



Optimal Price Discrimination for Randomized Mechanisms

SHAO-HENG KO, Computer Science Department, Duke University

KAMESH MUNAGALA, Computer Science Department, Duke University

We study the power of price discrimination via an intermediary in bilateral trade, when there is a revenue-maximizing seller selling an item to a buyer with a private value drawn from a prior. Between the seller and the buyer, there is an intermediary that can segment the market by releasing information about the true values to the seller. This is termed signaling, and enables the seller to price discriminate. In this setting, Bergemann et al. [7] showed the existence of a signaling scheme that simultaneously raises the optimal consumer surplus, guarantees the item always sells, and ensures the seller's revenue does not increase.

Our work extends the positive result of Bergemann et al. to settings where the type space is larger, and where optimal auction is randomized, possibly over a menu that can be exponentially large. In particular, we consider two settings motivated by budgets: The first is when there is a publicly known budget constraint on the price the seller can charge [12] and the second is the FedEx problem [19] where the buyer has a private deadline or service level (equivalently, a private budget that is guaranteed to never bind). For both settings, we present a novel signaling scheme and its analysis via a continuous construction process that recreates the optimal consumer surplus guarantee of Bergemann et al. and further subsumes their signaling scheme as a special case. In effect, our results show settings where even though the optimal auction is randomized over a possibly large menu, there is a market segmentation such that for each segment, the optimal auction is a simple posted price scheme where the item is always sold.

The settings we consider are special cases of the more general problem where the buyer has a private budget constraint in addition to a private value. We finally show that our positive results do not extend to this more general setting, particularly when the budget can bind in the optimal auction, and when the seller's mechanism allows for all-pay auctions. Here, we show that any efficient signaling scheme necessarily transfers almost all the surplus to the seller instead of the buyer.

CCS Concepts: • Theory of computation → Algorithmic game theory; • Mathematics of computing → Probabilistic algorithms.

Additional Key Words and Phrases: bilateral trade, auctions, signaling

ACM Reference Format:

Shao-Heng Ko and Kamesh Munagala. 2022. Optimal Price Discrimination for Randomized Mechanisms. In Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22), July 11–15, 2022, Boulder, CO, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3490486.3538335>

1 INTRODUCTION

A canonical problem in mechanism design is that of bilateral trade – a single seller selling an item to a buyer, or equivalently, an infinite supply of identical items to a stream of buyers. We assume the item has no value to the seller. Typically, the buyers directly interact with the seller, who given distributional knowledge of the buyer's private valuation, runs an incentive compatible mechanism in order to maximize its own revenue. This mechanism is termed the optimal auction, which in this

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '22, July 11–15, 2022, Boulder, CO, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9150-4/22/07...\$15.00

<https://doi.org/10.1145/3490486.3538335>

case is just a “take it or leave it” (or monopoly) price offered to the buyer [24]. Such a mechanism could potentially lead to loss in social welfare, since the item is unsold if the value of the buyer falls below the monopoly price.

Price Discrimination via an Intermediary. Now imagine a platform or exchange that mediates the interaction between the buyers and the seller. This intermediary observes the private value of each arriving buyer, and it uses this information to segment the market of buyers by providing additional information (or a signal) to the seller. The seller uses this signal (or additional information) to price discriminate between different types of buyers by running separate optimal auctions for each signal. Such intermediaries are motivated by modern platforms such as ad exchanges [1–4], which help buyers (in this case, advertisers) interact with sellers (in this case, publishers of content).

The ad exchange is usually run by a search engine or social media company that can use its own data to accurately learn values of advertisers for various ad slots, and selectively release this information to the publishers who then set the prices based on this information.

Such an intermediary clearly benefits the seller’s revenue; after all, the seller has more information that enables it to price discriminate. Counter-intuitively, as shown by Bergemann, Brooks, and Morris [7], it can also lead to more utility for the buyers, and hence larger social welfare! In fact, the main result of [7] is remarkable – there is a signaling scheme such that the item always sells (so that the social welfare is as large as possible), while the seller’s revenue is the same as that without signaling. Therefore, the entire extra social surplus due to signaling goes to the buyer as its utility (or its consumer surplus). This is the best possible outcome buyers can expect given that the seller controls the auction (or the pricing scheme).

Though this result is striking, the underlying setting is the simplest possible – there is one seller and one buyer (bilateral trade), so that the optimal auction given distributional information about the buyer’s valuation (either with or without signaling) is a posted price scheme that can be computed in closed form. Given a prior distribution G on the valuation of the buyer with a monopoly price \square the algorithms in [7] sequentially construct signals while maintaining the invariant that at any step, the monopoly price of the residual distribution after subtracting the signals constructed so far remains \square This strong invariant seems critical to the guarantee on social optimality achieved in [7]. This makes the positive results appear specific to this setting. The question we ask in this paper is:

Can the positive results in [7] be extended to significantly more general settings where the optimal auction need not be so simple?

In this paper, we answer this question in the affirmative by extending the positive results in [7] to settings where the optimal auction can be randomized, even with exponential menu complexity.¹

Concretely, we study the setting where the type space of the buyer is discrete, and includes not only their private value for the item, but also a budget or deadline. Our positive results concern two settings. In the first setting, there is a publicly known upper bound on the price any buyer can be charged; this is termed the public budget setting in literature [12, 23]. In the second setting, the buyer has a private deadline by which time they need to receive the item; receiving it later than the deadline yields the buyer no value. This can be equivalently viewed as a private service level for the product. The private values and deadlines are assumed to be drawn from an arbitrary two-dimensional discrete prior distribution. The auction thus needs to be incentive compatible in the sense that the buyer should not derive more utility by reporting a tighter deadline. This is termed the FedEx problem in literature [19, 26].

¹In randomized auctions where the outcome for each buyer type is a (payment, allocation) pair, one can equivalently view the collection of all such pairs as a menu, from which the buyers can choose the best one for them. This encodes incentive compatibility. The menu complexity refers to the size of this set.

In both settings, the optimal auction can be randomized. In the public budget case, the randomization is over two possible menu options [12], while for the FedEx case, the randomization can be over a menu that can be exponentially large in the number of deadlines [15, 19, 26].

1.1 Our Results

Our main contribution is a novel signaling scheme for price discrimination in bilateral trade for the two settings of public budgets and the FedEx problem mentioned above. We show that this scheme recreates the guarantee in [7] – it achieves full social welfare (that is, it always sells the item), while ensuring the seller's revenue is the same as without signaling, thereby transferring all excess social surplus to the buyer. In particular, this shows the following surprising corollary: For both these problems, even though the optimal auction is randomized in general, there is a decomposition of the prior distribution into a collection of signals such that for each signal, the optimal auction is a simple posted price scheme where the item is always sold.

The first technical highlight of our paper is a reinterpretation of the signaling schemes for bilateral trade in [7] as a continuous time process. In this process, an infinitesimal quantity of a signal is continuously removed from the prior distribution, and we maintain two invariants at any time instant: (a) An optimal auction for the signal being removed is efficient; and (b) the revenue for this signal is exactly equal to the rate of decrease in revenue of the optimal auction on the current prior distribution. We present a general proof technique based on convexity, which in essence shows that any algorithm that satisfies invariants (a) and (b) recreates the guarantee in [7], regardless of how complex the optimal auction for the setting in consideration is. The advantage of this approach is that it enables us to sidestep both the fine-grained characterization in [7] of how the prior changes when signals are removed from it, as well as proving their invariant that the optimal auction is preserved as signals are removed from the prior.

The continuous framework provides a unifying method to analyze signaling schemes for both the public budget and the private deadline settings. However, we still need a careful choice of how to run the continuous process so that the two invariants hold. This is particularly challenging for the FedEx problem, since the type space here is two-dimensional, representing the values and deadlines. As we show in the full paper [22], a naive approach that applies the scheme in [7] separately to the marginal induced by each deadline raises too little consumer surplus. We therefore need to develop an approach that carefully hides both the value and deadline information, and our main algorithmic contribution is the development of a novel signaling method in such spaces (Section 4) that achieves precisely this. This forms our second technical highlight.

Our signaling scheme and analysis require discrete (finite support) priors over valuations. Following [9, 25], such priors are also an arbitrary good approximation for continuous priors via discretization. Our analysis requires a characterization of the optimal auctions in this setting, which we present in Theorems 3.4 and 4.6. These are the discrete analogs of results in [12, 19] for

continuous priors, and show that the optimal auction is a distribution over posted prices that satisfy certain nice properties. The characterizations we require are much coarser than those in [12, 19] and we present stand-alone alternate proofs of these properties that are tailored to the discrete nature of the priors. In particular, the proof for the deadline setting (Theorem 4.6) uses convexity in the primal instead of duality, and this technique may be of independent interest.

Impossibility for Private Budgets. We finally ask how far we can push this positive result. Towards this end, we consider the generalization of the above settings to the private budget setting [13, 15]. Here, the buyer has a private budget, and the values and budgets are assumed to be drawn from an arbitrary two-dimensional discrete prior distribution. The buyer cannot over-report her budget, but an incentive compatible auction needs to prevent under-reporting it. We assume interim rationality

to allow for all-pay auctions, or equivalently, views the item as infinitely divisible; this is a standard assumption in economics literature [13, 23]. Note that the FedEx problem is a special case where the budgets are larger than all valuations.

For private budgets, there is no signaling scheme that satisfies both criteria (a) and (b) above. This leads to a strong lower bound: Even with two values and two budgets, any efficient signaling scheme (that always sells the item) transfers all surplus to the seller, leading to vanishingly small consumer surplus. Therefore, no efficient signaling scheme can reproduce the consumer surplus guarantee in [7] to any approximation. Furthermore, even if we sacrifice efficiency, we cannot hope to achieve better than a constant approximation to the consumer surplus guarantee.

1.2 Related Work

Our problem falls in the general framework of information design [8] where an information mediator can deliberately provide additional information to impact the behavior of agents in given mechanisms; this is also sometimes termed signaling or persuasion [16]. The Bayesian Persuasion model [21] is a special case of information design with only one agent (often called the receiver) receiving additional information that comes from a sender with more knowledge of the state of nature. Given the signal, the receiver chooses the actions to maximize her own utility based on her belief of the state of nature (which may be influenced by the signal). Therefore, the sender designs the signals so that the receiver, acting in her own interest, maximizes some utility function the sender cares about. This problem is studied from various theoretical perspectives [6, 17, 18] as well as in different application domains [7, 11, 20, 28].

Starting with the seminal work of Bergemann et al. [7], there has been a line of work [10, 14, 17, 20, 27] on Bayesian persuasion in the bilateral trade model and its extensions. In this context, the sender is an intermediary and the receiver is the seller, who given the signal, implements an incentive-compatible auction to maximize expected revenue. The sender, on the other hand, is interested in maximizing consumer surplus or social welfare. In the versions we study with budgets or deadlines, the receiver's action space is the set of all randomized pricing rules, instead of just the posted prices in the basic setting [7]. Our main contribution is to show the existence of socially efficient signaling schemes that preserve receiver utility (the revenue) and maximally increases sender utility (consumer surplus) despite this additional complexity. We note that for other non-trivial extensions of bilateral trade, for instance, the multi-buyer auction setting in [5] and the multi-item auction setting in [20], it may in general not be possible to find socially optimal signaling schemes that preserve seller revenue. This makes our positive results all the more surprising.

As mentioned before, our work crucially requires a characterization of optimal randomized auction in the respective settings. For public budgets, Laffont and Roberts [23] show that the optimal auction is a posted price scheme assuming regular distributions; for general priors, Chawla et al. [12] show it is a lottery over two options. Che and Gale [13] consider private budgets with a decreasing marginal revenue assumption, and show it is a different price curve for each budget. Fiat et al. [19] and subsequently Devanur and Weinberg [15] use duality to respectively generalize this characterization to private deadlines and private budgets with arbitrary priors; however, the characterization in the latter case is not closed form. Since we use finite support priors, we present stand-alone proofs of the required characterizations, and these may be of independent interest.

Organization. In Section 2, we present preliminaries for optimal auction design and signaling. In Section 3, we present the signaling scheme and analysis for the public budget case. In Section 4, we present our main result – the new signaling scheme for the FedEx problem, where the deadlines are private. In Section 5, we present the impossibility result for the private budget setting with interim rationality. All omitted proofs are in the full version of the paper [22].

2 PRELIMINARIES

2.1 Optimal Auctions with Budgets

We consider a single seller selling an item to a single buyer with private valuation \square and private budget \square as a hard upper bound of payment. It is known that optimal auctions with budgets require randomized allocations [12, 13, 15]: The buyer's utility is $(\square \cdot \square - \square)$ if she pays a price of $\square \leq \square$ to get the item with probability $\square \in [0, 1]$, and is $-\infty$ if $\square > \square$. Throughout the paper, we focus on interim IR auctions where the buyer pays at most \square before learning whether or not she receives the item. As mentioned before, this is the standard model for studying budget constrained auctions in economics literature [13, 23], and allows for all-pay auctions. Alternatively, it models ex-post rationality assuming the item is infinitely divisible, and $\square \in [0, 1]$ represents the fraction of the item the buyer obtains at price \square .

The joint distribution $(\square, \square) \in G$ is common knowledge and supported on a discrete set $\text{supp}(G) \subseteq \{\square_1, \dots, \square_n\} \times \{\square_1, \dots, \square_n\}$, where $0 < \square_1 < \dots < \square_n$ and $0 < \square_1 < \dots < \square_n$. For $\square = 1, 2, \dots, \square$, let G_{\square} represent the marginal distribution of \square given $\square = \square$, and define $\square_{G_{\square}}$ as the probability mass function of G_{\square} , i.e., $\square_{G_{\square}}(\square) = \Pr_{\square \in G}[\square = \square] = \Pr_{(\square, \square) \in G}[\square = \square \mid \square = \square]$. Let $\square_{G_{\square}}(\square) = \Pr_{\square \in G}[\square \leq \square]$ and $\square_{G_{\square}}(\square) = \Pr_{\square \in G}[\square \geq \square]$. We assume the item holds no value to the seller; therefore, the maximum social welfare is $SW^*(G) = E_{(\square, \square) \in G}[\square]$, and is achieved by any auction that always makes the trade happen (or sells the item).

Optimal Auctions. It is known [12, 13, 15, 23, 24] that the revenue maximizing auction for the seller can be described using lotteries or randomized allocation rules. Specifically, each buyer type with valuation \square and budget \square is associated to a payment $\square \geq 0$ and an allocation probability $\square \in [0, 1]$ to receive the item. Note that in the interim-IR setting the buyer pays \square upfront regardless of whether she receives the item.

Following [13, 15], we assume buyer with type (\square, \square) cannot report a budget larger than \square ; this can be enforced by collecting the entire reported budget with a small probability, or simply by a cash bond that requires the full reported budget. Further, the setting where the IC constraints are only enforced to smaller budgets is more challenging for designing optimal auctions [13]. By the revelation principle, it is sufficient to consider lotteries that are incentive compatible, i.e., for all \square and \square a buyer of type (\square, \square) receives maximum possible utility from reporting her true type (\square, \square) and thereby receiving the item with allocation probability \square at price \square .

The revenue optimal auction can be computed by the following LP from [15].

$$\begin{aligned}
 & \text{Budgets}(G) \subseteq \max_{\substack{\{\square_1, \dots, \square_n\} \times \{\square_1, \dots, \square_n\} \\ \square_1 < \dots < \square_n \\ \square_1 < \dots < \square_n}} \Pr_{\square \in G} \cdot \sum_{\square=1}^{\square} \square_{G_{\square}}(\square) \cdot \square \\
 \text{s.t.} \quad & \square \cdot \square - \square \geq \square \cdot \square - \square, \quad \square 1 \leq \square \leq \square 1 \leq \square \leq \square \quad (\text{Same-budget IC}) \\
 & \square \cdot \square - \square \geq \square \cdot \square_{(\square-1)} - \square, \quad \square 1 \leq \square \leq \square 2 \leq \square \leq \square \quad (\text{Inter-budget IC}) \\
 & \square_{(\square-1)}, \square \cdot \square - \square \geq 0, \quad \square 1 \leq \square \leq \square 1 \leq \square \leq \square \quad (\text{IR}) \\
 & 0 \leq \square \leq 1, \quad \square 1 \leq \square \leq \square 1 \leq \square \leq \square \quad (\text{Feasibility}) \\
 & \square \leq \square, \quad \square 1 \leq \square \leq \square 1 \leq \square \leq \square \quad (\text{Budgets})
 \end{aligned}$$

By transitivity, the same-budget and inter-budget IC constraints imply all necessary IC constraints so that the buyer with valuation \square and budget \square does not misreport with some valuation $\square \neq \square$ and/or some budget $\square < \square$.

Definition 2.1. For the revenue maximizing auction $(\{\mathbb{P}_i\}_{i=1}^n, \{\mathbb{Q}_i\}_{i=1}^n)$ that is the optimal solution to $\text{Budgets}(G)$, denote

$$R(G) = \sum_{i=1}^n \Pr[G = (\mathbb{Q}_i, \mathbb{P}_i)] \cdot \mathbb{P}_i$$

$$SW(G) = \sum_{i=1}^n \Pr[G = (\mathbb{Q}_i, \mathbb{P}_i)] \cdot \mathbb{Q}_i \cdot \mathbb{P}_i,$$

and

$$CS(G) = \Pr[G = (\mathbb{Q}_i, \mathbb{P}_i)] \cdot (\mathbb{Q}_i \cdot \mathbb{P}_i - \mathbb{P}_i)$$

as the expected revenue (generated by the seller), the expected social welfare, and the expected consumer surplus (generated for the buyer), respectively. Then we have $CS(G) + R(G) = SW(G)$.²

We now specify two special cases of the budgeted problem for which we derive positive results.

Optimal Auctions with Public Budget. The first special case we consider is the public budget setting [12, 23] where $\mathbb{Q} = 1$, the budget $\mathbb{P} = \mathbb{P}_1$ is public information, and the only marginal distribution is $G = G_1$. This setting is motivated by the seller having an upper bound on the price they can charge any buyer, say due to regulation or other considerations.

In this case we omit the subscripts by referring to $\mathbb{P}_G(\mathbb{Q})$, $\mathbb{Q}_G(\mathbb{Q})$, and $\mathbb{P}_G(\mathbb{Q})$, and use \mathbb{Q} and \mathbb{P} as shorthand for the payment variables \mathbb{P}_1 and allocation variables \mathbb{Q}_1 , respectively. For this case, the optimal auction is captured by the following special case of Budgets with $\mathbb{Q} = 1$:

$$\begin{aligned} \text{Public}(G) \quad & \max_{\{\mathbb{P}\}, \{\mathbb{Q}\}} \quad \sum_{i=1}^n \mathbb{P}_G(\mathbb{Q}_i) \cdot \\ & \text{s.t.} \quad \mathbb{Q}_i \cdot \mathbb{Q} - \mathbb{P}_i \geq \mathbb{Q}_i \cdot \mathbb{Q}_i - \quad \quad \quad \mathbb{Q} \leq \mathbb{Q}_i \leq \mathbb{Q} \quad (\text{IC}) \\ & \quad \mathbb{Q}_i \cdot \mathbb{Q} - \mathbb{P}_i \geq 0, \quad \quad \quad \mathbb{Q} \leq \mathbb{Q}_i \leq \mathbb{Q} \quad (\text{IR}) \\ & \quad 0 \leq \mathbb{Q}_i \leq 1, \quad \quad \quad \mathbb{Q} \leq \mathbb{Q}_i \leq \mathbb{Q} \quad (\text{Feasibility}) \\ & \quad \mathbb{P}_i \leq \mathbb{P} \quad \quad \quad \mathbb{Q} \leq \mathbb{P} \leq \mathbb{Q} \quad (\text{Budget}) \end{aligned}$$

We devise price discrimination schemes for the public budget setting in Section 3.

Optimal Auctions with Deadlines. In this setting [19, 26], we consider a single seller selling an identical item with different levels of service quality to a single buyer. The buyer now has private valuation \mathbb{Q} (conditioned on getting the item with at least her desired level of quality) and a private desired level of quality \mathbb{Q} . One can think of \mathbb{Q} as either a personal deadline for shipping options, or as a level of service quality for a product. Keeping with previous work, we will refer to \mathbb{Q} as deadlines throughout.

The buyer's utility is $(\mathbb{Q} \cdot \mathbb{Q} - \mathbb{P})$ if she pays a price of \mathbb{P} to get the item with a probability of \mathbb{Q} at some point before or right at her deadline. She incurs utility $-\mathbb{P}$ if she gets the item later than her deadline, since in this case, she accrues no value from the item. As observed in [19], it is sufficient to consider auctions that, for each buyer with deadline \mathbb{Q} , only allocates the item right at the \mathbb{Q} -th deadline (if at all). This is because a buyer does not get any additional utility if she receives the item at some point earlier than her own deadline. Furthermore, the buyer weakly prefers getting nothing over getting the item after her own deadline for some price.

This setting is a special case of the private budget setting with large budgets, that is, when $\mathbb{P}_i < \mathbb{Q}_i$ holds for all budget types $\mathbb{Q} = 1, \dots, \mathbb{Q}$. As every budget is above the highest possible valuation, by

²If there are multiple optimal auctions maximizing $R(G)$, we break ties by defining $(\{\mathbb{P}_i\}_{i=1}^n, \{\mathbb{Q}_i\}_{i=1}^n)$ to be the auction that maximizes $SW(G)$ among the optimal solutions. This auction must maximize $CS(G)$ as well.

the IR constraint, the optimal auction never sets a price above \bar{v}_i for any buyer with budget \bar{b}_i , and thus the budget constraint $\bar{b}_i \leq \bar{v}_i$ in Budgets can be omitted.

For this case, we simplify the notations by denoting the joint distribution $(\bar{v}, \bar{b}) \in G$ supported on $\text{supp}(G) \subseteq \{\bar{v}_1, \dots, \bar{v}_n\} \times \{1, \dots, \bar{b}_i\}$, where $0 < \bar{v}_1 < \dots < \bar{v}_n$. The deadlines can be represented as $\{1, \dots, \bar{b}_i\}$ since their cardinal values do not matter. For $i = 1, 2, \dots, n$, G_i now represents the marginal distribution of \bar{v}_i given $\bar{b}_i = \bar{b}_i$, and the corresponding probability mass function of G_i is

$$q_{G_i}(\bar{v}_i) = \Pr_{\bar{v}, \bar{b} \in G} [\bar{v}_i = \bar{v}_i] = \Pr_{(\bar{v}, \bar{b}) \in G} [\bar{v}_i = \bar{v}_i \mid \bar{b}_i = \bar{b}_i].$$

Let $\bar{q}_{G_i}(\bar{v}_i) = \Pr_{\bar{v}, \bar{b} \in G} [\bar{v}_i \leq \bar{v}_i]$ and $\bar{q}_{G_i}(\bar{v}_i) = \Pr_{\bar{v}, \bar{b} \in G} [\bar{v}_i \geq \bar{v}_i]$. We again assume the item holds no value to the seller; therefore, the maximum social welfare is $\text{SW}^*(G) = E_{(\bar{v}, \bar{b}) \in G} [\bar{v}_i]$, and is achieved by any auction that always allocates the item to each buyer right at her personal deadline.

The revenue maximizing randomized incentive compatible auction for the deadlines setting is thus the following:

$$\begin{aligned} \text{Deadlines}(G) \in \max_{\substack{\{q_{G_i}(\bar{v}_i)\}_{i=1}^n, \{q_{G_i}(\bar{v}_i)\}_{i=1}^n}} & \sum_{i=1}^n \Pr_{(\bar{v}, \bar{b}) \in G} [\bar{v}_i = \bar{v}_i] \cdot q_{G_i}(\bar{v}_i) \cdot \bar{v}_i \\ \text{s.t.} & \bar{v}_i \cdot q_{G_i}(\bar{v}_i) \geq \bar{v}_i \cdot q_{G_i}(\bar{v}_i) - q_{G_i}(\bar{v}_i), \quad \bar{v}_1 \leq \bar{v}_i \leq \bar{v}_n, 1 \leq i \leq n \quad (\text{Same-deadline IC}) \\ & \bar{v}_i \cdot q_{G_i}(\bar{v}_i) \geq \bar{v}_i \cdot q_{G_i}(\bar{v}_{i-1}) - q_{G_i}(\bar{v}_{i-1}), \quad \bar{v}_1 \leq \bar{v}_i \leq \bar{v}_n, 2 \leq i \leq n \quad (\text{Inter-deadline IC}) \\ & \bar{v}_i \cdot q_{G_i}(\bar{v}_i) \geq 0, \quad \bar{v}_1 \leq \bar{v}_i \leq \bar{v}_n, 1 \leq i \leq n \quad (\text{IR}) \\ & 0 \leq q_{G_i}(\bar{v}_i) \leq 1, \quad \bar{v}_1 \leq \bar{v}_i \leq \bar{v}_n, 1 \leq i \leq n \quad (\text{Feasibility}) \end{aligned}$$

Note that Deadlines is a special case of Budgets where the budget constraint is omitted, and the IC constraints in Deadlines prevent misreporting a lower deadline.

Remarks. We first note that though our scheme require discrete priors over valuations, these also serve as arbitrarily good approximations to continuous priors via simple discretization [9, 25]. Secondly, note that the optimal auctions with interim IR coincides with that for ex-post IR for both the public budget and the deadline setting; the former follows from Theorem 3.4 (or from [12]), while the latter follows because the prices are not really constrained by any budget. Therefore, our positive results in Sections 3 and 4 extend as is to ex-post IR. Our negative results in Section 5 do require interim IR.

2.2 Price Discrimination

We next introduce price discrimination via signaling by an information intermediary for the general private budget setting; specializing it to deadlines or public budgets is straightforward. The intermediary knows the type (\bar{v}_i, \bar{b}_i) of the buyer, and can propose a signaling scheme that maps the buyer's private information (i.e., a value-budget pair (\bar{v}_i, \bar{b}_i)) to a distribution over signals that conveys additional information to the seller. This makes the seller update her belief of the buyer's information via Bayes' rule. The signaling scheme thus can be seen as segmenting the market of buyers, each segment representing the conditional distribution of buyer type given the signal. Therefore, we can overload terminology and simply define a signal \bar{s} as the posterior distribution of (\bar{v}_i, \bar{b}_i) given the signal.

Signaling Scheme. Formally, a signaling scheme $\Theta = \{(\bar{s}_i, \bar{q}_i)\}_{i=1}^n$ is a collection of signals $\bar{s}_1, \dots, \bar{s}_n$ and probability weights $\bar{q}_1, \dots, \bar{q}_n > 0$, where $\sum_{i=1}^n \bar{q}_i = 1$. Here \bar{q}_i represents the

posterior distribution of the type given the \square -th signal. We also require Θ being Bayes plausible [21],

$$\sum_{\square=1}^{\infty} \Pr[\square = \square] = G, \quad (1)$$

i.e., the average signal is just the prior G . The intermediary commits to this signaling scheme before she observes the buyer type, and this scheme is public knowledge to all parties.

Upon observing private information (\square, \square) , the intermediary sends the \square -th signal with probability $\Pr[\square = (\square, \square)]$, and given this signal, if the seller uses Bayes rule to update the prior on the buyer's type, the posterior will be precisely \square . The seller then implements the revenue maximizing auction based on the updated prior \square .

Buyer Optimal Schemes. Abusing the notation defined before, we let $R(\Theta) = \sum_{\square=1}^{\infty} \Pr[\square] \cdot R(\square)$, $SW(\Theta) = \sum_{\square=1}^{\infty} \Pr[\square] \cdot SW(\square)$, and $CS(\Theta) = \sum_{\square=1}^{\infty} \Pr[\square] \cdot CS(\square)$ denote the expected revenue, the expected social welfare, and the expected consumer surplus, respectively, achieved by the signaling scheme Θ , where the expectation is now taken over all signals. As before, we have $CS(\Theta) + R(\Theta) = SW(\Theta)$.

Furthermore, $R(\Theta) \geq R(G)$; otherwise, the seller can ignore the signaling scheme Θ and implement the revenue maximizing auction based on G instead. Hence, for any possible signaling scheme Θ , we have

$$CS(\Theta) = SW(\Theta) - R(\Theta) \leq SW^*(G) - R(G)$$

as an upper bound of the expected consumer surplus. Recall that $SW^*(G)$ is the maximum possible social welfare assuming the item always sells. We define this bound on maximum achievable consumer surplus as

$$CS^*(G) \triangleq SW^*(G) - R(G) = \mathbb{E}_{(\square, \square) \sim G} [\square - R(\square)].$$

To achieve $CS(\Theta) = CS^*(G)$, the signaling scheme Θ thereby needs to satisfy (a) the item always sells, and (b) the revenue $R(\Theta)$ generated by Θ is exactly $R(G)$, i.e., the expected revenue without signaling. We call a signaling scheme buyer optimal if it achieves this upper bound.

2.3 Buyer Optimal Signaling without Budgets or Deadlines

In this case, as shown by Bergemann et al. [7] the optimal consumer surplus $CS^*(G)$ is indeed achieved by a signaling scheme Θ .

Theorem 2.1 (Bergemann et al. 's signaling schemes [7]). Suppose $\square = 1$ and $\square_1 \geq \square$. Then for any arbitrary prior G , there exists a signaling scheme Θ_G that guarantees:

- (1) efficiency: $SW(\Theta_G) = SW^*(G)$ (i.e., the item always sells);
- (2) minimum revenue: $R(\Theta_G) = R(G)$ (i.e., the seller's revenue does not increase);
- (3) maximum consumer surplus: $CS(\Theta_G) = CS^*(G) = SW^*(G) - R(G)$ (i.e., the scheme maximizes the expected consumer surplus among all possible signaling schemes.)

Note that the third property is implied from the first two. There are multiple constructions of Θ_G given in [7], and one of these is equivalent to our scheme for public budgets presented in Section 3. These schemes proceed via the notion of Equal Revenue Signals. We now introduce this notion since it is essential to our signaling scheme as well.

Definition 2.2 (Equal Revenue Signals). A valuation distribution \square over its support set $\text{supp}(\square) = \{\square_1, \dots, \square\}$ is equal revenue if it satisfies:

$$\bar{\square}_1(\square_1) \cdot \square_1 = \bar{\square}_2(\square_2) \cdot \square_2 = \dots = \bar{\square}_n(\square_n) \cdot \square_n = R(\square).$$

In other words, assuming no budgets, every valuation with nonzero probability mass in $\text{supp}(\mathbf{q})$ is an optimal monopoly price for \mathbf{q} . This distribution is unique and can be obtained as follows:

$$\mathbf{q}_1(\mathbf{q}_1) = 1 - \frac{\mathbf{q}_1}{\mathbf{q}_1; 2} \quad \mathbf{q}_2(\mathbf{q}_2) = 1 - \frac{\mathbf{q}_2}{\mathbf{q}_1} (\mathbf{q}_{1-1}) \cdot \left(1 - \frac{\mathbf{q}_2}{\mathbf{q}_1; \mathbf{q}_2}\right), \quad 2 \leq \mathbf{q} \leq \mathbf{q} - \quad \mathbf{q}_n(\mathbf{q}_n) = 1 - \frac{\mathbf{q}_n}{\mathbf{q}_n} (\mathbf{q}_{n-1}).$$

3 WARMUP: SIGNALING SCHEME FOR PUBLIC BUDGETS

In this section, we prove the analog of Theorem 2.1 when there is a public budget. We show that there is a signaling scheme that is buyer optimal with a public budget. We will show this via reinterpreting the algorithm in [7] as a continuous time process (Algorithm 1 below). The nice aspect of this interpretation is that it leads to a different proof of optimality (than [7]) via a general convexity property of the revenue of the residual prior as a function of time (see Lemma 3.5). This continuous time interpretation and convexity property will form the building blocks for our main result for the version with deadlines (the FedEx problem) in Section 4.

Interestingly, our signaling scheme for public budgets is the same as the no-budget signaling scheme in [7]; this is easy to check and we omit the proof. However, our analysis is entirely different and more generalizable to the more complex deadline setting considered later.

3.1 Signaling Algorithm

Throughout this section, the buyer's budget $\mathbf{q} = \mathbf{q}_1$ is public information, and we use \mathbf{G} to denote the prior over the buyer values $\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n\}$. We view the progress of the algorithm as continuously decreasing this prior into a residual prior, and continuously placing the remaining probability mass into the constructed signals.

We use the function $\mathbf{f}(\mathbf{q}) = \langle \mathbf{q}_1(\mathbf{q}), \dots, \mathbf{q}_n(\mathbf{q}) \rangle$ to represent the residual prior, where $\mathbf{q}_i(\mathbf{q})$ represents the remaining probability mass on type \mathbf{q}_i in the residual prior distribution at time \mathbf{q} . Strictly speaking, $\mathbf{f}(\mathbf{q})$ is not a distribution since the process we describe only guarantees $\mathbf{q}_i(\mathbf{q}) < 1$ for $\mathbf{q} > 0$. To make this a valid distribution, we place the remaining probability mass $1 - \sum_{i=1}^n \mathbf{q}_i(\mathbf{q})$ at a dummy value $\mathbf{q}_0 = 0$. We call the resulting distribution $\mathbf{G}(\mathbf{q})$. In the subsequent discussion, the notation $\mathbf{f}(\mathbf{q})$ represents the probability mass of $\mathbf{G}(\mathbf{q})$ at non-zero valuations, and we omit explicitly considering the dummy value $\mathbf{q}_0 = 0$ as part of the support of $\mathbf{G}(\mathbf{q})$. We define $\text{supp}(\mathbf{G}(\mathbf{q})) \subseteq \{\mathbf{q}_i > 0 \mid \mathbf{q}_i(\mathbf{q}) > 0\}$ and $\mathbf{q}_{\min}(\mathbf{G}(\mathbf{q})) \subseteq \min\{\mathbf{q}_i > 0 \mid \mathbf{q}_i(\mathbf{q}) > 0\}$.

We start with the prior $\mathbf{G}(0) = \mathbf{G}$ and let $\mathbf{q}_i(0) = \mathbf{q}_G(\mathbf{q}_i)$, i.e., $\mathbf{f}(0)$ is just the probability vector associated with \mathbf{G} . Our algorithm continuously takes away probability mass from $\mathbf{f}(\mathbf{q})$ and transfers it to the constructed signals, terminating when $\mathbf{f}(\mathbf{q})$ becomes 0; denote the latter time as \mathbf{q}_* .

At any time \mathbf{q} such that $\mathbf{f}(\mathbf{q}) \neq 0$, we denote $\mathbf{s}(\mathbf{q})$ as the probability distribution associated with the equal revenue distribution (see Definition 2.2) $\mathbf{q}(\mathbf{q})$ over the set of values in $\text{supp}(\mathbf{G}(\mathbf{q}))$. Note that $\mathbf{s}(\mathbf{q})$ depends on $\text{supp}(\mathbf{G}(\mathbf{q}))$ but not the $\mathbf{q}_i(\mathbf{q})$; therefore, it is fixed as long as $\text{supp}(\mathbf{G}(\mathbf{q}))$ does not change.³ Our algorithm continuously reduces $\mathbf{f}(\mathbf{q})$ at rate $\mathbf{s}(\mathbf{q})$ until $\mathbf{f}(\mathbf{q})$ becomes 0. Formally:

$$\frac{d\mathbf{f}}{d\mathbf{q}} = -\mathbf{s}(\mathbf{q}). \tag{2}$$

Since $\sum_{i=1}^n \mathbf{q}_i(\mathbf{q}) = 1$, the rate of decrease of $\sum_{i=1}^n \mathbf{q}_i(\mathbf{q})$ is also 1. Since $\mathbf{q}_1(0) = 1$, this means the process terminates at time $\mathbf{q}_* = 1$.

³Note that there exist other equal revenue distributions over different support sets; for example, any distribution with a support size of one is equal revenue. However, for the purpose of our algorithm, the equal revenue distribution must use all remaining nonzero valuations with nonzero probability mass in the residual prior.

Signals constructed. We say the type- \square -valuation \square is exhausted at time \square if $\square(\square) = 0$ but $\square(\square') > 0$ for all $\square' < \square$. The algorithm therefore terminates once all types are exhausted. Consider a maximal time interval $\square \in [\square_1, \square_2]$ in which $\text{supp}(G(\square))$ remains fixed; denote the equal revenue signal in this interval by \mathbf{s} . Therefore, $\square(\square) = \mathbf{s}$ for $\square \in [\square_1, \square_2]$. Then we have:

$$\mathbf{f}(\square_1) - \mathbf{f}(\square_2) = - \int_{\square=\square_1}^{\square=\square_2} -\mathbf{s}(\square) d\square = (\square_2 - \square_1) \cdot \mathbf{s} \quad (3)$$

Therefore, the final scheme includes a signal \mathbf{s} with weight $(\square_2 - \square_1)$. This holds for every such interval $[\square_1, \square_2]$. Since $\mathbf{s}(\square)$ changes only if some element in $f(\square)$ becomes zero, the number of signals constructed is at most \square . The signaling scheme is now formally described in Algorithm 1.

Algorithm 1 Continuous Algorithm for Public Budget Setting

Input: G

Output: $\Theta = \Theta_G^{\mathbb{R}}$

- 1: $\mathbf{f}(0) \leftarrow \langle \square_G(\square_1), \dots, \square_G(\square_b) \rangle$
- 2: Compute $\mathbf{f}(\square)$, $\square(\square)$, and $\mathbf{s}(\square)$ for all $\square \in [0, 1]$ (as defined in Equation 2)
- 3: Compute $0 < \square_1 < \dots < \square_b = 1$ so that some type is exhausted at each $\square = \square_b$; let $\square_0 = 0$
- 4: for $\square \in \{1, \dots, b\}$ do
- 5: $\square^{\mathbb{R}} \leftarrow \square - \square_{b-1}$; $\square^{\mathbb{R}} \leftarrow \square(\square_{b-1})$; $\Theta \leftarrow \Theta \cup \{(\square^{\mathbb{R}}, \square^{\mathbb{R}})\}$
- 6: return Θ

We have $\sum_{\square=1}^b \square^{\mathbb{R}} = \sum_{\square=1}^b (\square - \square_{b-1}) = \square_b - \square_0 = 1$. Further, we have Bayes plausibility (as defined in Eq. (1) in Section 2):

Observation 3.1. For any arbitrary G , let $\Theta_G^{\mathbb{R}} = \{(\square^{\mathbb{R}}, \square^{\mathbb{R}})\}_{\square \in [0, 1]}$ be the set of signals output by Algorithm 1 taking G as input. Then we have $\square \leq |\text{supp}(G)|$, and $\sum_{\square=1}^b \square^{\mathbb{R}} = G$.

3.2 Optimal Auction For Signals

We start with the easy step. We characterize the revenue-optimal auctions in the signals created by Algorithm 1 in Lemma 3.2. As an easy consequence, $\Theta_G^{\mathbb{R}}$ always sells the item, and therefore guarantees efficiency. This is the first necessary condition for buyer optimality.

Let $\square \in \Theta_G^{\mathbb{R}}$ denote a signal created by Algorithm 1 for the prior G , and let $\square_{\min}(\square)$ denote the minimum $\square > 0$ such that $\Pr_{\square \in \Theta_G^{\mathbb{R}}}[\square = \square] > 0$.

Lemma 3.2. The optimal auction for \square has the following structure:

- If $\square \leq \square_{\min}(\square)$, there is an optimal auction that posts a price of \square and raises a revenue of \square
- If $\square > \square_{\min}(\square)$, there is an optimal auction that posts price $\square_{\min}(\square)$ and raises a revenue $\square_{\min}(\square)$. Further, for every $\square \in \text{supp}(\square)$, we have $\square \cdot \square_{\min}(\square) = \square_{\min}(\square)$.

By the characterizations of the optimal auctions in the above two cases, we have the following claim that states the item always sells in $\Theta_G^{\mathbb{R}}$.

Lemma 3.3 (Efficiency of $\Theta_G^{\mathbb{R}}$). For each signal $\square \in \Theta_G^{\mathbb{R}}$, there exists a revenue optimal auction that always allocates the item. As a consequence, $\text{SW}(\Theta_G^{\mathbb{R}}) = \text{SW}^{\mathbb{R}}(G)$.

3.3 Characterization of Optimal Auction for $G(\square)$

We analyze the revenue of the signals by showing that the rate of decrease in revenue of the optimal auction for $G(\square)$ is equal to the revenue of the signal $\square(\square)$. (See Theorem 3.6.) This when integrated over time shows that the optimal revenue of the signals is exactly equal to the optimal revenue for prior G , hence showing buyer optimality.

Continuous Constraints. For the purpose of analysis, we make the constraints in $\text{Public}(G(\square))$ hold not just for $\text{supp}(G(\square))$, but for all continuous values $\square \geq 0$, where the prior possibly has zero probability mass.⁴ Among other things, this formulation allows us to argue that the revenue changes continuously as the prior changes while constructing our signals.

Formally, fix some time \square , and let $A = G(\square)$ so that $\text{supp}(A) \subseteq \text{supp}(G)$. Recall that the decision variables in $\text{Public}(A)$ are \square (the payment) and \square (the allocation probability) for all buyer types with valuation \square . We augment the variables by extending the domain to $[0, \square]$; for all $\square \in [0, \square]$, we let $\square(\square)$ and $\square(\square)$ denote the expected payment and allocation probability at $\square \in [0, \square]$. This yields the following LP, where the IC and IR constraints are extended to this domain.

$$\begin{aligned}
 \text{PublicContinuous}(A) \quad & \max_{\square(\cdot), \square(\cdot)} \sum_{\square=1}^{\square} \square(\square) \cdot \square(\square) \\
 \text{s.t.} \quad & \square \cdot \square(\square) - \square(\square) \geq \square \cdot \square(\square') - \square(\square'), \quad \square, \square' \in [0, \square], \quad (\text{Cont. IC}) \\
 & \square \cdot \square(\square) - \square(\square) \geq 0, \quad \square \in [0, \square], \quad (\text{Cont. IR}) \\
 & 0, 0 \leq \square(\square) \leq 1, \quad \square \in [0, \square]. \quad (\text{Feasibility}) \\
 & \square(\square) \leq \square \quad (\text{Budget})
 \end{aligned}$$

From Definition 2.1, $R(A)$ is the optimal revenue achievable by $\text{Public}(A)$. Denote by $\tilde{R}(A)$ the optimal revenue achievable by $\text{PublicContinuous}(A)$. Clearly, $R(A) \geq \tilde{R}(A)$.

We now present the main characterization result for the optimal solution to this LP. This can be viewed as a discrete analog of the characterization for continuous priors in [12]. We present a stand-alone proof for our discrete setting in the full paper [22] via convexity of the utility curve.

Theorem 3.4. For any prior $A = G(\square)$ with $\square > \square_{\min}(A)$, there exists a set of valuations $\{\square_1, \square_2, \dots, \square_{\square}\} \subseteq \text{supp}(A)$ and weights $\square_1, \square_2, \dots, \square_{\square} \in (0, 1]$ such that $\sum_{\square=1}^{\square} \square = 1$, and the optimal revenue of $\text{PublicContinuous}(A)$ is $\tilde{R}(A) = \sum_{\square=1}^{\square} \square \cdot \square \cdot \tilde{R}_A(\square)$.

3.4 Revenue Preservation in Algorithm 1

We are now ready to prove the second necessary criterion for buyer optimality: $R(\Theta_G) = R(G)$, i.e., Algorithm 1 minimizes the expected seller revenue through signaling. As mentioned above, the key step (Theorem 3.6) is to argue that the rate of decrease of revenue of $\text{PublicContinuous}(G(\square))$ exactly equals the optimal revenue of the signal $\square(\square)$. This when combined with Lemma 3.3 gives buyer optimality. As a side effect, this will also show the optimal objectives of $\text{Public}(G)$ and $\text{PublicContinuous}(G)$ are identical.

Convexity of Revenue in $\text{PublicContinuous}(G(\square))$. The next lemmas bound the continuous flow of revenue being transferred into the signals. At any time \square , let $\tilde{R}(\square)$ denote the optimal revenue of $\text{PublicContinuous}(G(\square))$.

⁴It follows from [9, 25] that this formulation is equivalent to $\text{Public}(G(\square))$, though we will not need this equivalence.

Our first key step is fairly generic, and shows that $\tilde{R}(\square)$ is convex⁵ in any time interval $(\square_{-1}, \square_0)$ where $\text{supp}(G(\square))$ (and hence $\square(\square)$) does not change for $\square \in (\square_{-1}, \square_0)$. Let this signal be \square and its corresponding probability vector to be $\mathbf{s} = \{\square_i\}$ over $\square \in (\square_{-1}, \square_0)$.

Consider the decrease in revenue of any feasible solution $M \in \mathcal{B}(\square(\cdot), \square(\cdot))$ of $\text{PublicContinuous}(G(\square))$. The revenue of the solution M in $\text{PublicContinuous}(G(\square))$ is given by $\tilde{R}_M(\square)$

$\square(\square) \cdot \square(\square)$. Fixing this M , for each $\square \in (\square_{-1}, \square_0)$, we have:

$$\frac{d\tilde{R}_M(\square)}{d\square} = \frac{d}{d\square} \sum_{i=1}^{\square} \square_i(\square) \cdot \square(\square_i) = - \sum_{i \in \text{supp}(G(\square))} \square_i \cdot \square(\square_i), \quad (4)$$

where the last equality uses Eq. (2) and the fact that $\square_i = 0$ for any $\square_i \notin \text{supp}(G(\square))$. Therefore, for each M , its revenue $R_M(\square)$ decreases linearly with time. Further, since the constraints of $\text{PublicContinuous}(G(\square))$ do not change with time, M remains feasible at all points in time $\square \in [0, 1]$ (the duration of the algorithm), and its revenue $R_M(\square)$ is continuous at all $\square \in [0, 1]$ since each $\square_i(\square)$ changes continuously. Since $R(\square) = \max_M R_M(\square)$, we have:

Lemma 3.5. In Algorithm 1, for any interval $(\square_{-1}, \square_0)$ where $\text{supp}(G(\square))$ does not change, the function $\tilde{R}(\square)$ is convex. Further, the function $\tilde{R}(\square)$ is continuous for all $\square \in [0, 1]$, that is, the entire duration of the algorithm.

Revenue of Signals. We are now ready to prove our main theorem quantifying the rate of decrease of $\tilde{R}(\square)$, and thus bounding the revenue of the signals.

Theorem 3.6. In any interval $(\square_{-1}, \square_0)$ where $\text{supp}(G(\square))$ does not change, the function $\tilde{R}(\square)$ is linear, and $\frac{d\tilde{R}(\square)}{d\square} = -R(\square(\square))$, where $R(\square(\square))$ is the optimal revenue of $\text{Public}(\square(\square))$. Furthermore, at the end of Algorithm 1, it holds that $R(\Theta_G^0) = R(G)$.

Proof. Fix some $\square \in (\square_{-1}, \square_0)$, and assume $\text{supp}(G(\square))$ does not change in $(\square_{-1}, \square_0)$ and thus $\square(\square) = \square$. Note that $\text{supp}(G(\square)) = \text{supp}(\square)$. First consider the case where $\square_{\min}(\square) \geq \square$. By Lemma 3.2, every revenue maximizing auction for \square must have $\square_i = \square$ for all $\square_i \in \text{supp}(\square)$, and thus $R(\square) = \square$. The same proof implies every revenue maximizing auction M for $\text{PublicContinuous}(G(\square))$ must have $\square_i = \square$ for all $\square_i \in \text{supp}(G(\square))$. Therefore, we have

$$\frac{d\tilde{R}(\square)}{d\square} = \sum_{i=1}^{\square} \frac{d\square_i(\square)}{d\square} \cdot \square = \square \cdot \sum_{i \in \text{supp}(\square)} -\square_i = -\square = -R(\square) = -R(\square(\square)) \cdot \square \in \text{supp}(\square)$$

Next consider the case when $\square_{\min}(\square) < \square$ and let $\text{supp}(\square) = \{\square_1, \dots, \square_n\}$. Since \square is an equal revenue distribution, by Lemma 3.2, we have

$$\square_1 \cdot \bar{\square}_1(\square_1) = \dots = \square_n \cdot \bar{\square}_n(\square_n) = \square_1 = \square_{\min}(\square) = R(\square). \quad (5)$$

Let M^\square be an optimal solution to $\text{PublicContinuous}(G(\square))$. Note that M^\square is not necessarily unique, and further, can change as \square changes. Since $\text{supp}(\square) = \text{supp}(G(\square))$, by Theorem 3.4, the revenue $\tilde{R}(\square)$ achieved by M^\square in $\text{PublicContinuous}(G(\square))$ is the revenue of a distribution of posted prices $\{\square_1, \dots, \square_n\} \subseteq \text{supp}(\square)$ with weights $\square_1, \dots, \square_n \in (0, 1]$, where $\sum_{i=1}^n \square_i = 1$.

⁵This is slightly misleading: As we show later, the function $\tilde{R}(\square)$ is actually linear as long as $\text{supp}(G(\square))$ does not change. The overall function over $\square \in [0, 1]$ turns out to be piece-wise linear and concave.

At any time $\square \in (\square_{-1}, \square)$, the function $\tilde{R}_{M^*}(\square)$ – the revenue of M^* over $G(\square)$ – is linearly decreasing. We now calculate the rate of decrease at time \square . By Theorem 3.4,

$$\begin{aligned} \frac{d\tilde{R}_{M^*}(\square)}{d\square} &= -\frac{d}{d\square} \sum_{\square=1}^{\square'} \square \cdot \square \cdot \sum_{\square \geq \square} \square \cdot \square(\square) \\ &= -\sum_{\square=1}^{\square'} \square \cdot \square(\square) \quad (\text{by definition that } \frac{d\square}{d\square} = -\frac{d}{d\square}) \\ &= -\sum_{\square=1}^{\square'} \square \cdot R(\square) = -R(\square), \end{aligned}$$

where the second last equality is by Eq. (5), and the last because $\sum_{\square=1}^{\square'} \square = 1$.

Therefore, $\tilde{R}_{M^*}(\square)$ is a linear function for $\square \in (\square_{-1}, \square)$. By Lemma 3.5, $\tilde{R}(\square)$ is a maximum of linear revenue functions, one for each feasible solution, and hence convex. Since the solution M^* achieves this maximum at time $\square = \square$, the function $\tilde{R}_{M^*}(\cdot)$ is a subgradient of $\tilde{R}(\cdot)$ at time \square with slope $-R(\square)$. But the same holds for all $\square \in (\square_{-1}, \square)$ (although the corresponding optimal solution M^* may be different). Therefore, at every point $\square \in (\square_{-1}, \square)$, the function $\tilde{R}(\square)$ has a subgradient of slope $-R(\square)$. Since $\tilde{R}(\square)$ is convex in this interval, it must be a linear function, and hence differentiable. It also implies $\frac{d\tilde{R}(\square)}{d\square} = -R(\square) = -R(\square(\square))$, completing the proof of this case.

Note that $\square(\square)$ and thus $-R(\square(\square))$ changes only when some valuation in the residual prior $G(\square)$

is exhausted. This happens finitely many times throughout the process. Hence $\frac{d\tilde{R}(\square)}{d\square} = -R(\square(\square))$ is a piecewise constant function with finitely many discontinuities, and is thus Riemann integrable. Recall $\tilde{R}(0) = \tilde{R}(G)$ and $\tilde{R}(1) = 0$. Also recall every signal \square in Algorithm 1 is associated with weight $\square = \square - \square_{-1}$. Therefore we have

$$\begin{aligned} \tilde{R}(G) &= \tilde{R}(0) - \tilde{R}(1) = \int_{\square=0}^1 -\frac{d\tilde{R}(\square)}{d\square} d\square = \int_{\square=0}^1 R(\square(\square)) d\square \\ &= \sum_{\square=1}^{\square'} \int_{\square=\square_{-1}}^{\square} \square \cdot \square(\square) d\square = \sum_{\square=1}^{\square'} \square \cdot R(\square) = R(\Theta_G^{\square}). \end{aligned} \tag{6}$$

We then observe that it is always feasible for the seller to ignore the signals: Consider any arbitrary optimal auction $\hat{M} = (\{\square\}, \{\square\})$ for $\text{Public}(G)$. If the seller implements \hat{M} as the auction for each signal \square created by Algorithm 1, by Bayes plausibility, the resulting revenue is given by

$$R(G) = \sum_{\square=1}^{\square'} \sum_{\square=1}^{\square} \square \cdot \square(\square) \cdot \square = \sum_{\square=1}^{\square} \sum_{\square=1}^{\square} \square \cdot \square \cdot \square(\square) \cdot \square = \sum_{\square=1}^{\square} \square(\square) \cdot \square =$$

Since the above is the total revenue raised by some auction \hat{M} over all signals $\{\square\}$, the total revenue $R(\Theta_G^{\square})$ raised by implementing the optimal auction for each signal \square is at least as much. Therefore, we have $R(G) \leq R(\Theta_G^{\square})$. Combining this with Eq. (6), and observing that $\text{PublicContin-$

uous}(G) relaxes $\text{Public}(G)$, we have:

$$R(G) \geq \tilde{R}(G) = R(\Theta_G^{\square}) \geq R(G).$$

Hence, all inequalities must be equalities, which proves the theorem. \square

In Theorem 3.6 above, we have proved that the process in Algorithm 1 preserves the seller's expected revenue. By Lemma 3.3, Θ_G^{\square} also achieves efficiency, and thus maximizes social welfare.

Hence Θ_G^* must maximize the consumer surplus. This implies that the analog of Theorem 2.1 on buyer optimality holds for the public budget case:

Theorem 3.7 (Buyer optimality for public budgets). Suppose $\square = 1$ in $\text{Budgets}(G)$ for some prior G . Then there exists a signaling scheme Θ_G^* (given by Algorithm 1) that guarantees:

(1) $\text{SW}(\Theta_G^*) = \text{SW}^*(G)$; (2) $R(\Theta_G^*) = R(G)$; and (3) $\text{CS}(\Theta_G^*) = \text{CS}^*(G) = \text{SW}^*(G) - R(G)$.

Discussion. The nice aspect of our proof approach is twofold. First, we can derive the following corollary showing the existence of a common revenue-optimal auction throughout the process.

Corollary 3.1. Any optimal auction $M^*(\square')$ for $\text{PublicContinuous}(G(\square'))$ at time \square' in Algorithm 1 stays revenue optimal throughout the course of the algorithm (i.e., for all $\square > \square'$).

Secondly, our approach is extensible to more complex settings in the following sense: Theorem 3.4 yields a characterization of the revenue optimal auction in the specific case of public budgets. We use this to prove the first claim in Theorem 3.6, that the rate of decrease of revenue of the optimal auction is equal to the revenue of the signal constructed. The rest of the proof is generic in that it involves the convexity of $\tilde{R}(\square)$, which only relies on this function being the maximum of linear functions, one for each solution to the revenue maximizing LP. Our proof for the FedEx case simply reuses the generic portion, along with a specialized characterization of the optimal auction there.

4 MAIN RESULT: SIGNALING SCHEME WITH PRIVATE DEADLINES

We now focus on the case with private valuation-deadline pairs $(\square, \square) \in G$, where G_\square denotes the marginal distribution of \square given $\square = \square$. This is the so-called FedEx problem [19]. The non-trivial aspect now is the construction of the signals themselves. We first generalize the notion of equal revenue signals in Section 4.1, and outline the corresponding signaling algorithm (Algorithm 2). We then proceed to show that Algorithm 2 is buyer optimal, using the same plan of attack in Section 3.

We begin with some notation. For an arbitrary prior A with $\text{supp}(A) \subseteq \text{supp}(G) = \{\square_1, \dots, \square_n\} \times \{1, \dots, \square\}$, we denote by A_\square the conditional distribution of A given $\square = \square$. Further, let $\square_{A_\square}(\square)$ $\Pr_{(\square, \square) \in A}[\square = \square \mid \square = \square]$, and $\square_{A_\square}(\square) = \Pr_{(\square, \square) \in A}[\square \leq \square \mid \square = \square]$. Further, define $\text{values}(A) = \{\square > 0 \mid \square_{A_\square}(\square) > 0\}$ as the set of values with non-zero support in A , and similarly, $\text{values}(A_\square) = \{\square > 0 \mid \square_{A_\square}(\square) > 0\}$. Finally, let $\square_{\min}(A) = \min\{\square > 0 \mid \square_{A_\square}(\square) > 0\}$, and $\square_{\min}(A_\square) = \min\{\square > 0 \mid \square_{A_\square}(\square) > 0\}$.

4.1 Generalized Equal Revenue Signals

One natural approach to solving this problem is to apply the algorithm in [7] to the marginal G_\square induced by each deadline separately. However, as we show in the full paper [22], such an approach does not raise optimal consumer surplus. The main reason is that such a signaling scheme must reveal the deadline of the buyer, which provides the seller with too much information. We therefore need a different and novel signaling scheme that can “blur” the deadline information in addition to the value information. Our key idea is to define signals that continuously pull mass from all marginals G_\square at once, albeit in an equal revenue fashion.

Definition 4.1 (Lower Envelope). Given prior A with $\text{supp}(A) \subseteq \text{supp}(G)$, for all $\square \in [1, \square]$, let

$$\hat{\square}_\square = \max\{\square : \square_{A_\square}(\square) = 0 \mid \square \in [1, \square]\}$$

denote the largest \square such that no buyer with valuation at most \square and deadline at least \square exists in A . Note that $\hat{\square}_\square = 0$ if $\square_{A_\square}(\square) > 0$ for some $\square' \in [1, \square]$. Let $\hat{\square}_{\square+1} = \square$. The lower envelope of A is defined as

$$\text{LE}(A) = \{(\square, \square) \mid \square_{A_\square}(\square) > 0 \mid \square < \square \leq \hat{\square}_{\square+1}\}.$$

We say a value-deadline pair $(\underline{v}, \underline{d})$ is on the lower envelope of A if $(\underline{v}, \underline{d}) \in \text{LE}(A)$. Two pairs $(\underline{v}, \underline{d}), (\underline{v}', \underline{d}')$ in $\text{LE}(A)$ where $\underline{d} < \underline{d}'$ and $\underline{v}' \geq \underline{v}$ are consecutive points on $\text{LE}(A)$ if there is no $\underline{d}'' \in (\underline{d}, \underline{d}')$ and some \underline{v} such that $(\underline{v}, \underline{d}'')$ is on the lower envelope of A .

An immediate observation is that for any A , its lower envelope $\text{LE}(A)$ does not contain two different valuation-deadline pairs with the same valuation:

Observation 4.1. For any prior A , if $(\underline{v}, \underline{d}), (\underline{v}', \underline{d}')$ in $\text{LE}(A)$ for some \underline{v} then $\underline{d} = \underline{d}'$.

Definition 4.2 (Equal revenue Lower Envelope LE_A). For an arbitrary prior A with $\text{supp}(A) \subseteq \text{supp}(G)$, let $\text{LE}(A)$ be supported on $\{(\underline{v}_1, \underline{d}_1), (\underline{v}_2, \underline{d}_2), \dots, (\underline{v}_n, \underline{d}_n)\}$, where $0 < \underline{v}_{\min}(A) = \underline{v}_1 \leq \underline{v}_2 \leq \dots \leq \underline{v}_n \leq \underline{v}$.⁶ We define the Equal Revenue Lower Envelope signal LE_A for A to be the equal revenue distribution over $\{(\underline{v}_1, \underline{d}_1)\}$, i.e.,

$$\Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(A)}} [\underline{v} \geq \underline{v}_1] \cdot \underline{d}_1 = \Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(A)}} [\underline{v} \geq \underline{v}_2] \cdot \underline{d}_2 = \dots = \Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(A)}} [\underline{v} \geq \underline{v}_n] \cdot \underline{d}_n = \underline{d}_1 = \underline{d}_{\min}(A).$$

In other words, when disregarding the deadlines (and thus treating LE_A as a distribution of \underline{v}), every valuation with nonzero probability mass in its support is an optimal monopoly price. Analogous to Definition 2.2, this distribution is unique given $\text{LE}(A)$. We also have the following observation:

Observation 4.2. For an arbitrary prior A with $\text{supp}(A) \subseteq \text{supp}(G)$, for any $\underline{v}, \underline{d}$ such that $(\underline{v}, \underline{d}) \in \text{supp}(\text{LE}_A)$, it holds that $\text{LE}_A(\underline{v}, \underline{d}) = 0$ for all $\underline{d} < \underline{d}' \leq \underline{d}$.

Algorithm. For any time $\underline{t} \in [0, 1]$, we now let $F(\underline{t}) = [\underline{v}_1 \ \underline{v}_2 \ \dots]$ be an $n \times 1$ matrix function representing the residual prior, where $\underline{v}_i \in \underline{v}$ represents the remaining probability mass on type $(\underline{v}_i, \underline{d}_i)$ at time \underline{t} . Similar to the public budget case, let $G(\underline{t})$ denote the probability distribution obtained by placing the remaining probability mass $1 - \Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(A)}} [\underline{v} \geq \underline{v}_1]$ at $(0, 0)$. We omit considering $(0, 0)$ as part of the support of $G(\underline{t})$. Therefore we define $\text{supp}(G(\underline{t})) \subseteq \{(\underline{v}, \underline{d}) \mid \underline{v} > 0, \underline{d} > 0\}$. For each deadline $\underline{d} = \underline{d}_i$ we denote the marginal distribution of $G(\underline{t})$ as $G_{\underline{d}}(\underline{t})$. We therefore have: $\text{values}(G(\underline{t})) \subseteq \{\underline{v} > 0 \mid \Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(A)}} [\underline{v} > 0] > 0\}$, $\text{values}(G_{\underline{d}}(\underline{t})) \subseteq \{\underline{v} > 0 \mid \Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(A)}} [\underline{v} > 0] > 0\}$, and $\underline{v}_{\min}(G_{\underline{d}}(\underline{t})) \subseteq \min\{\underline{v} > 0 \mid \Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(A)}} [\underline{v} > 0] > 0\}$.

We now start with the prior distribution $G(0) = G$ and let $\underline{v}_i(0) = \Pr[G = (\underline{v}_i, \underline{d}_i)]$ for every i . Our algorithm continuously takes away probability mass from $F(\underline{t})$ and transfers it to the constructed signals, terminating when $F(\underline{t})$ becomes 0 at time $\underline{t} = \underline{t}_*$. At any time $\underline{t} \in [0, \underline{t}_*]$, denote $\underline{v}(\underline{t}) = \underline{v}_{\text{LE}(G(\underline{t}))}$ over $\text{LE}(G(\underline{t}))$ (Definition 4.2), $\underline{v}_i(\underline{t}) = \Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(G(\underline{t}))}} [(\underline{v}, \underline{d}) = (\underline{v}_i, \underline{d}_i)]$, and $S(\underline{t}) \subseteq [\underline{v}_i(\underline{t})]$. As $S(\underline{t})$ depends on $\text{values}(G(\underline{t}))$ but not $\text{values}(G(\underline{t}))$, it is fixed as long as $\text{values}(G(\underline{t}))$ does not change. Our algorithm continuously reduces $F(\underline{t})$ at rate $S(\underline{t})$ until $F(\underline{t})$ becomes 0:

$$\frac{dF}{d\underline{t}} = -S(\underline{t}). \quad (7)$$

Since $\Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(G(\underline{t}))}} [\underline{v} > \underline{v}_i(\underline{t})] = 1$, the rate of decrease of $\Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(G(\underline{t}))}} [\underline{v} > \underline{v}_i(\underline{t})]$ is 1. Since $\Pr_{\substack{(\underline{v}, \underline{d}) \in \text{LE}(G(0))}} [\underline{v} > \underline{v}_i(0)] = 1$, we have $\underline{v}_i(\underline{t}_*) = 1$.

Signals constructed. We say the type $(\underline{v}_i, \underline{d}_i)$ is exhausted at time \underline{t} if $\underline{v}_i(\underline{t}) = 0$ but $\underline{v}_i(\underline{t}') > 0$ for all $\underline{t}' < \underline{t}$. Therefore, $\text{values}(G(\underline{t}))$ changes only when some type is exhausted. For each maximal time interval $[\underline{t}_1, \underline{t}_2]$ in which $\text{values}(G(\underline{t}))$ remains fixed, the final scheme includes a corresponding signal \underline{v} with weight $(\underline{t}_2 - \underline{t}_1)$ so that $\underline{v}(\underline{t}) = \underline{v}$ for $\underline{t} \in [\underline{t}_1, \underline{t}_2]$. Since $\underline{v}(\underline{t})$ changes only if some element in $F(\underline{t})$ becomes zero, the number of signals constructed is finite.

The overall signaling scheme is described in Algorithm 2.

⁶Note that all \underline{v}_i 's are unique by Observation 4.1.

Algorithm 2 Continuous Algorithm for Deadlines Setting

Input: G
Output: $\Theta = \Theta_G^{\mathbb{R}}$

- 1: $F(0) \leftarrow [\Pr_{(\square, \square) \in G}[(\square, \square) = (\square_0, \square_0)]]$
- 2: Compute $F(\square), \square(\square), S(\square)$ for all $\square \in [0, \square]$
- 3: Compute $0 < \square_1 < \dots < \square_n = 1$ so that some type (\square, \square) is exhausted at each $\square = \square_i$; let $\square_0 = 0$: for $\square \in \{1, \dots, \square\}$ do
- 5: $\square_i \leftarrow \square_i - \square_{i-1}; \square_i^2 \leftarrow \square(\square_{i-1}); \Theta \leftarrow \Theta \cup \{\square_i^2, \square_i\}$
- 6: return Θ

Similar to Algorithm 1, the signals created by Algorithm 2 are Bayes plausible:

Observation 4.3. For any arbitrary G , let $\Theta_G^{\mathbb{R}} = \{(\square_i^2, \square_i)\}_{\square \in [0, \square]}$ be the set of signals output by Algorithm 2 taking G as input. Then we have $\bigcup_{\square=1}^{\square} \square \square_i^2 = G$.

4.2 Optimal Auction for Signals

In the following, we show the counterparts of Lemmas 3.2 and 3.3 in the deadlines context, and that Algorithm 2 guarantees efficiency. Let $\square \in \Theta_G^{\mathbb{R}}$ denote a signal created by Algorithm 2 for the prior G .

Lemma 4.4. There is an optimal auction for \square that posts a price of $\square_{\min}(\square)$. Further, for every $\square \in \text{values}(\square)$, we have $\square \cdot \Pr_{(\square, \square) \in G}[\square \geq \square] = \square_{\min}(\square)$.

The characterization of the optimal auction above implies the item always sells in $\Theta_G^{\mathbb{R}}$:

Lemma 4.5 (Efficiency of $\Theta_G^{\mathbb{R}}$). For each signal $\square \in \Theta_G^{\mathbb{R}}$, there exists a revenue optimal auction that always sells the item. As a consequence, $\text{SW}(\Theta_G^{\mathbb{R}}) = \text{SW}^{\mathbb{R}}(G)$.

4.3 Characterization of Optimal Auction for $G(\square)$

In the following, we analyze the revenue of the signals using the same technique in Section 3.3: We make the constraints in $\text{Deadlines}(G(\square))$ hold not just for values in $\text{values}(G(\square))$ but for all continuous values $\square > 0$. Fix some time \square and let $A = G(\square)$ so that $\text{supp}(A) \subseteq \text{supp}(G)$. We describe the linear program with extended domain and IC/IR constraints:

$$\begin{aligned}
 \text{DeadlinesContinuous}(A) \quad & \max_{\{\square(\cdot), \square(\cdot)\}, \{\square(\cdot), \square(\cdot)\}} \Pr_{(\square, \square) \in A} [\square = \square] \cdot \sum_{\square=1}^{\square} \square_{\square}(\square) \cdot \square(\square) \\
 \text{s.t.} \quad & \square \cdot \square(\square) - \square(\square) \geq \square \cdot \square(\square') - \square(\square'), \quad \square, \square' \in [0, \square], 1 \leq \square \leq \square, \\
 & \square \cdot \square(\square) - \square(\square) \geq \square \cdot \square_{i-1}(\square) - \square_{i-1}(\square), \quad \square, \square' \in [0, \square], 2 \leq \square \leq \square, \\
 & \square \cdot \square(\square) - \square(\square) \geq 0, \quad \square, \square' \in [0, \square], 1 \leq \square \leq \square, \\
 & 0 \leq \square(\square) \leq 1,
 \end{aligned}$$

We denote $R(A)$ the optimal revenue achievable by $\text{Deadlines}(A)$ and $\tilde{R}(A)$ the optimal revenue achievable by $\text{DeadlinesContinuous}(A)$. Clearly, $R(A) \geq \tilde{R}(A)$.⁷

We now present a characterization result for the optimal auction that is a discrete analog of the characterization for continuous priors in [15, 19]. We present a stand-alone and elementary proof

⁷It follows from [9, 25] that these two revenues are equal; however, we will not need this fact in our proof.

for the discrete setting in the full paper [22]. We note that unlike [15, 19], our proof uses convexity of the utility curve in the primal solution instead of invoking duality, and may be of independent interest.

Theorem 4.6. For any prior $A = G(\square)$ such that $\text{values}(A) = \{\square_1, \square_2, \dots, \square_n\}$, where $\square_1 < \square_2 < \dots < \square_n$, there exists weights $\alpha_1, \alpha_2, \dots, \alpha_n \in [0, 1]$ for all $\square \in [1, \square_n]$ such that the optimal revenue of $\text{DeadlinesContinuous}(A)$ is

$$\tilde{R}(A) = \Pr_{\substack{\square = 1 \\ \square \in \text{supp}(G(\square))}} \left[\square = \square \right] \cdot \sum_{\square=1}^{\square_n} \alpha_{\square} \cdot \square_A(\square) \cdot \square_{\square} \quad !$$

Furthermore, we have the following properties about the lower envelope $LE(A)$:

- If $(\square, \square) \in LE(A)$, then $\alpha_{\square} = \alpha_{\square}$ for all $\square \geq \square$
- If (\square, \square) and (\square, \square') are consecutive points on $LE(A)$ where $\square < \square$ and $\square' \geq \square$, then $\alpha_{\square} = 0$ for all $\square \in (\square, \square)$ and $\square \geq \square$
- $\sum_{\square=1}^{\square_n} \alpha_{\square} \cdot \square_{\square} = 1$, where α_{\square} equals 1 if $(\square, \square) \in LE(A)$ for some \square , and 0 otherwise.

4.4 Revenue Preservation in Algorithm 2

We now prove that Algorithm 2 preserves the expected seller revenue, following the same roadmap as in Section 3.4: We argue that the rate of decrease of revenue of $\text{DeadlinesContinuous}(G(\square))$ equals the optimal revenue of the signal $\square(\square)$.

Convexity of Revenue in $\text{DeadlinesContinuous}(G(\square))$. At any time \square , let $\tilde{R}(\square)$ denote the optimal revenue of $\text{DeadlinesContinuous}(G(\square))$; since $\text{DeadlinesContinuous}(G(0)) = \text{DeadlinesContinuous}(G)$ has more constraints than $\text{Deadlines}(G)$, we have $\tilde{R}(0) \leq R(G)$. Also, $\tilde{R}(1) = 0$.

Similar to Section 3.4, we consider any time interval (\square_{-1}, \square) in which $\text{supp}(G(\square))$, and hence $\square(\square)$, does not change for $\square \in (\square_{-1}, \square)$, and let the signal be \square and its corresponding probability matrix to be $S = [\square_{\square}]$ over $\square \in (\square_{-1}, \square)$. Then for any feasible solution $M \in \mathbb{B}^{\{\square_{-1}(\cdot)\}, \{\square_{\square}(\cdot)\}}$, its revenue in $\text{DeadlinesContinuous}(G(\square))$ is given by $\tilde{R}_M(\square) = \sum_{\square=1}^{\square_n} \square_{\square} \cdot \square_{\square}(\square) \cdot \square_{\square}(\square)$. Fixing this M , for each $\square \in (\square_{-1}, \square)$ we have

$$\frac{d\tilde{R}_M(\square)}{d\square} = \frac{d}{d\square} \sum_{\square=1}^{\square_n} \sum_{\square=1}^{\square_n} \square_{\square} \cdot \square_{\square}(\square) \cdot \square_{\square}(\square) = - \sum_{\square=1}^{\square_n} \square_{\square} \cdot \square_{\square}(\square) \quad !, \quad (8)$$

where the last equality uses Eq. (7) and the fact that $\square_{\square} = 0$ for any $\square \in \text{values}(G_{\square}(\square))$ for all $\square \in [1, \square]$. Therefore, for each M , its revenue $\tilde{R}_M(\square)$ decreases linearly with time. Further, since the constraints of $\text{DeadlinesContinuous}(G(\square))$ do not change with time, M remains feasible at all points in time $\square \in [0, 1]$ (the duration of the algorithm), and its revenue $\tilde{R}_M(\square)$ is continuous at all $\square \in [0, 1]$ since each $\square_{\square}(\square)$ changes continuously. Since $R(\square) = \max_M \tilde{R}_M(\square)$, we have:

Lemma 4.7. For any interval (\square_{-1}, \square) where $\text{supp}(G(\square))$ does not change, the function $\tilde{R}(\square)$ is convex. Further, the function $R(\square)$ is continuous for all $\square \in [0, 1]$, the entire duration of the algorithm.

Revenue of Signals. We now quantify the rate of decrease of $\tilde{R}(\square)$ in the deadlines setting.

Theorem 4.8. In any interval (\square_{-1}, \square) where $\text{supp}(G(\square))$ does not change, the function $\tilde{R}(\square)$ is linear, and $\frac{d\tilde{R}(\square)}{d\square} = -R(\square(\square))$, where $R(\square(\square))$ is the optimal revenue of $\text{Deadlines}(\square(\square))$. Furthermore, at the end of Algorithm 2, it holds that $R(\Theta_G^{\square}) = R(G)$.

Proof. Fix some $\square \in (\square_{-1}, \square)$, and assume $\text{supp}(G(\square))$ does not change in (\square_{-1}, \square) and thus $\square(\square) = \square$. Let $\text{values}(G(\square)) = \{\square_1, \square_2, \dots, \square_d\}$, where $\square_1 < \square_2 < \dots < \square_d$. Recall that \square is an equal revenue lower envelope distribution supported on $\text{LE}(G(\square))$ (see Definition 4.2), and thus $\text{values}(\square) \subseteq \text{values}(G(\square))$. By Lemma 4.4, for all $\square \in \text{values}(\square)$ we have:

$$\Pr_{(\square, \square) \in \square} [\square \geq \square] \cdot \square = \square_{\min}(\square) = R(\square). \quad (9)$$

Let M^* be the (not necessarily unique) revenue maximizing solution to $\text{DeadlinesContinuous}(G(\square))$. At any time $\square \in (\square_{-1}, \square)$, the function $\tilde{R}_{M^*}(\square)$ – the revenue of M^* over $G(\square)$ – is

linearly decreasing. We now calculate the rate of decrease at time \square . By Theorem 4.6, for each deadline $\square = \square$, the revenue $\tilde{R}(\square)$ achieved by M^* in $\text{DeadlinesContinuous}(G(\square))$ is a weighted combination of the revenue for posted prices $\{\square_1, \dots, \square_d\}$ with weights $\square_1^{\square}, \dots, \square_d^{\square} \in [0, 1]$, where

$$\sum_{\square=1}^d 1_{\square}^{\text{LE}(G(\square))} \cdot \square^{\square} = 1. \text{ Therefore,}$$

$$\begin{aligned} \frac{d\tilde{R}_{M^*}(\square)}{d\square} &= \frac{d}{d\square} \sum_{\square=1}^d \Pr_{(\square, \square) \in \square} [\square = \square] \cdot \square^{\square} \cdot \square_{\min}(\square) \\ &= \frac{d}{d\square} \sum_{\square=1}^d \sum_{\square=1}^d \square^{\square} \cdot \square^{\square} \cdot \Pr_{\square \geq \square} [\square = \square] = \sum_{\square=1}^d \sum_{\square=1}^d \square^{\square} \cdot \square^{\square} \cdot \Pr_{\square \geq \square} [\square = \square] \\ &= - \sum_{\square=1}^d \square^{\square} \cdot \square^{\square} \cdot \Pr_{\square \geq \square} [\square = \square]. \quad (\star) \quad \square=1 \quad \square=1 \quad \square \geq \square \end{aligned}$$

Consider some $\square \in \text{values}(\square)$. Suppose there exist some $\square < \square' < \square$ and $\square \leq \square'$ such that $(\square, \square), (\square, \square')$ are consecutive points in $\text{LE}(G(\square)) = \text{supp}(\square)$ are consecutive points in $\text{LE}(G(\square))$. By Theorem 4.6, this implies $\square_t = 0$ for all $\square \geq \square$. Also, for all $\square \geq \square$ we have $\square_t(\square) = 0$ for all $\square < \square$. Otherwise, there is no $\square > \square$ such that $\square \in \text{supp}(\square)$; in this case we have $\square_t(\square) = 0$ for all $\square \geq \square$ and all $\square \in [1, \square]$.

In both cases, the summation $\sum_{\square=1}^d \square^{\square} \cdot \Pr_{\square \geq \square} [\square = \square]$ evaluates to 0.

On the other hand, for each $\square \in \text{values}(\square)$, there is a unique \square such that $(\square, \square) \in \text{LE}(G(\square)) = \text{supp}(\square)$, which (by Theorem 4.6) implies $\square_t = \square$ for all $\square \in [\square, \square]$. Since for all $\square \geq \square$ it still holds that $\square_t(\square) = 0$ for all $\square < \square$, we have

$$\sum_{\square=1}^d \square^{\square} \cdot \Pr_{\square \geq \square} [\square = \square] = \sum_{\square=1}^d \square^{\square} \cdot \Pr_{\square \geq \square} [\square = \square] = \square \cdot \Pr_{(\square, \square) \in \square} [\square \geq \square].$$

Combining the above with Eq. (9), the expression (\star) evaluates to

$$- \sum_{\square \in \text{values}(\square)} \square \cdot \Pr_{(\square, \square) \in \square} [\square \geq \square] = - \sum_{\square \in \text{values}(\square)} \square \cdot R_{\square} = -R(\square),$$

where the last equality is by $\Pr_{(\square, \square) \in \square} [\square \geq \square] = 1_{\square}^{\text{LE}(G(\square))} \cdot \square^{\square} = 1$.

Therefore, $\tilde{R}_{M^*}(\square)$ is a linear function for $\square \in (\square_{-1}, \square)$. By Lemma 4.7 and the same arguments in the proof of Theorem 3.6, the above implies $R(\square)$ is linear and hence differentiable in the interval $\square \in (\square_{-1}, \square)$, and $\frac{dR(\square)}{d\square} = -R(\square) = -R(\square)$.

Note that $R(\square)$ and thus $-R(\square)$ changes only when some (\square, \square) -type in the residual prior $G(\square)$ is exhausted; this happens finitely many times throughout the process. Hence $\frac{dR(\square)}{d\square} = -R(\square)$ is a piecewise constant function with finitely many discontinuities, and is thus Riemann integrable.

Similar to that in the proof of Theorem 3.6, we have $\tilde{R}(0) = \tilde{R}(G)$, $\tilde{R}(1) = 0$, and every signal \square in Algorithm 2 is associated with weight $\square = \square - \square_{b-1}$. Therefore Eq. (6) holds exactly as is.

We observe that it is still feasible for the seller to ignore the signals by implementing some revenue optimal auction for $\text{Deadlines}(G)$ as the auction for each signal \square and achieve $R(G)$ as the total revenue. Thus, the theorem follows analogously to the public budget case. \square

Theorem 4.8 shows that Algorithm 2 preserves expected revenue. By Lemma 4.5, $\Theta_G^{\mathbb{B}}$ also achieves efficiency, and thus maximizes social welfare. Hence $\Theta_G^{\mathbb{B}}$ must maximize the expected consumer surplus, and the analog of Theorem 3.7 on buyer optimality holds for the deadlines case as well:

Theorem 4.9 (Buyer optimality for deadlines). In the private deadlines setting, there exists a signaling scheme $\Theta_G^{\mathbb{B}}$ for prior G that guarantees: (1) $\text{SW}(\Theta_G^{\mathbb{B}}) = \text{SW}^{\mathbb{B}}(G)$; (2) $R(\Theta_G^{\mathbb{B}}) = R(G)$; and (3) $\text{CS}(\Theta_G^{\mathbb{B}}) = \text{CS}^{\mathbb{B}}(G) = \text{SW}^{\mathbb{B}}(G) - R(G)$.

5 IMPOSSIBILITY OF OPTIMAL SIGNALING FOR PRIVATE BUDGETS

We now consider the setting with private budgets. Recall the program $\text{Budgets}(G)$ from Section 2, where the type space has valuation and budget, with the IR constraint being interim. We show that there are instances with just two budget types in which achieving full social welfare via signaling requires sacrificing almost all consumer surplus.

Theorem 5.1. For $\square = \square = 2$, for any given constant $\square > 0$, there exists a prior G in which any signaling scheme Θ that achieves efficiency (i.e., item always sells) has $\text{CS}(\Theta) \leq \square \cdot \text{CS}^{\mathbb{B}}(G)$, where $\text{CS}^{\mathbb{B}}(G) = \text{SW}^{\mathbb{B}}(G) - R(G)$ is the maximum achievable consumer surplus with respect to prior G .

Furthermore, a similar proof shows a lower bound of 2 on approximating the consumer surplus even when it is no longer required that the signaling scheme retains full social welfare.

Theorem 5.2. For $\square = \square = 2$, for any given constant $\square > 0$, there exists a prior G in which any signaling scheme Θ has $\text{CS}(\Theta) \leq (\frac{1}{2} + \square) \cdot \text{CS}^{\mathbb{B}}(G)$.

Both these theorems use the following family of instances, and are proved in the full paper [22].

Definition 5.1. For any $\square > 1$ and $\square < \frac{1}{\square}$, let the prior $G_{\square, \square}$ be supported on $\{(\square_1, \square_1), (\square_2, \square_2)\}$, where $\square_1 = 1$, $\square_1 = 1 - \square$ and $\square_2 = \square$, $\square_2 = \square$. Let $\square_1 = \Pr_{(\square, \square) \in G_{\square, \square}}[(\square, \square) = (\square_1, \square_1)] = 1 - \square$, and $\square_2 = \Pr_{(\square, \square) \in G_{\square, \square}}[(\square, \square) = (\square_2, \square_2)] = \square$.

6 CONCLUSION

Observe that our positive results hold for two budgeted settings where the optimal auctions with interim and ex-post IR constraints coincide, while our negative result holds for the most general budgeted setting where imposing ex-post IR constraints does reduce optimal revenue. In effect, our work points to a separation between auctions with ex-post IR constraints, where optimal signaling is possible (public budget or deadlines), and interim IR constraints, where it is not possible (private budget setting). The main open question is whether this separation can be formalized. We conjecture that there is indeed an optimal signaling scheme for the general private budget setting, when the mechanism is required to be ex-post IR instead of interim IR.

Several other open questions arise from our work. For instance, for private budgets with interim IR, is there an inefficient signaling scheme that extracts a constant factor of the optimal consumer surplus, thereby providing a positive counterpart to Theorem 5.2? Finally, can our results be generalized to larger type spaces, for instance, spaces with three dimensions such as value, deadline, and amount required?

Acknowledgment: This work is supported by NSF grant CCF-2113798.

REFERENCES

- [1] [n.d.]. DoubleClick Ad Exchange. <https://static.googleusercontent.com/media/www.google.com/en//adexchange/AdExchangeOverview.pdf>.
- [2] [n.d.]. Microsoft Ad Exchange. <https://ads.microsoft.com/>.
- [3] [n.d.]. Pubmatic Ad Exchange. <https://pubmatic.com/>.
- [4] [n.d.]. Yahoo Ad Exchange. <https://policies.yahoo.com/us/en/yahoo/terms/yahooadexchange/index.htm>.
- [5] Reza Alijani, Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang. 2020. The Limits of an Information Intermediary in Auction Design. *CoRR* abs/2009.11841 (2020).
- [6] Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi. 2021. Regret-Minimizing Bayesian Persuasion. In Proceedings of the 22nd ACM Conference on Economics and Computation. 128.
- [7] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2015. The limits of price discrimination. *American Economic Review* 105, 3 (2015), 921–57.
- [8] Dirk Bergemann and Stephen Morris. 2019. Information design: A unified perspective. *Journal of Economic Literature* 57, 1 (2019), 44–95.
- [9] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. 2021. A Duality-Based Unified Approach to Bayesian Mechanism Design. *SIAM J. Comput.* 50, 3 (2021).
- [10] Yang Cai, Federico Echenique, Hu Fu, Katrina Ligett, Adam Wierman, and Juba Ziani. 2020. Third-Party Data Providers Ruin Simple Mechanisms. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 1 (2020), 1–31.
- [11] Archishman Chakraborty and Rick Harbaugh. 2014. Persuasive puffery. *Marketing Science* 33, 3 (2014), 382–400.
- [12] Shuchi Chawla, David L. Malec, and Azarakhsh Malekian. 2011. Bayesian mechanism design for budget-constrained agents. In Proceedings of the 12th ACM Conference on Electronic Commerce. 253–262.
- [13] Yeon-Koo Che and Ian L. Gale. 2000. The Optimal Mechanism for Selling to a Budget-Constrained Buyer. *Journal of Economic Theory* 92, 2 (2000), 198–233.
- [14] Rachel Cummings, Nikhil R Devanur, Zhiyi Huang, and Xiangning Wang. 2020. Algorithmic price discrimination. In Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms. 2432–2451.
- [15] Nikhil R. Devanur and S. Matthew Weinberg. 2017. The Optimal Mechanism for Selling to a Budget Constrained Buyer: The General Case. In Proceedings of the 2017 ACM Conference on Economics and Computation. 39–40.
- [16] Shaddin Dughmi. 2017. Algorithmic information structure design: A survey. *ACM SIGecom Exchanges* 15, 2 (2017), 2–24.
- [17] Shaddin Dughmi, David Kempe, and Ruixin Qiang. 2016. Persuasion with limited communication. In Proceedings of the 2016 ACM Conference on Economics and Computation. 663–680.
- [18] Shaddin Dughmi and Haifeng Xu. 2021. Algorithmic Bayesian Persuasion. *SIAM J. Comput.* 50, 3 (2021).
- [19] 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. 21–22.
- [20] Nima Haghpanah and R. Siegel. 2020. Pareto Improving Segmentation of Multi-product Markets.
- [21] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.
- [22] Shao-Heng Ko and Kamesh Munagala. 2022. Optimal Price Discrimination for Randomized Mechanisms. <https://arxiv.org/abs/2205.04687>.
- [23] Jean-Jacques Laffont and Jacques Robert. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2 (1996), 181–186.
- [24] Roger B. Myerson. 1981. Optimal Auction Design. *Mathematics of Operational Research* 6, 1 (1981), 58–73.
- [25] Aviad Rubinstein and S. Matthew Weinberg. 2018. Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity. *ACM Trans. Economics and Comput.* 6, 3-4 (2018), 19:1–19:25.
- [26] Raghuvansh R. Saxena, Ariel Schwartzman, and S. Matthew Weinberg. 2018. The menu complexity of “one-and-a-half-dimensional” mechanism design. In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. 2026–2035.
- [27] Weiran Shen, Pingzhong Tang, and Yulong Zeng. 2018. A closed-form characterization of buyer signaling schemes in monopoly pricing. In Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems. 1531–1539.
- [28] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. 2015. Exploring Information Asymmetry in Two-Stage Security Games. In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence. 1057–1063.