be in the interval $[v^*, \bar{v}]$ as $r \leq b \leq v^*$, and it cannot be in any other ironed interval due to the work above as $\varphi(r-1) \leq 0$). If $b > x$, then we also have $r \leq x$: x is clearly not in the ironed interval $[v^*, \bar{v}]$ as $x < b \leq v^*$. Moreover, for all $v \in [x, v^*)$, we have (before ironing) $\Phi^\Lambda(v) \geq \varphi(v) \geq 0$. So every ironed interval (except possibly $[v^*, \bar{v}]$) contains only values with $\Phi^\Lambda(v) \geq 0$ before ironing, and therefore $\Phi^\Lambda(v) \geq 0$ after ironing as well. Together, this implies that $\Phi^\Lambda(x) \geq 0$, and therefore $r \leq x$. As we have $\lambda(v-1, v) = 0$ for all $v \leq x$, we conclude that $\lambda(r-1, r) = 0$ as well. In both cases, similarly to Proposition 3.9, if $\gamma > 0$, we must have $r = b$ due to complementary slackness, as no IC mechanism can simultaneously optimize LP2(Λ) and have $p(\bar{v}) = b\pi(\bar{v})$ (this is because $r \neq b \Rightarrow \Phi^{\Lambda}(b-1) > 0$). □

Note that Corollary B.1 completes the proof of Proposition 3.9.

Proof of Theorem 3.12. If $x \leq b$, then the optimal mechanism clearly sets reserve x (as this is optimal even without the budget constraint). If $x > b$, then the optimal dual necessarily has $\gamma > 0$ (the budget constraint is binding), in which case Proposition 3.11 tells us that the optimal mechanism sets price b. □

C OMITTED PROOFS FROM SECTION 4

Proof of Observation 3. Simply check through the definition of $\Phi_i^{\Lambda}(\cdot)$. Decreasing $\lambda_i(w+1, w)$ by c increases $\Phi_i^{\Lambda}(w)$ by $c/f_i(w)$. Any change in $\alpha_i(v)$ doesn't directly affect any virtual values (the effect is indirect due to how $\vec{\lambda}$ necessarily changes to conserve flow). Increasing $\lambda_{i-1}(w+1, w)$ by c decreases $\Phi_{i-1}^{\Lambda}(w)$ by $c/f_i(w)$. □

Proof of Observation 4. Again the proof just requires checking the definition of $\Phi_i^{\Lambda}(\cdot)$. Increasing $\lambda_i(v, v-1)$ by c decreases $\Phi_i^{\Lambda}(v-1)$ by $c/f_i(v-1)$. The change in $\alpha_i(v)$ and $\alpha_i(v-1)$ doesn't directly affect any virtual values. Increasing $\lambda_i(v-1, v)$ by c inreases $\Phi_{i-1}^{\Lambda}(v)$ by $c/f_{i-1}(v)$. □

Proof of Observation 5. Again, just check the definition of $\Phi_i^{\Lambda}(\cdot)$. (i, v) gets $c/2$ less flow in from both $(i, v+1)$ and $(i, v-1)$, so the effects cancel. $(i-1, v)$ gets $c/2$ more flow in from both $(i-1, v+1)$ and $(i-1, v-1)$, so again the effects cancel. □

Proof of Observation 6. First, observe that $b_i < \bar{v}$ always (otherwise the budget constraint is already implied by IR). The remainder of the proof again just chases through the definition of $\Phi_j^{\Lambda}(\bar{v})$. Observe that decreasing γ_i by $\frac{c}{\bar{v}-b_i}$ increases $\Phi_i^{\Lambda}(\bar{v})$ by c, while decreasing $\lambda_i(\bar{v}-1, \bar{v})$ by c decreases $\Phi_i^{\Lambda}(\bar{v})$ by c, so $\Phi_i^{\Lambda}(\bar{v})$ is unchanged. Similarly, the change to $\Phi_{i-1}^{\Lambda}(\bar{v})$ is $c - \frac{c(\bar{v}-b_{i-1})}{\bar{v}-b_i}$. Observe that as $b_{i-1} < b_i$, we do in fact guarantee that $\Phi_{i-1}^{\Lambda}(\bar{v})$ decreases. □

Proof of Proposition 4.6. The intuition is similar to the proof of Proposition 4.5: this operation decreases $f_i(v-1)\cdot\Phi_i^{\Lambda}(v-1)$, which is definitely helping to lower $\mathcal{L}_{\max}$ because $f_i(v-1)\cdot\Phi_i^{\Lambda}(v-1) > 0$. The possible catch is that we are increasing $f_{i-1}(v)\cdot\Phi_{i-1}^{\Lambda}(v)$, which could hurt. But we can show that the hurt can never outweigh the help. A little more formally, Observation 1 shows that $\mathcal{L}_{\max}(\Lambda) - \mathcal{L}_{\max}(\Lambda') = \max\{0, f_i(v-1)\cdot\Phi_i^{\Lambda}(v-1)\} - \max\{0, f_i(v-1)\cdot\Phi_i^{\Lambda}(v-1)-\epsilon\} + \max\{0, f_{i-1}(v)\cdot\Phi_{i-1}^{\Lambda}(v)\} - \max\{0, f_{i-1}(v)\cdot\Phi_{i-1}^{\Lambda}(v)+\epsilon\}$. As $f_i(v-1)\cdot\Phi_i^{\Lambda}(v-1) = \epsilon$, the difference between the first two terms is exactly ϵ. The difference between the last two terms is clearly at least $-\epsilon$, so the total difference is non-negative, and $\mathcal{L}_{\max}(\Lambda) \leq \mathcal{L}_{\max}(\Lambda)$, and we have only improved. □

Proof of Corollary 4.9. Suppose that $\Phi_i^\Lambda(\bar{v}) = 0$. Simply observe that $(\vec{0}, \vec{0})$ satisfies all the necessary IC constraints as well as the budget constraint. Therefore, by Lemma 4.8, it is the unique possibility.

Now suppose that $\Phi_i^\Lambda(\bar{v}) > 0$. Again, simply observe that the proposed menu satisfies all the necessary IC constraints as well as the budget constraint (due to budget-feasibility). Therefore, by Lemma 4.8, it is the unique possibility. □

Proof of Theorem 4.10. The first two bullets follow directly from Corollary 4.1. The third bullet follows from Proposition 4.5: once Λ is properly ironed, if there is some $v, i > 0$ with $\Phi_i^\Lambda(v-1) < 0$, we can always boost $\alpha_i(v)$ to make it 0 instead. The fourth bullet follows from Proposition 4.6, and Observations 5 and 6. Assume first that there is some $\alpha_i(v) > 0$ with $\Phi_i^\Lambda(v-1) > 0$ as well. Then Proposition 4.6 says that we can reroute $\alpha_i(v)$ and only improve Λ. Now assume that $\alpha_i(v) > 0$, $\Phi_i^\Lambda(v-1) = 0$, and $\lambda_i(v-1, v) > 0$. Observe that becasuse Λ is properly ironed, that $\Phi_i^\Lambda(v) = 0$ as well. Next, observe that in order to possibly have $\Phi_i^\Lambda(v) = 0$, we must have $\lambda_i(v+1, v) > 0$ (or $\gamma_i > 0$ if $v = \bar{v}$), so $\alpha_i(v)$ is a candidate for splitting. The fifth bullet follows from the following observation: If $\Phi_{i-1}^\Lambda(\bar{v}) > 0$, then any IC and budget-respecting solution to LP2(Λ) must have $\pi_{i-1}(\bar{v}) = 1$ and $p_{i-1}(\bar{v}) \leq b_{i-1}$, and therefore also have $p_i(\bar{v}) \leq b_{i-1}$ (as otherwise, $(i, \bar{v})$ would clearly rather lie and report $(i-1, \bar{v})$). So no IC, ex-post budget-respecting solution to LP2(Λ) can possibly satisfy complementary slackness, and by Strong Duality (Λ) cannot possibly solve Min-Max2. To summarize, bullet points two and five necessarily hold in any solution to Min-Max2. Whenever bullet points one, three, or four don't hold, there exist elementary operations (proper ironing, boosting, re-reouting, splitting) that can only improve the dual. So these operations can be repeatedly performed until all bullet points hold. □

Proposition C.1 ([Fiat et al., 2016]). *Let $(\vec{\pi}_i, \vec{p}_i)$ satisfy all leftwards/rightwards IC constraints, and $\pi_i(v) \in [0,1], p_i(v) \geq 0$ for all v, and $(\vec{\pi}_{i+1}, p_{i+1}) =$ Generate-Menu$(i+1, \vec{\pi}_i, \vec{p}_i)$. Then $\pi_{i+1}(v) \in [0,1]$ and $p_{i+1}(v) \geq 0$ for all $v < v_i^+$.*

Proof. Consider any ironed interval $[w, v]$. Note that $(v+1)\pi_i(v+1) - w\pi_i(w) - p_i(v+1) + p_i(w)$ is just the utility that $(i, v+1)$ receives minus the utility that (i, w) receives. As $v \geq w$, this is clearly positive. Therefore, $\pi_i(v) = \frac{(v+1)\pi_i(v+1) - w\pi_i(w) - p_i(v+1) + p_i(w)}{v+1-w} \geq 0$. Also, the difference in utility is maximized when $\pi_i(w) = 1$ (over all menus that satisfy leftwards/rightwards IC), in which case $\pi_i(v) = \frac{(v+1)\pi_i(v+1) - w\pi_i(w) - p_i(v+1) + p_i(w)}{v+1-w} = 1$. So no matter what, $\pi_{i+1}(v) \in [0,1]$. Also,

$$p_i(v) = w\frac{(v+1)\pi_i(v+1) - w\pi_i(w) - p_i(v+1) + p_i(w)}{v+1-w} - w\pi_i(w) + p_i(w)$$

$$= \frac{w(v+1)\pi_i(v+1) - w^2\pi_i(w) - wp_i(v+1) + wp_i(w) - w(v+1)\pi_i(w) + w^2\pi_i(w) + (v+1)p_i(w) - wp_i(w)}{v+1-w}$$

$$= \frac{w(v+1)(\pi_i(v+1) - \pi_i(w)) - wp_i(v+1) + (v+1)p_i(w)}{v+1-w}$$

$$\geq \frac{w(v+1)(\pi_i(v+1) - \pi_i(w)) - wp_i(w) - w(v+1)(\pi_i(v+1) - \pi_i(w)) + (v+1)p_i(w)}{v+1-w}$$

$$= \frac{(v+1-w)p_i(w)}{v+1-w} \geq 0.$$

Where the penultimate line follows because $p_i(v+1) \leq p_i(w) + (v+1)(\pi_i(v+1) - \pi_i(w))$. □

Corollary C.2 ([Fiat et al., 2016]). *Let $(\vec{\pi}_i, \vec{p}_i)$ satisfy all leftwards/rightwards IC constraints, $\pi_i(v) \in [0,1], p_i(v) \geq 0$ for all v, and $\Phi^{\Lambda}_{i+1}(\bar{v}) > 0$. Then $\pi_{i+1}(v) \in [0,1]$ and $p_{i+1}(v) \geq 0$ for all v when $(\vec{\pi}_{i+1}, p_{i+1}) =$ Generate-Menu$(i+1, \vec{\pi}_i, \vec{p}_i)$.*

Proof. As $\pi_i(v_i^+) \in [0,1]$, this immediately follows from the penultimate line in the description of Generate-Menu combined with Proposition C.1. □

Corollary C.3. *Let $(\vec{\pi}_i, \vec{p}_i)$ satisfy all leftwards/rightwards IC constraints, $\pi_i(v) \in [0,1], p_i(v) \geq 0$ for all v, and $\Phi^{\Lambda}_{i+1}(\bar{v}) = 0$. Then if Λ is budget-feasible at i, $\pi_{i+1}(v), p_{i+1}(v) \geq 0$ for all v when $(\vec{\pi}_{i+1}, p_{i+1}) =$ Generate-Menu$(i+1, \vec{\pi}_i, \vec{p}_i)$.*

Moreover, if $v_i^+ \cdot \pi_i(v_i^+) - p_i(v_i^+) \leq v_i^+ - b_{i+1}$, $\pi_{i+1}(v) \in [0,1]$ and $p_{i+1}(v) \geq 0$ for all v when $(\vec{\pi}_{i+1}, p_{i+1}) =$ Generate-Menu$(i+1, \vec{\pi}_i, \vec{p}_i)$.

Proof. Again, immediately follows from Proposition C.1 combined with the final line in the description of Generate-Menu. When $v_i^+ \geq b_i$, the allocation and price are clearly non-negative (although perhaps infinite). For the "moreover" portion: intuitively, $v_i^+ \cdot \pi_i(v_i^+) - p_i(v_i^+) \leq v_i^+ - b_{i+1}$ ensures that $(i+1, v_i^+)$ at least prefers to receive the item with probability 1 and pay b_{i+1} than receive $(\pi_i(v_i^+), p_i(v_i^+))$. If this condition is satisfied, there is certainly some $\pi_{i+1}(v_i^+) \leq 1$ with which we can award the item to $(i+1, v_i^+)$, charge her $b_{i+1} \cdot \pi_{i+1}(v_i^+)$, and have the desired IC constraints be tight. If this condition fails to hold, then we might need to charge a negative price or allocate at a "probability" > 1. As a sanity check, observe that there is absolutely no way for the hypotheses of the corollary statement to be satisfied if Λ isn't budget-feasible at i. We know that $\Phi^{\Lambda}_i(v) = 0$ for all v, so the first and second conditions of budget-feasibility are certainly satisfied. So if there's an issue, it's with the third condition. If the third condition is violated, it essentially means that there's one very long ironed interval going from $(i, \bar{v})$ all the way below b_i, and therefore $v_i^+ < b_i$, and the RHS in the corollary's hypothesis is negative. □

Proof of Observation 7. Just observe that by budget-feasibility at i, v_i^+ will be $\geq b_i$. It is then easy to check that $(\vec{0}, \vec{0})$ satisfies all the hypotheses of Corollary C.3. □

Proof of Proposition 4.12. Observe that all leftwards/rightwards IC constraints within budget b_1 are strongly satisfied by assumption. Now let's proceed by induction and assume that all leftwards/rightwards IC constraints are strongly satisfied within budget b_i and consider all IC constraints involved with budget b_{i+1}. First, observe that all rightwards IC constraints are certainly strongly satisfied, as everyone in the same ironed interval receives the same allocation/price. Next, observe that for duals of the form guaranteed by Theorem 4.10 that $\alpha_i(v) > 0 \Rightarrow \lambda_i(v-1, v) = 0$. In other words, if $\alpha_i(v) > 0$, then v is at the bottom of an ironed interval. Checking through the algebra/English description of the Generate-Menu procedure, we see that $(\pi_i(\cdot), p_i(\cdot))$ are set exactly so that all $(i+1, v)$ at the bottom of an ironed interval are indifferent between what they receive and what (i, v) receives, *and* indifferent between what (i, v) receives and what $(i+1, v-1)$ receives. Together, this implies that $(i+1, v)$ is also indifferent between what she receives and what $(i+1, v-1)$ receives. Therefore, all leftwards IC constraints as well as all downwards IC constraints are strongly satisfied as well. □

Proof of Corollary 4.13. Observe that whenever $\Phi_i^\Lambda(v) > 0$, we have hard-coded $\pi_i(v) = 1$ into Generate-Menu. Additionally, for all $i > 1$, $\Phi_i^\Lambda(v) \geq 0$, and any value is acceptable when $\Phi_i^\Lambda(v) = 0$ (and for all $i > 1$ such that $\Phi_i^\Lambda(\bar{v}) = 0$ for all v). However, we still can't guarantee that the resulting menu is feasible without additional hypotheses (e.g. Corollary C.3), so we just leave it as a hypothesis in the corollary statement. □

Proposition C.4. *Let $\Phi_1^\Lambda(\bar{v}) > 0$. If $\gamma_1 = 0$, let $v^+ = v^- = \min\{v|\Phi_1^\Lambda(v) \geq 0\}$. Else, let $[v^-, v^+ - 1]$ denote the ironed interval containing b_1 (could be $[b_1, b_1]$). Then the following mechanism solves LP2(Λ) and strongly satisfies complementary slackness.*

- $(\pi_1(v), p_1(v)) = (0, 0)$ *for all* $v < v^-$.
- *If* $\gamma_1 = 0$, $(\pi_1(v), p_1(v)) = (1, v^-)$ *for all* $v \geq v^-$.
- *Else,* $(\pi_1(v), p_1(v)) = (1, b_1)$ *for all* $v \geq v^+$, *and* $(\pi_1(v), p_1(v)) = (\frac{v^+-b_1}{v^+-v^-}, v^- \cdot \frac{v^+-b_1}{v^+-v^-})$ *for all* $v \in [v^-, v^+)$.
- $(\vec{\pi}_i, \vec{p}_i) =$ *Generate-Menu*$(i, \vec{\pi}_{i-1}, \vec{p}_{i-1})$ *for all* $i > 1$.

Proof of Proposition C.4. Quickly note that by Theorem 4.10, Λ is budget-feasible at 1, so the above algebra always results in a feasible menu. Also, note that as there are no j such that $\Phi_j^\Lambda(\bar{v}) = 0$, there's no possibility that such j have $\pi_j(\bar{v}) > 1$. So the first concern is immediately addressed - the resulting menu is certainly feasible. Finally, observe that certainly the first budget constraint is strongly satisfied - essentially becuase Λ is budget-feasible at 1 and we have chosen the correct menu to strongly satisfy the budget constraint. If $\gamma_1 = 0$, then we charge the minimum possible price, which is guaranteed to be at most b_1. If $\gamma_1 > 0$, we have hard-coded $(\vec{\pi}_1, \vec{p}_1)$ so that $p_1(\bar{v}) = b_1$. So both concerns are addressed, and by the previous observations/corollaries/propositions, the proposition is proved. □

Proposition C.5. *Let $\Phi_1^\Lambda(\bar{v}) = 0$. Then the following mechanism solves LP2(Λ) and* almost *strongly satisfies complementary slackness. There might exist one i with $\gamma_i > 0$ but $p_i(\bar{v}) > b_i \cdot \pi_i(\bar{v})$, but complementary slackness is strongly satisfied for all other constraints. Additionally, there might exist some $j \geq i$ with $p_j(\bar{v}) > \pi_j(\bar{v}) \cdot b_j$.*

- $(\pi_1(v), p_1(v)) = (0, 0)$ *for all* v.
- $(\vec{\pi}_i, \vec{p}_i) =$ *Generate-Menu*$(i, \vec{\pi}_{i-1}, \vec{p}_{i-1})$ *for all* $i > 1$.

Proof of Proposition C.5. Again, recall that we only need to check two concerns. Note that in this case, we will have $\pi_j(\bar{v}) = p_j(\bar{v}) = 0$ for all j such that $\Phi_j^\Lambda(\bar{v}) = 0$, so the first concern is addressed. However, we can't make any claims about $p_i(\bar{v})$ for the unique (if it exists) i such that $\gamma_i > 0$ and $\Phi_i^\Lambda(\bar{v}) > 0$. Note that the payment can only be too large (and not too small), as $v_i^+ \geq b_i$. □

Proposition C.6. *Let Λ be budget-feasible at all budgets, and $\Phi_1^\Lambda(\bar{v}) = 0$. Let also $[v^-, v^+ - 1]$ be the ironed interval containing $(1, b_1)$. There exists a $q \in [0, 1]$ such that the following mechanism solves LP2(Λ) and strongly satisfies complementary slackness.*

- $(\pi_1(v), p_1(v)) = (0, 0)$ *for all* $v < v^-$.

- $(\pi_1(v), p_1(v)) = (q, b_1 q)$ *for all* $v \geq v^+$, *and* $(\pi_1(v), p_1(v)) = (q\frac{v^+ - b_1}{v^+ - v^-}, v^- q \cdot \frac{v^+ - b_1}{v^+ - v^-})$ *for all* $v \in [v^-, v^+)$.
- $(\vec{\pi}_i, \vec{p}_i) =$ *Generate-Menu*$(i, \vec{\pi}_{i-1}, \vec{p}_{i-1})$ *for all* $i > 1$.

Proof of Proposition C.6. The high-level idea is that as we increase q all the way to 1, we get an IC mechanism that charges b_1 to receive the item deterministically. In this case, clearly $p_i(\bar{v}) < b_i$. If we keep q all the way at 0, we might instead have $p_i(\bar{v}) > b_i$. Note that $p_i(\bar{v})$ is a continuous function in q, so by the intermediate value theorem, there is some q such that $p_i(\bar{v}) = b_i$. Observe that for this q, we necessarily have $\pi_j(\bar{v}) < 1$ for all $j < i$, as otherwise we couldn't possibly have IC, $p_j(\bar{v}) = b_j \cdot \pi_j(\bar{v})$, and $\pi_i(\bar{v}) = 1$, all three of which are guaranteed to hold for *all* q by Corollary 4.13. Finally, we just need to ensure that $\pi_j(\bar{v}) \geq 0$ for all $j < i$. This is a potential concern if Λ is not budget-feasible at j (i.e. if $v_j^+ < b_j$, then we might wind up with negative probabilities/prices). But due to Corollary C.3, this will never result when Λ is budget-feasible at j. □

Proof of Lemma 4.15. First, consider the case that $\Phi^{\Lambda}_{s^*}(\bar{v}) > 0$. Then immediately from the second bullet of Theorem 4.10, the corollary follows. If instead $\Phi^{\Lambda}_{s^*} = 0$, it is because we went backwards in step above, which means that $s^* + x$ is linked to $s^* + x - 1$ for all x such that $\Phi^{\Lambda}_{s^*+x-1}(\bar{v}) = 0$. Let y be the largest such x. So if there exists some $x \in [0, y]$ where Λ isn't budget-feasible at $s^* + x$, consider the consequences:

- If $x = y$, there is an immediate contradiction to the second bullet in Theorem 4.10.
- Otherwise, we must have $\pi_{s^*+x}(v) = p_{s^*+x}(v) = 0$ for all v by Lemma 3.5.
- By Corollary 4.9, we necessarily have $\pi_{s^*+z}(v) = p_{s^*+z}(v) = 0$ for all v and all $z \in [x, y)$.
- Because Λ links $s^* + y$ and $s^* + y - 1$, by Corollary ?? there is a unique menu that can possibly be offered to budget $s^* + y$. Because $\pi_{s^*+y-1}(v) = p_{s^*+y-1}(v) = 0$ for all v, this menu necessarily offers a single non-trivial option, and to the highest ironed interval, $[v^+, \bar{v}]$ to receive the item with probability 1 at price v^+.
- Because Λ links $s^* + y$ and $s^* + y - 1$, we *must* have $v^+ \neq b_{s^*+y}$ (this is the reason for the exception in the definition). So the budget constraint won't bind, which violates complementary slackness as we necessarily had $\gamma_{s^*+y} > 0$ to start going backwards in the first place.

So all together, this says that we reach a contradiction if Λ isn't budget-feasible above the seed budget. □

Proof of Proposition 4.16. Identical to that of Proposition C.6, combined with Lemma 4.15 to guarantee that Λ is budget-feasible wherever the above proposes a non-trivial menu. We also observe that because (by definition), Λ does not link s^* to $s^* - 1$, $\alpha_i(v^+) = 0$. This means that we are free to set $\pi_{s^*}(\cdot), p_{s^*}(\cdot)$ as above without violating complementary slackness, as v^+ need not be indifferent between their (allocation, price) and $(0, 0)$. □

D THE GENERAL CASE: VALUE-REGULAR DISTRIBUTIONS

In this section, we provide a tighter characterization in the event that each F_i is value-regular. The result matches that of Che and Gale.

Theorem D.1 ([Che and Gale, 2000]). *Let each F_i be value-regular. Then the optimal mechanism takes the following form, where ℓ_i denote the number of non-trivial menu options for types with budget b_i.*

- *There exists an i such that $b_j \cdot \pi_j(\bar{v}) = p_j(\bar{v})$ for all $j \leq i$, and there exists an a such that $\pi_j(\bar{v}) = 1$ for all $j \geq a$.*
- *Above, $a \in \{i, i+1\}$. So there is at most one j for which both $b_j \cdot \pi_j(\bar{v}) = p_j(\bar{v})$ and $\pi_j(\bar{v}) = 1$.*
- *$\ell_1 \leq 1$, and $\ell_j \leq \ell_{j-1} + 1$ for all $j \leq a$. $\ell_j \leq \ell_{j-1}$ for all $j > a$.*
- *The menu for types with budget b_j is obtained by Generate-Menu with input equal to the menu for types with budget b_{j-1}.*

Proof. Let Λ denote an optimal dual of the form from Theorem 4.10. Let's first consider the flow that sets $\gamma_i' = \alpha_i'(v) = \lambda_i'(v-1, v) = 0$ for all i, v, and $\lambda_i'(v, v-1) = 1 - F_i(v)$ for all i, v. This results in $\Phi_i^{\lambda'}(v) = v - \frac{1-F_i(v)}{f_i(v)}$, and as each F_i is value-regular, we have $f_i(\cdot)\Phi_i^{\lambda'}(\cdot)$ monotone non-decreasing at this point. Now, let's consider changing γ_i' and $\alpha_i'(v)$ one i at a time, starting from $i = k$. First change γ_i' from 0 to γ_i. This decreases $f_i(\bar{v}) \cdot \Phi_i^{\lambda'}(\bar{v})$ by $\gamma_i(\bar{v} - b_i)$, and increases $f_i(v) \cdot \Phi_i^{\lambda'}(v)$ by γ_i for all $v < \bar{v}$. So $f_i(\cdot)\Phi_i^{\lambda'}(\cdot)$ is *almost* monotone non-decreasing, except perhaps up at $\bar{v}$. Ironing will therefore create a single ironed interval at the top, $[v_i^*, \bar{v}]$, and there will be no other ironing. Next, increase $\alpha_i'(v)$ to $\alpha_i(v)$ as in Proposition 4.5. This won't create any new ironed intervals within budget b_i. Moreover, $f_{i-1}(w) \cdot \Phi_{i-1}^{\lambda'}(w)$ will decrease by $\sum_{v>w} \alpha_i(v)$, so we maintain monotonicity within budget $i-1$. So at the end of the above procedure, we wind up with the dual exactly Λ, of the form guaranteed by Theorem 4.10. Each i has only one ironed interval, at the top. Chasing through what this means for the procedure Generate-Menu, we see that none of the menu options for budget i get "split" into two options on the menu for budget $i+1$. Moreover, as F_1 is itself value-regular, no ironing is needed in the menu for budget b_1 either. The reason that no new menu options arise as we go from j to $j+1$ when $j \geq a$ is because one of the options on the menu for j was to receive the item with probability 1 at some price. The "new" option on the menu for $j+1$ is to receive the item with probability 1 at a possibly lower price. So if there is a "new" option on the menu for $j+1$, it is because it "kicks out" an option on the menu for j. □